

Pseudo- C_ℓ approach to kinematic Sunyaev-Zel'dovich stacking

Lea Harscouet,^{1,*} Kevin Wolz^{1,†}, Amy Wayland¹, David Alonso¹ and Boryana Hadzhiyska^{2,3}

¹*Department of Physics, University of Oxford,*

Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom

²*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, United Kingdom*

³*Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom*



(Received 13 January 2026; accepted 2 June 2026; published 26 June 2026)

We present a harmonic-space estimator for the cross-correlation between the kinematic Sunyaev-Zel'dovich (kSZ) effect and the reconstructed galaxy momentum field that offers several practical advantages over the traditional stacking approach. The estimator is easy to deploy using relatively modest computational resources and recovers all information available in the galaxy-kSZ cross-correlation. In particular, by using well-understood power spectrum estimation techniques, its statistical uncertainties, including potential correlated uncertainties with other large-scale structure observables, can be easily and accurately estimated. Moreover, standard kSZ stacking measurements can be reconstructed exactly from the estimator at a lower computational cost, employing harmonic-space, catalog-level techniques to recover all small-scale information.

DOI: [10.1103/ctv8-qn22](https://doi.org/10.1103/ctv8-qn22)

I. INTRODUCTION

As the size and sensitivity of new large-scale structure (LSS) and cosmic microwave background (CMB) datasets have grown over the last decade, our measurements have become sensitive to complex astrophysical effects that can no longer be modeled via simplified parametrizations [1–5]. Our poor understanding of these processes, particularly on small scales [e.g. the impact of active galactic nuclei (AGN) feedback on the suppression of small-scale matter fluctuations [6] or the detailed clustering of galaxies [7–9]], requires us to marginalize over astrophysical models with an increasing number of degrees of freedom, leading to a severe degradation in the cosmological constraints that could be recovered otherwise. This, as well as the need to obtain data-driven priors over these unknown astrophysical processes, has pushed the community toward the analysis of canonical cosmological probes (e.g. the clustering of galaxies and weak gravitational lensing measurements) in combination with other LSS tracers with a complementary sensitivity to the main astrophysical processes that dominate our uncertainties [3,10–18].

The kinematic Sunyaev-Zel'dovich (kSZ) effect is a perfect example of this paradigm [19,20]. The inverse-Compton scattering of CMB photons by free electrons in the intergalactic medium (IGM) causes an additional contribution to the CMB anisotropies. While the thermal motion of these electrons modifies the CMB spectrum, in the so-called “thermal” Sunyaev-Zel'dovich effect (tSZ), the kSZ effect is sourced by the bulk peculiar motion of the structures in which these electrons reside, which leaves the CMB spectrum unchanged. Crucially, the kSZ effect is directly sensitive to the density of ionized gas in the IGM, as well as the peculiar velocity field. As a direct probe of gas, the kSZ effect can provide a direct measurement of the distribution of baryons, including the impact of AGN feedback, and thus serve as an external calibrator for baryonic effects [3,21–25]. This is typically achieved by measuring the amplitude of the kSZ effect around galaxies, in the so-called “kSZ stacking” approach [21,23,26–29]: CMB maps, filtered to suppress the large-scale contribution from the uncorrelated primary CMB fluctuations, are stacked at the positions of sources weighted by an estimate of the galaxy velocities. These kSZ stacking measurements have provided evidence for the strength of AGN feedback to be larger than previously expected or predicted by hydrodynamical simulations, leading to a stronger suppression in the small-scale matter power spectrum as measured in cosmic shear analyses [3,21,22,28,30].

To ensure that uncertainties in astrophysical processes are robustly propagated in cosmological inference and self-calibrated by including external measurements (e.g. kSZ in the context of baryonic effects), an ideal analysis would

*Contact author: lea.harscouet@physics.ox.ac.uk

†Contact author: kevin.wolz@physics.ox.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

combine these measurements with the canonical cosmological probes and analyze them simultaneously within a unified physics-based model (e.g. one linking the stacked kSZ signal and the power spectrum suppression), taking into account all observational effects and cross-probe correlations [31–33]. Such joint analyses also improve in robustness when all different probes are processed under a common framework, making it possible to apply self-consistent analysis choices (e.g. scale cuts for different tracers), enabling the use of a coherent error model for all measurements, including correlated uncertainties and simplifying the production of theoretical predictions for all probes under a common scheme. In this sense, incorporating the standard kSZ stacking measurements into cosmological analyses poses a number of practical challenges: accounting for interpolation and pixelization effects inherent to stacking, the impact on scale cuts of the filtering used, the development of a robust covariance estimate, masking and survey geometry effects, and the relatively high computational cost of real-space correlation estimators [34–36]. While all these challenges can be addressed (e.g. using simulation-based covariances, forward-modeling pixelization effects, avoiding masking effects via catalog-level cuts), doing so increases the technical complexity of the analysis pipelines and the level of expertise needed to reliably incorporate kSZ measurements in any given study.

The integration of kSZ measurements into the widely used framework of correlation functions is not out of reach, however. Many traditional, real-space kSZ measurements are in fact indirect probes of the galaxy-galaxy-temperature (ggT) bispectrum; Ref. [37] establishes a direct equivalence between ggT and—among other estimators—the pairwise velocity method of [38], which led to the first detection of the kSZ effect, and the stacking estimator of [26,27], described above. Other works have also measured the kSZ effect via the galaxy-temperature-temperature gTT bispectrum (recently in [39] for projected fields) or its compressed proxy, the $C_\ell^{T^2g}$ skew spectrum in [40]. Our endeavor is therefore part of a wider effort to turn long-established kSZ detection methods into reliable tools that can be used in concordance with current methods for large-scale structure studies.

In this paper we present an angular power spectrum approach to measure the kSZ signal around galaxies which, as we will show, is at least equivalent to the kSZ stacking estimator in terms of information content, while resolving the practical challenges listed above. This new estimator $C_\ell^{\pi T}$ probes the ggT bispectrum by cross-correlating the momentum field π_g , built from the galaxy density and reconstructed velocity fields, and the CMB temperature field T (an approach originally advocated in [41]). Using well-tested methods developed in the context of pseudo- C_ℓ estimation [42–44], the resulting power spectrum measurements can be obtained at a low computational cost, are robust against masking effects, and their statistical

uncertainties can be accurately and easily estimated. The estimator may be seamlessly combined with power spectrum measurements involving other LSS tracers using the same framework, easily enabling the multitracers analysis needed for a robust self-calibration of astrophysical effects. Recently, we learned of similar work presented in upcoming publications [45,46], where this approach is applied to new datasets. Other works have also recently begun to model and study the kSZ signal around galaxies in terms of the harmonic-space spectrum [47,48].

This paper is structured as follows. Section II introduces the kSZ-galaxy power spectrum estimator and describes its advantages and its relation with kSZ stacking. Section III validates our implementation of this framework, demonstrates its ability to recover the stacking measurements at the estimator level, and showcases its application to real data. We then conclude in Sec. IV.

II. METHODS

A. kSZ angular power spectra

1. Galaxy momentum maps

The kSZ contribution to the temperature fluctuations is

$$\Delta T_{\text{kSZ}}(\hat{\mathbf{n}}) = \int d\chi q_{\text{kSZ}}(\chi) [1 + \delta_e(\chi\hat{\mathbf{n}}, z)] \hat{\mathbf{n}} \cdot \mathbf{v}(\chi\hat{\mathbf{n}}, z), \quad (1)$$

$$q_{\text{kSZ}}(\chi) \equiv T_{\text{CMB}} \sigma_T \bar{n}_{e,c} (1+z)^2, \quad (2)$$

where $\bar{n}_{e,c}$ is the comoving mean number density of electrons, T_{CMB} is the CMB temperature, and σ_T is the Thomson scattering cross section. δ_e is the electron overdensity, and $\mathbf{v}$ is the peculiar velocity field, in units where the speed of light is $c = 1$. Redshifts z and comoving distances χ are implicitly related to one another on the light cone. The kSZ effect is therefore a line-of-sight projection of the radial “electron momentum density” $q_e \equiv (1 + \delta_e) \hat{\mathbf{n}} \cdot \mathbf{v}^1$ and is sensitive to the large-scale structure both through the density and velocity fields.

Given a sample of galaxies, we may attempt a tomographic measurement of the kSZ effect by first constructing a proxy for the corresponding projected “galaxy momentum density” using the galaxy information, before correlating it with ΔT_{kSZ} . Since galaxies themselves trace the large-scale structure, this may be achieved by simply adding the radial velocities of all galaxies located in a given sky patch. That is, for a given sky pixel p , we may construct the following quantity:

$$\pi_{g,p} \equiv \frac{1}{N_p} \sum_{i \in p} w_i \hat{\mathbf{n}}_i \cdot \mathbf{v}_i^g, \quad (3)$$

¹Note that, in this definition, the momentum density has units of velocity, being normalized by the mean number density.

where the sum is over all objects with coordinates $\hat{\mathbf{n}}_i$ lying inside pixel p , w_i are the associated galaxy weights, and $\mathbf{v}_i^g$ is an estimate of their velocity. $\bar{N}_p$ is an estimate of the expected mean number of galaxies that could have been observed in pixel p and is included in order to account for potential variations in survey completeness and pixel area. π_g thus defined corresponds to the quantity

$$\pi_g(\hat{\mathbf{n}}) \equiv \int dz p(z) [1 + \delta_g(\chi\hat{\mathbf{n}}, z)] \hat{\mathbf{n}} \cdot \mathbf{v}^g(\chi\hat{\mathbf{n}}, z), \quad (4)$$

where δ_g is the galaxy overdensity, and $p(z)$ is the redshift distribution of the sample.

Note that we purposely denote the galaxy velocities $\mathbf{v}^g$, distinct from the true peculiar velocity field $\mathbf{v}$. This is because it is rarely possible to estimate the true galaxy velocities and, instead, it is most common to employ reconstruction techniques to estimate $\mathbf{v}^g$ from δ_g . For our purposes here, we will simply assume that an estimate of $\mathbf{v}_i^g$ exists for each galaxy, and its origin will be otherwise irrelevant.

Once π_g has been constructed, our final aim is to estimate the cross-spectrum,

$$C_\ell^{\pi T} \equiv \langle \Delta T_{\text{kSZ}, \ell m} \pi_{g, \ell m}^* \rangle. \quad (5)$$

We can do this following the standard pseudo- C_ℓ estimator [42,43]. In this case, the cross-power spectrum between two fields a and b is calculated assuming a diagonal pixel-space inverse covariance matrix for each field, defined by their masks. In most cases we will aim to extract ΔT_{kSZ} from either a component-separated CMB + kSZ map or from a simple frequency map, under the assumption that, out of all Galactic and extragalactic components contributing to it, the only one that correlates with the velocity-weighted density field is the kSZ. Either way, this map will normally be accompanied by its own mask $w_{\text{kSZ}}(\hat{\mathbf{n}})$, typically constructed through a combination of point-source and diffuse Galactic masks in addition to any mask associated with the actual survey footprint of the data. In contrast, a typical choice for the mask for π_g is $\bar{n}_p \equiv \bar{N}_p \Omega_p^{-1}$, the expected 2D-projected number density of galaxies in each pixel (with the dimension of inverse solid angle), where Ω_p is the pixel area. For noise-dominated data, and assuming that the velocity reconstruction uncertainties are homogeneous across the survey footprint, this choice is close to an inverse-variance weighting of the field.² In this case, the masked galaxy momentum field is

²For simplicity, we assume that $\bar{n}_p$ is the mask of choice in this work. This would correspond to the inverse-variance map of an inhomogeneous Poisson-sampled galaxy field. If an estimate of the spatially varying statistical uncertainties in the reconstructed velocities is available, this could be incorporated into a more optimal mask.

$$\tilde{\pi}_{g,p} = \frac{1}{\Omega_p} \sum_{i \in p} w_i v_{r,i}^g, \quad (6)$$

where $v_{r,i}^g \equiv \hat{\mathbf{n}}_i \cdot \mathbf{v}_i^g$.

In most cases, $\bar{n}_p$ will be constructed from a continuous map m_p^g , tracing the survey's completeness across the sky. In this case, $\bar{n}_p$ can be calculated by requiring that its integral across the sky should be the total (weighted) number of galaxies,

$$\bar{n}_p = \alpha_N m_p^g, \quad \alpha_N \equiv \frac{\sum_{i \in \mathcal{D}} w_i}{\sum_{p'} m_{p'}^g \Omega_{p'}}, \quad (7)$$

where the sums in α_N run over all galaxies in the sample and all pixels in the sky. Another common alternative is to define the survey completeness in terms of a random catalog, with the completeness at a given point in the sky given by the local weighted number density of random objects. In this case $\bar{n}_p$ may be constructed as

$$\bar{n}_p = \frac{\alpha_r}{\Omega_p} \sum_{i \in p} w_i^r, \quad \alpha_r = \frac{\sum_{j \in \mathcal{D}} w_j}{\sum_{k \in \mathcal{R}} w_k^r}, \quad (8)$$

where $\sum_{i \in p}$ involves summing over all random objects in pixel p , with weights w_i^r , and the sums $\sum_{i \in \mathcal{D}}$ and $\sum_{i \in \mathcal{R}}$ run over all data and random objects, respectively.

With the masked fields and associated masks for both fields ΔT_{kSZ} and π_g constructed as we have just described, the pseudo- C_ℓ estimate of their cross-spectrum can be extracted efficiently following the steps outlined in e.g. [43]. Moreover, its covariance matrix may also be estimated accurately and easily using the so-called improved narrow-kernel approximation (iNKA) as described in [49,50], using the measured pseudospectra of the fields involved as input.

2. Avoiding pixels: Catalog-based C_ℓ 's

As described in [44,51], the galaxy density is ultimately determined by the positions of the galaxies themselves, and traditionally the only reason to bin them into some form of regular pixel grid is the requirement of most numerical spherical harmonic transform (SHT) routines to use such a grid to achieve both efficiency and accuracy. However, the recent development of SHT algorithms on irregularly shaped grids [52] has made it possible to estimate the power spectra of fields sampled at the positions of discrete sources without pixelization. Since much of the kSZ signal resides at small angular scales ($\ell \gtrsim 3000$, the larger scales often being dominated by primary CMB “noise”), we can take advantage of these methods to efficiently recover the kSZ-galaxy power spectrum at these scales.

For example, the observed number density of galaxies and its spherical harmonic coefficients are given by

$$n_g(\hat{\mathbf{n}}) = \sum_{i \in \mathcal{D}} w_i \delta^{\mathcal{D}}(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (9)$$

$$n_{g,\ell m} = \sum_{i \in \mathcal{D}} w_i Y_{\ell m}^*(\hat{\mathbf{n}}_i), \quad (10)$$

where $\delta^{\mathcal{D}}(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i)$ is the Dirac delta function on the sphere. As in the previous section, let us assume that a map of the expected number density of galaxies $\bar{n}(\hat{\mathbf{n}})$ exists. Then, following this formalism, we may express the masked galaxy momentum field (and its SHT) as

$$\tilde{\pi}_g(\hat{\mathbf{n}}) = \sum_{i \in \mathcal{D}} w_i v_{r,i}^g \delta^{\mathcal{D}}(\hat{\mathbf{n}}, \hat{\mathbf{n}}_i), \quad (11)$$

$$\tilde{\pi}_{g,\ell m} = \sum_{i \in \mathcal{D}} w_i v_{r,i}^g Y_{\ell m}^*(\hat{\mathbf{n}}_i), \quad (12)$$

assuming $\bar{n}$ to be its mask.

The only other ingredient needed to estimate the power spectra of π_g is the spherical coefficients of the mask $\bar{n}_{\ell m}$. Again, as in the previous section, $\bar{n}$ may be provided either in terms of a completeness map m^g or through a random catalog. In the first case, $\bar{n}$ is simply given by $\bar{n}_{\ell m} = \tilde{\alpha}_n m_{\ell m}^g$, where

$$\tilde{\alpha}_n \equiv \frac{\sum_{i \in \mathcal{D}} w_i}{\int d\hat{\mathbf{n}} m^g(\hat{\mathbf{n}})}. \quad (13)$$

In the case of a random catalog, we may again take advantage of irregular-grid SHTs to calculate $\bar{n}_{\ell m}$,

$$\bar{n}_{\ell m} = \alpha_r \sum_{i \in \mathcal{R}} w_i^r Y_{\ell m}^*(\hat{\mathbf{n}}_i), \quad (14)$$

avoiding any finite-pixel effects.

The power spectra of two discretely sampled fields $\{a, b\}$, with masks $\{v, w\}$, can then be estimated following the standard pseudo- C_ℓ estimator [43,44]

- (1) Let $\tilde{a} \equiv va$ and $\tilde{b} \equiv wb$ be the masked fields. First, estimate their pseudo- C_ℓ , defined as the naive C_ℓ estimator ignoring the impact of masking,

$$\tilde{C}_\ell^{ab} \equiv \frac{1}{2\ell + 1} \sum_m \tilde{a}_{\ell m} \tilde{b}_{\ell m}^*. \quad (15)$$

Since masking leads to mode coupling, $\tilde{C}_\ell^{ab}$ is also sometimes called the “mode-coupled” C_ℓ .

- (2) Calculate the pseudo- C_ℓ of the masks $\tilde{C}_\ell^{vw}$ and, from it, the mode-coupling matrix (MCM) $M_{\ell\ell'}^{vw}$, encoding the statistical coupling between angular multipoles induced by masking and relating the ensemble average of $\tilde{C}_\ell^{ab}$ with the true underlying power

spectrum of the fields,

$$\langle \tilde{C}_\ell^{ab} \rangle = \sum_{\ell'} M_{\ell\ell'}^{vw} C_\ell^{ab}. \quad (16)$$

- (3) Invert the MCM to obtain an unbiased estimate of C_ℓ^{ab} . This step is often preceded by binning the power spectrum in ℓ into so-called “band powers,” which avoids numerical inversion issues when the masks are complex or cover a small portion of the sky and shortens the final data vector if a high resolution in ℓ is not necessary.

In addition to these steps, one caveat must be borne in mind in the case of autocorrelations of catalog-based fields where, as noted in [44,53], the shot-noise contributions to $\tilde{C}_\ell^{ab}$ and $\tilde{C}_\ell^{vw}$ must be subtracted to avoid aliasing effects, simultaneously hardening the estimator against any white noise contribution. This subtlety is of less importance in the case of $C_\ell^{\pi T}$, since we will rarely want to calculate the autocorrelation of π_g .

Finally, since this will be instructive for our discussion in the next section, let us consider the particular case in which the temperature map entering $C_\ell^{\pi T}$ is masked with a relatively simple binary mask w_T that encompasses the footprint over which π_g has been measured (i.e. $w_T \bar{n} = \bar{n}$). As we will see in the next section, this is in fact a common practice in kSZ stacking studies. Approximating $C_\ell^{\pi T}$ to be a slowly varying function, and that $M_{\ell\ell'}^{w_T \bar{n}}$ is sharply peaked around $\ell \simeq \ell'$, we can approximate Eq. (16) as

$$\langle \tilde{C}_\ell^{\pi T} \rangle \simeq C_\ell^{\pi T} \sum_{\ell'} M_{\ell\ell'}^{w_T \bar{n}} = \langle w_T \bar{n} \rangle_{\text{sky}} C_\ell^{\pi T}, \quad (17)$$

where $\langle \cdot \rangle_{\text{sky}}$ denotes averaging over the sky, and in the second equality we used a well-known property of the MCM (see e.g. [42,50]). The prefactor in the final expression is then

$$\langle w_T \bar{n} \rangle_{\text{sky}} = \langle \bar{n} \rangle_{\text{sky}} = \int \frac{d\hat{\mathbf{n}}}{4\pi} \bar{n}(\hat{\mathbf{n}}) = \frac{\sum_{i \in \mathcal{D}} w_i}{4\pi}. \quad (18)$$

Thus, we may define the following approximate estimator of the kSZ-galaxy power spectrum:

$$C_\ell^{\pi T} \equiv \frac{\tilde{C}_\ell^{\pi T}}{\langle \bar{n} \rangle_{\text{sky}}}. \quad (19)$$

The factor $1/\langle \bar{n} \rangle_{\text{sky}}$ corrects the overall amplitude of $\tilde{C}_\ell^{\pi T}$ for the effects of masking, although no attempt is made to account for mode coupling in $C_\ell^{\pi T}$.

B. The connection with kSZ stacking

The more standard method to extract the kSZ-galaxy cross-correlation signal used in the literature [21,23,27–29]

is the so-called kSZ stacking approach, in which a suitably filtered CMB map is stacked around the positions of galaxies weighted by their reconstructed velocities. The aim of this section is to show that the resulting kSZ stacking estimator can be recovered from the kSZ-galaxy angular power spectrum described in the previous section *at the estimator level*: i.e. all of the information recovered by the kSZ stack is encoded in the power spectrum measurements and can be recovered from it if so desired.

Let us start by describing the stacking estimator. Specifically, we will follow the prescription of [28], employing an aperture-photometry filter to reduce the impact of the large-scale primary CMB noise. Consider a galaxy at position $\hat{\mathbf{n}}_i$, and let $T(\hat{\mathbf{n}})$ be a map of the CMB temperature fluctuations. We define the aperture-photometry-filtered CMB map on an angular scale θ_d around this galaxy as

$$T(\hat{\mathbf{n}}_i, \theta_d) = \int d\hat{\mathbf{n}} \tilde{T}(\hat{\mathbf{n}}) W_{\text{AP}}(\hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}, \theta_d), \quad (20)$$

where the compensated aperture-photometry (CAP) filter is

$$W_{\text{AP}}(\cos \theta, \theta_d) = \begin{cases} 1 & \theta \leq \theta_d \\ -1 & \theta_d < \theta \leq \sqrt{2}\theta_d \\ 0 & \theta > \sqrt{2}\theta_d \end{cases}, \quad (21)$$

and, in Eq. (20) we have added a tilde over T to denote that it is a masked CMB map. The CAP filter is designed to suppress the impact of large-scale ($\ell \lesssim 100$) CMB fluctuations in the measurements, by estimating the largest-scale mode of the field in the inner disk of radius θ_d as the mean in an outer ring of equal area. The kSZ stacking estimator is then defined as the velocity-weighted average of T_i over a sample of galaxies,

$$\Delta \hat{T}_{\text{kSZ}}(\theta_d) \equiv Q \sum_{i \in \mathcal{D}} w_i v_i T_i(\hat{\mathbf{n}}_i, \theta_d), \quad (22)$$

where Q is a normalization factor. A typical choice for Q is

$$Q = \frac{1}{r_v v_{\text{rms}}} \frac{1}{\sum_{i \in \mathcal{D}} w_i}, \quad (23)$$

where $v_{\text{rms}}^2 \equiv \sum_{i \in \mathcal{D}} w_i (v_{r,i}^g)^2 / \sum_{i \in \mathcal{D}} w_i$ is the variance of the reconstructed velocities. r_v is the correlation coefficient between the reconstructed and true velocities and is usually included in Q to account for the mismatch between both. It is common in kSZ stacking analyses to only include in the estimator galaxies that are located far enough from the edges of the mask defining the T field to fit the largest disk-ring pair defining the CAP filters used in the measurement. We will assume this to be the case in what follows.

To rewrite the kSZ stacking estimator in terms of the masked galaxy momentum field we introduced in Eq. (11),

we insert a Dirac delta function for every object entering the sum in Eq. (22) and integrate over the sphere,

$$\begin{aligned} \Delta \hat{T}_{\text{kSZ}}(\theta_d) &= Q \int d\hat{\mathbf{n}} \tilde{\pi}_g(\hat{\mathbf{n}}) T(\hat{\mathbf{n}}, \theta_d) \\ &= Q \int d\hat{\mathbf{n}}_1 d\hat{\mathbf{n}}_2 \tilde{\pi}_g(\hat{\mathbf{n}}_1) \tilde{T}(\hat{\mathbf{n}}_2) W_{\text{AP}}(\mu_{12}, \theta_d), \end{aligned} \quad (24)$$

where $\mu_{12} \equiv \hat{\mathbf{n}}_1 \cdot \hat{\mathbf{n}}_2$ and, in the second line, we have made use of Eq. (20). We can now expand $\tilde{\pi}_g$ and $\tilde{T}$ in terms of their spherical harmonic coefficients

$$\tilde{T}(\hat{\mathbf{n}}) = \sum_{\ell m} \tilde{T}_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \quad \tilde{\pi}_g(\hat{\mathbf{n}}) = \sum_{\ell m} \tilde{\pi}_{g,\ell m} Y_{\ell m}(\hat{\mathbf{n}}),$$

and do likewise for the aperture-photometry filter

$$\begin{aligned} W_{\text{AP}}(\mu_{12}, \theta_d) &= \sum_{\ell} \frac{2\ell+1}{4\pi} W_{\ell}(\theta_d) P_{\ell}(\mu_{12}), \\ &= \sum_{\ell m} W_{\ell}(\theta_d) Y_{\ell m}^*(\hat{\mathbf{n}}_1) Y_{\ell m}(\hat{\mathbf{n}}_2), \end{aligned} \quad (25)$$

where W_{ℓ} is the harmonic-space CAP filter, discussed at the end of this subsection, $P_{\ell}(\mu)$ are the Legendre polynomials, and, in the second line, we have used the addition theorem for spherical harmonics. Finally, using the orthogonality of the spherical harmonic coefficients, we find

$$\Delta \hat{T}_{\text{kSZ}}(\theta_d) = Q \sum_{\ell} (2\ell+1) W_{\ell}(\theta_d) \tilde{C}_{\ell}^{\pi T}, \quad (26)$$

where, as in Eq. (15), we have defined

$$\tilde{C}_{\ell}^{\pi T} \equiv \frac{1}{2\ell+1} \sum_m \tilde{\pi}_{g,\ell m} \tilde{T}_{\ell m}^*. \quad (27)$$

Finally, using the definition of Q [Eq. (23)] and $C_{\ell}^{\pi T}$ [Eq. (19)], we may rewrite our last result in a more appealing form,

$$\Delta \hat{T}_{\text{kSZ}}(\theta_d) = \frac{1}{r_v v_{\text{rms}}} \sum_{\ell} \frac{2\ell+1}{4\pi} W_{\ell}(\theta_d) C_{\ell}^{\pi T}. \quad (28)$$

The relation between the velocity-weighted kSZ stack and the kSZ-galaxy power spectrum is thus very similar to that existing between the latter and the real-space correlation function of two scalar fields, with Legendre polynomials replaced by the harmonic-space CAP window functions $W_{\ell}(\theta_d)$. The latter are defined as

$$W_{\ell}(\theta_d) = 2\pi \int_{-1}^1 d\mu P_{\ell}(\mu) W_{\text{AP}}(\mu, \theta_d) \quad (29)$$

$$\simeq 2\pi \int_0^{\infty} d\theta \theta J_0(\ell\theta) W_{\text{AP}}(\cos \theta, \theta_d), \quad (30)$$




FIG. 1. Harmonic transform of the compensated aperture-photometry filter for different apertures θ_d . The solid lines show the calculation using the flat-sky approximation [Eq. (32)], with the points showing the exact calculation [Eq. (31)]. The flat-sky estimate is accurate enough on all relevant scales and significantly faster to calculate. The CAP filters peak at a scale $\ell_{\text{CAP}} \sim 2.5/\theta_d$, with an amplitude proportional to θ_d^2 .

where, in the second line, we have made use of the flat-sky approximation (appropriate for the small scales, $\ell \gtrsim 500$, over which the kSZ effect is normally studied) and $J_\alpha(x)$ is the cylindrical Bessel function of order α . Given the simple form of the CAP filter [Eq. (21)], $W_\ell(\theta_d)$ can be calculated analytically,

$$W_\ell(\theta_d) = \frac{2\pi}{2\ell + 1} [2\Delta P_\ell(\mu_d) - \Delta P_\ell(\mu_r)] \quad (31)$$

$$\simeq \frac{2\pi}{\ell^2} \ell \theta_d [2J_1(\ell \theta_d) - \sqrt{2} J_1(\sqrt{2} \ell \theta_d)], \quad (32)$$

where we have defined

$$\Delta P_\ell(\mu) \equiv \frac{1}{2} (P_{\ell-1}(\mu) - P_{\ell+1}(\mu)), \quad (33)$$

$$\mu_d \equiv \cos \theta_d, \quad \mu_r \equiv 2\mu_d - 1, \quad (34)$$

and, in Eq. (32), we have again adopted the flat-sky approximation. As shown in Fig. 1, the latter is sufficiently accurate on all relevant scales. The CAP filter peaks at $\ell_{\text{CAP}} \sim 2.5/\theta_d$, suppressing all large-scale modes (i.e. those dominated by the CMB primary fluctuations) below $\sim \ell_{\text{CAP}}$, and oscillates above ℓ_{CAP} with an amplitude that decays like $\propto \ell^{-3/2}$.

C. Advantages of the C_ℓ approach

As shown in the previous section, the kSZ stacking estimator can be reconstructed from the kSZ-galaxy power spectrum at the estimator level. That is, all the information encoded in the stack can be exactly recovered from $C_\ell^{\pi T}$ by convolving it with the CAP filters following Eq. (28). Moreover, since the CAP filters are not a complete set of

orthogonal basis functions, Eq. (28) selects only some of the modes in $C_\ell^{\pi T}$ and thus necessarily incurs some information loss. We will see that this is likely a small fraction, depending on the angular scales included and the impact of the primary CMB anisotropies as a source of irreducible noise. Nevertheless, $C_\ell^{\pi T}$ encodes all information recoverable from the π_g and T maps at the level of two-point statistics, making it an optimal observable to use.

Furthermore, the power spectrum estimator is fast, with the computational cost of spherical harmonic transforms scaling like $\propto \ell_{\text{max}}^3$. In contrast, stacking is effectively a two-point correlation function, which, for a given maximum θ_d , scales like $\propto N_{\text{gal}} N_{\text{pix}}$, where N_{gal} and N_{pix} are the number of galaxies and the number of sky pixels in a square that fits the largest ring for each galaxy, respectively. As we report in Sec. III B, the power spectrum approach is significantly faster (by a factor of ~ 30 for $\sim 3 \times 10^6$ objects) than the stacking approach employing similar computational resources.

Further advantages are offered by the extensive infrastructure developed for power spectrum estimation using the pseudo- C_ℓ approach. First, the estimator is able to account for all masking effects exactly and therefore avoids the need to discard galaxies close to the mask edges (as commonly done in stacking studies), potentially increasing the sensitivity of the measurements. Additionally, highly accurate and efficient methods have been developed and tested for the calculation of power spectrum covariances [49,50], allowing us to compute the statistical uncertainties of $C_\ell^{\pi T}$ in the approximation of Gaussian fields, without the need to rely on potentially inaccurate bootstrap/jackknife estimates or large numbers of costly mock realizations. This also enables us to easily calculate the correlated uncertainties of $C_\ell^{\pi T}$ with any other power spectrum of cosmological interest. This would, for example, enable a joint analysis of kSZ and 3×2 -point data from overlapping experiments, ensuring that all correlated uncertainties are accurately taken into account. Moreover, by exploiting statistical isotropy, the power spectra measured on different angular multipoles are largely uncorrelated, and therefore the covariance matrix of $C_\ell^{\pi T}$ is close to diagonal and numerically stable. This is in contrast to the stacking estimator where, due to the “cumulative” nature of the CAP filters, measurements on different angular scales θ_d are highly correlated. This also allows for a quick assessment of the detection significance or the goodness of fit of power spectrum measurements (i.e. performing “ χ^2 by eye”).

Finally, it is often the case that theoretical models for cosmological observables are naturally expressed in terms of power spectra. For instance, the galaxy-kSZ correlation (either in terms of stacking or $C_\ell^{\pi T}$) is largely determined by the cross-power spectrum between the galaxy and electron overdensities $P_{ge}(k)$ [47]. Thus, the use of $C_\ell^{\pi T}$ may also be

convenient when comparing against theoretical predictions, removing the need for Hankel transforms [e.g. Eq. (28)] connecting harmonic- and real-space quantities or enabling a quick assessment of the relevance of different contributions to the measurements [e.g. 1-halo and 2-halo terms in a halo model calculation of $P_{ge}(k)$].

III. RESULTS

A. Implementation and validation

We have implemented the general methodology described in the previous section to calculate the power spectrum between galaxy momentum fields and maps in the public pseudo- C_ℓ code NAMASTER³ [43]. We validate this implementation here using a toy model.

We generate Gaussian realizations of a “velocity” field v_r from a known power spectrum

$$C_\ell^{vv} = \frac{1}{(\ell + 10)^\alpha}, \quad (35)$$

with $\alpha = 4$. We then generate an uncorrelated galaxy overdensity map by first constructing a Gaussian field $g(\hat{\mathbf{n}})$ from a power spectrum following the same functional form with $\alpha = 3$,⁴ and define the overdensity δ_g from it via a log-normal transformation

$$\delta_g(\hat{\mathbf{n}}) = \exp \{g(\hat{\mathbf{n}}) - \sigma_g^2/2\} - 1, \quad (36)$$

where σ_g is the standard deviation of the map g . This ensures that the field is positive definite ($\delta_g \geq -1$) everywhere. From these maps we obtain the ground truth velocity momentum map via $\pi_g(\hat{\mathbf{n}}) = v_r^g(\hat{\mathbf{n}})[1 + \delta_g(\hat{\mathbf{n}})]$. To create a catalog equivalent of π_g , we construct a galaxy probability map given by $p(\hat{\mathbf{n}}) \propto (1 + \delta_g(\hat{\mathbf{n}}))m^g(\hat{\mathbf{n}})$, where m^g is the survey completeness map. For this, we use the selection function of the *Quia* survey [54], shown in Fig. 2. We then generate a Poisson realization of $p(\hat{\mathbf{n}})$ with a desired number density $\bar{n} \equiv \bar{N}/4\pi$. We explore catalogs with source numbers $\bar{N} \in \{10^3, 10^4, 10^5\}$. We assign to each source i the value $v_r(\hat{\mathbf{n}}_i)$ of the velocity field at the exact source position, avoiding any intermediate pixelization effects by direct transformation from the spherical harmonic coefficients $v_{r,\ell m}$ using an irregular-grid inverse SHT [52].

Having constructed the base momentum map $\pi_{g,p}$, representing the kSZ map, and the galaxy momentum




FIG. 2. Completeness mask used for the mock survey described in Sec. III A.

catalog $\pi_{g,i}$, we calculate their cross-correlation using the methods described in the previous section. Specifically, we construct a standard NAMASTER NmtField object, representing a continuous sky map with a given mask, from $\pi_{g,p}$ masked by the *Quia* survey selection map. In contrast, the galaxy momentum field is represented by an NmtFieldCatalogMomentum object, which uses direct SHTs to obtain the harmonic coefficients of $\pi_{g,i}$ [Eq. (12)], and the field’s mask is constructed from the completeness map as described in Eq. (13). We also validated our implementation of the NmtFieldCatalogMomentum field using a random catalog sampled from m^g , finding results similar to those presented here. The power spectrum between these two fields is then calculated using the standard pseudo- C_ℓ steps (see e.g. [43,44] for details). To check for estimator bias, we compare against the true input spectrum, obtained from the simulated full-sky momentum map⁵ and convolving with the band-power window functions associated with the mode-coupling matrix.

We generate 1000 of these simulations using maps defined by the HEALPix pixelization scheme [55] with resolution $N_{\text{side}} = 128$. The results of this validation exercise, in terms of the mode-coupled pseudo- C_ℓ , are shown in Fig. 3 for the three different number densities considered here.⁶ The mode-coupled estimator is unbiased in all cases, with Poisson noise becoming visible by eye at $\ell > 100$ in the low-density $\bar{N} = 10^3$ case. Figure 4 shows the mode-decoupled cross-power spectrum of the same set of simulations, compared against the

³<https://github.com/LSSTDESC/NaMaster>. An example notebook showcasing the kSZ implementation presented here can be found in https://github.com/LSSTDESC/NaMaster/blob/master/doc/kSZ_tutorial.ipynb.

⁴This is motivated by the fact that, at linear order, the continuity equation makes the velocity power spectrum scale as $P(k)/k$.

⁵To avoid pixelization effects in this comparison, we use this simulation-based estimate of the true spectrum, averaged over all simulations, on large scales $\ell < 150$, while, on small scales $150 \leq \ell < 2N_{\text{side}}$, we employ the analytical result derived in the Appendix for the linearized overdensity field $\delta_g \approx e^{-\sigma_g^2/2}(1 + g) - 1$.

⁶Note that, since the mask is the mean source number density, the amplitude of the pseudo- C_ℓ ’s (before accounting for mode coupling due to the mask) scales with the average number density, and hence the three cases explored here (and their theoretical expectations) differ by factors of 10.




FIG. 3. Coupled cross-spectrum between a momentum field computed from a simulated source catalog and its smooth map-based counterpart at $N_{\text{side}} = 128$, for catalog sizes of 10^5 , 10^4 , 10^3 . We compare against the theoretical input computed from the ground truth coupled with the catalog-based mode-coupling matrix. Upper: mean over 1000 simulations with bars showing the expected standard error on the mean. Lower: relative bias, in units of the error on the mean.




FIG. 4. Mode-decoupled catalog $\times$ map cross-spectrum for a simulated source catalog and its smooth map-based counterpart at $N_{\text{side}} = 128$, for catalog sizes of 10^5 , 10^4 , 10^3 . Upper: mean over 1000 simulations with bars showing the expected standard error on the mean. Lower: relative bias, in units of the error on the mean.

band-power-convolved ground truth spectrum. The decoupled estimator remains unbiased for all cases.

Although not tested here, in practice, one might also wish to bin a catalog-momentum field with inputs $(\hat{\mathbf{n}}_i, w_i, v_{g,i}^g)$ into a pixelized momentum map [see Eq. (3)] and use it as a map-based field, with the mask constructed from a completeness map as described in Sec. II A 1. In this case, one must remain mindful of pixelization, aliasing, and other finite-resolution effects (see [44] for further details).

B. Comparison against stacking

Having validated our implementation of the kSZ-galaxy power spectrum estimator, we now apply it to a realistic simulation with three aims:

- (i) to verify that the stacking estimator can indeed be recovered from the kSZ-galaxy pseudo- C_{ℓ} following Eq. (28);
- (ii) to quantify the information content in the power spectrum as opposed to the kSZ stack in scenarios with varying levels of realism;
- (iii) to benchmark the computational efficiency of both estimators (C_{ℓ} and stacking).

For this exercise, we use a mock galaxy catalog constructed from the AbacusSummit simulations [56,57]. Specifically, the catalog was designed to mimic the clustering properties of the photometric luminous red galaxy sample (LRG) targeted by the Dark Energy Spectroscopic Instrument (DESI) Legacy Survey. This was constructed using the halo occupation distribution parameters presented in [58]. We select all galaxies in the sample, covering the redshift range $0.2 \lesssim z \lesssim 1.1$ and peaking at $z \sim 0.5$ (see Fig. 5). We also construct a kSZ map by first building a three-dimensional map of the electron overdensity, constructed by scaling the matter




FIG. 5. Redshift distribution of the galaxy samples used here. The simulated photometric LRG galaxy sample of [58], used in Sec. III B, is shown in purple. The spectroscopic DESI Y1 LRG sample [61], used in Sec. III C, is shown in green.



FIG. 6. Simulation footprints: the blue region corresponds to the footprint of the galaxy catalog and is fully enclosed in the footprint of the temperature map (in yellow).

overdensity by a baryonic transfer function (see [59]). The resulting electron density field is scaled by the local value of the radial peculiar velocity and then integrated along the line of sight to construct a map of the kSZ temperature fluctuation [Eq. (1)]. As shown in Fig. 6, the mock LRG catalog was constrained to a sky footprint mimicking the area covered by the sixth data release of the Atacama Cosmology Telescope Collaboration (ACT DR6 [60]), with the kSZ map built over a larger region fully covering this footprint. The map was constructed using equirectangular pixels following the plate carrée pixelization as implemented in `PIXELL`,⁷ with pixels of size $0.5'$. For simplicity we weight all galaxies by their true peculiar velocities, both when stacking and when building the galaxy momentum field in the power spectrum estimator, enabling us to set $r_v = 1$ in Eq. (23).

Figure 7 shows the mode-decoupled kSZ-galaxy power spectrum computed from the simulation. We bin the multipole range $\ell \in [2; 12, 281]$ into 20 band powers of equal width. The figure includes the statistical uncertainties for this measurement, estimated using the `iNKA` of [50]. Note that within this approximation, the covariance only depends on two-point correlators of the fields involved, which can be estimated from the data, including the impact of shot noise. Since the simulated kSZ map contains only the pure kSZ signal, the spectrum is measured at very high precision on all scales, with a total signal to noise ratio (SNR) in the range $\ell \leq 12, 281$ of ~ 374 . We compute the SNR as [62]

$$\text{SNR} = \sqrt{\mathbf{d} \mathbf{C}^{-1} \mathbf{d}^T - N_{\text{d.o.f.}}}, \quad (37)$$

where $\mathbf{d}$ represents the data vector (here, either the $C_b^{\pi T}$ or the stacked profile), $\mathbf{C}$ its corresponding covariance matrix, and $N_{\text{d.o.f.}}$ is the number of degrees of freedom (here, the length of the data vector). For comparison, we show also the error bars calculated after including the contribution




FIG. 7. The decoupled pseudo- C_ℓ measured on the kSZ simulation. We measure the $C_b^{\pi T}$ on the pure kSZ realization at a SNR of 373.9 (turquoise error bars), which drops to 17.7 (purple error bars) when including the impact of other CMB contributions, foregrounds, and SO-like noise [63].

from primary CMB anisotropies, secondary anisotropies, extragalactic foregrounds at 90 GHz, and instrumental noise compatible with the sensitivity that the Simons Observatory (SO) will achieve after 5 years of observations [63]. These contributions reduce the SNR significantly to 17.7.

We then use the galaxy-kSZ pseudo- C_ℓ to reconstruct the CAP-stacked kSZ signal using Eq. (28) and compare the result with the direct stack measurements. For the latter, we compute the stacked kSZ signal using a curved-sky parallelized code. For each galaxy, the code extracts the “rectangle” defined by an interval of right ascension and declination that encompasses a spherical cap of radius $\sqrt{2}\theta_d^{\text{max}}$, enclosing the outer ring of the largest CAP filter used in the analysis. We consider 11 CAP filters with equispaced radii in the range $\theta_d \in [1', 10']$, but will discard $\theta_d > 6'$ to prevent edge effects from contaminating our measurements. As shown in Fig. 8, we recover the stack measurement with excellent precision (below 0.1%), showing that all information contained in kSZ stacking is encoded in the C_ℓ estimator. We find that this stacking algorithm matches the prediction from the C_ℓ better than other algorithms previously used in the literature.⁸ These differences are most prominent (up to $\sim 15\%$) on the smaller CAP apertures ($\theta_d \sim 1$ arc min). We ascribe this to interpolation and projection effects, related to the flat-sky projection and recentering operations commonly applied on the temperature map cutouts for each source [28]. This is in addition to the impact of the finite pixel size, which distorts the shape of the disk and ring components of the CAP filter. Although pixelization, projection, and interpolation effects

⁷<https://github.com/simonsobs/pixell>.

⁸E.g. <https://github.com/EmmanuelSchaan/ThumbStack>.




FIG. 8. Comparison of the stacking and C_ℓ estimators. Upper: stacked profiles obtained from the stacking estimator used in this work (orange) and from converting the kSZ-galaxy power spectrum $C_\ell^{\pi T}$ (whose decoupled version is shown in Fig. 7) to a stacked profile using Eq. (28) (turquoise). Note that all error bars have been multiplied by 20 for better readability. Lower: relative difference between the stacked profile and its reconstruction from the pseudo- C_ℓ . We highlight the $\pm 1\%$ and $\pm 0.1\%$ ranges in gray; the residuals lie largely within the $\pm 0.1\%$ range, showing that the stacks can be reconstructed from the pseudo- C_ℓ at the estimator level.

may be taken into account by including the associated kernels in the theoretical prediction used to interpret the measurements, the C_ℓ approach can avoid these issues altogether.

We can also estimate the covariance matrix of the stacks from the Gaussian covariance of the power spectrum, which we calculate using the iNKA of [50]. For this we can use Eq. (28), according to which the covariance of the stacks is related to the power spectrum covariance via

$$\mathbf{C}_{\theta\theta'} = N^2 \sum_{\ell, \ell'} (2\ell + 1)(2\ell' + 1) W_\ell(\theta) W_{\ell'}(\theta') \mathbf{C}_{\ell\ell'} \quad (38)$$

where $N = (4\pi v_{\text{rms}})^{-1}$. Figure 9 shows the covariances of the mode-decoupled power spectrum band powers $\mathbf{C}_{bb'}$ and the stacks $\mathbf{C}_{\theta\theta'}$, both derived from the pseudo- C_ℓ covariance $\mathbf{C}_{\ell\ell'}$. As expected, the $C_\ell^{\pi T}$ covariance is mostly diagonal and well behaved under inversion. The stack covariance, on the other hand, shows large correlation in the off-diagonal entries, as expected from the cumulative nature of the CAP filter. The error bars associated with this covariance are shown in Fig. 8 and can be compared there with the uncertainties calculated via bootstrap resampling (estimated from 1000 samples of the individual stacks with replacement). The Gaussian C_ℓ covariance, once converted to its stack form, recovers smaller errors than the bootstrap method at the smallest angular scales, while both estimates agree well at larger separations. A similar result is found for a covariance estimated by shuffling the velocities of different sources in Eq. (22), effectively canceling out the individual kSZ contributions. The differences found on small scales could be caused by inaccuracies in the bootstrap/reshuffling covariances or by non-Gaussian




FIG. 9. Comparison of the kSZ-galaxy power spectrum covariance $\mathbf{C}_{bb'}$ (left) with the stack covariance $\mathbf{C}_{\theta\theta'}$ (right), both derived from the full $C_\ell^{\pi T}$ covariance, $\mathbf{C}_{\ell\ell'}$. $\mathbf{C}_{bb'}$ is mostly diagonal, while $\mathbf{C}_{\theta\theta'}$ shows large cross-correlations in the off-diagonal entries. Note that the colors are shown using a logarithmic scale and that we show the absolute value of the $\mathbf{C}_{bb'}$ covariance in the left panel.

contributions to the covariance matrix not captured by the Gaussian power spectrum covariance. We defer a more detailed study of the kSZ power spectrum covariance to future work.

Using this new Gaussian estimate of the stack covariance, we can easily compare the information content in the C_ℓ and the stacking estimator. We compute the SNR for each of these using the two covariance matrices used in Fig. 7, which we label “kSZ-only” and “realistic.” To compute the SNR of the stacks, we first convert these two C_ℓ covariances to stack covariances using Eq. (38). For the realistic case, we find similar SNR values for both C_ℓ ’s and stacks of approximately $\text{SNR} \simeq 18$ ($\text{SNR} = 17.7$ for the C_ℓ ’s, and $\text{SNR} = 18.2$ for the C_ℓ -derived stack). The fact that both observables recover the same SNR is not entirely surprising, as we should expect the stacks to recover most of the small-scale information, whereas most of the large-scale signal, which the C_ℓ could recover more easily, is swamped by the large contribution from the primary CMB. Perhaps more surprisingly, we also find that the stacks and C_ℓ recover compatible SNRs in the kSZ-only scenario, where this large-scale information is more readily available, achieving $\text{SNR} \simeq 374$ in both cases. We therefore conclude that the level of information loss inherent in compressing the C_ℓ data into CAP-filtered stacks is largely negligible.

We note two subtleties when comparing the SNR of both of these estimators. First, binning the power spectrum into band powers may incur a small loss of information. Second, inverting the stack covariance can be numerically unstable, due to its large off-diagonal entries. The combination of these two effects can lead to a marginally larger SNR for the stacks than the C_ℓ ’s in some cases, even though the former can be constructed from the latter, as we have shown. We also estimated the SNR of the pseudo- C_ℓ before binning it into mode-decoupled band powers.⁹ In this case we find an SNR that is marginally larger than either the stacks or the C_ℓ band powers (e.g. $\text{SNR} = 378$ in the kSZ-only case), confirming that a similar amount of information is lost in binning or convolving with the CAP filters.

Finally, we compare the computational efficiency of both the stacking and C_ℓ estimators. We repeat measurements for different sample sizes N ($N = 10^3, 10^4, 10^5, 10^6$ and the full simulated catalog, i.e. $N = 3.6 \times 10^6$ objects). For the stacking estimator, we only compute the individual profiles at the nine angular separations shown in Fig. 8; for the C_ℓ estimator, we compute the pseudo- C_ℓ up to the same $\ell_{\text{max}} = 12,281$, but we do not include the decoupling and binning operations in the run-time performance. This is because these operations are discretionary, need only be performed once for a given set of sky masks, and can become a large run-time contribution as they depend on the




FIG. 10. Computational efficiency of both estimators as a function of the number of objects included in the catalog. The stacking estimator run-time performance (orange line) exhibits a linear trend, while the C_ℓ estimator performance (turquoise line) remains roughly constant with an increasing number of objects. Note that the performance of the stacking estimator depends on the number of angular separation values at which the profile is computed (here, the nine values of θ_d from Fig. 8), and that of the C_ℓ estimator depends on the ℓ_{max} up until which the power spectrum is computed (here, $\ell_{\text{max}} = 12,281$). For reference, we also show, as an orange band, the expected stacking performance for 5 angular separations (lower bound) and 15 angular separations (upper bound), and, as a turquoise band, the C_ℓ performance for $\ell_{\text{max}} = 10,000$ (lower bound) and $\ell_{\text{max}} = 15,000$ (upper bound). For simplicity, times are provided for runs on a single computer core, but both estimators can be fully parallelized.

binning scheme used.¹⁰ The result of this comparison is shown in Fig. 10. For both estimators, we only present run-time that pertains to the actual measurement (the loading of data products, common in both cases, is not included). The stacking estimator’s run-time evidently increases linearly with the number of objects in the sample and depends on the radial binning—as illustrated by the orange band in Fig. 10. The C_ℓ estimator, on the other hand, shows a roughly constant run-time performance, nearly independent of the size of the catalog and scaling mainly with ℓ_{max} (as shown by the turquoise band). For the radial binning used in this work, the C_ℓ estimator outperforms the stacking estimator at around a sample size of $N = 10^5$ objects, corresponding to a rather small galaxy catalog. For the size of the simulated catalog, the C_ℓ estimator is more than 30 times faster than the stacking estimator. It must be noted that, although the times reported in Fig. 10 correspond to runs on a single core, both tasks (pseudo- C_ℓ and stacking) are easily and fully parallelizable.

⁹We use singular-value decomposition in this case, since the unbinned inverse covariance is numerically unstable.

¹⁰For reference, for the power spectra shown in Fig. 7, the mode-coupling matrix takes about 1.7 core hours to compute up to $\ell_{\text{max}} = 12,281$ —or about 10 min on a 10-core machine.




FIG. 11. Dataset footprints: the dark blue footprint corresponds to the apodized ACT mask, which varies between 0 and 1. The yellow-green footprint is that of the DESI Y1 spectroscopic LRG catalog; variations in this mask track the number density of galaxies observed. We limit the galaxy catalog to those sources overlapping with the ACT mask.

C. Demonstration on real data

As a first demonstration of the kSZ power spectrum estimator on real data, we apply it to the cross-correlation between the component-separated “CMB + kSZ” ACT DR6 temperature map, presented in [64], and the DESI Y1 LRG spectroscopic galaxy sample [61]. LRGs are an ideal sample for kSZ measurements due to their higher clustering bias and optical depth and low satellite fraction. The latter makes them less sensitive to observational systematics such as off-centering [47,65]. Most importantly, the ACT $\times$ DESI LRG Y1 kSZ stacked profile has already been measured [65], allowing us to compare our results against those found in the literature. The redshift distribution of the DESI LRG Y1 sample is shown in Fig. 5 (green histogram, spanning $0.4 \lesssim z \lesssim 1.1$), and the sky masks of both datasets are shown in Fig. 11. Note that, for our fiducial measurements, we select only galaxies in the LRG sample that fall within the ACT DR6 footprint. We use reconstructed velocities computed using the PYRECON package¹¹ [66,67].

We use the same binning scheme employed in Sec. III B, both for the multipole band powers and angular separations. The kSZ-galaxy power spectrum measurement is shown in the left panel of Fig. 12. We detect a clear correlation between the kSZ map and the galaxy momentum field, with a total SNR of 8.7 [estimated via Eq. (37)], using angular scales $\ell \leq 12, 281$. The stacking measurements from the same data are shown in the right panel of the same figure (orange points with error bars), with the stack inferred from the C_ℓ measurement using Eq. (28)¹² shown in turquoise.

The two approaches agree well at large angular separation and, despite some differences, are still consistent

within the errors at small angular separations. As we described in Sec. III B, these differences are expected due to the impact of finite pixel effects, which are exacerbated in the case of real data by the presence of small-scale noise, leading to larger variations between nearby pixels. We verified that the small differences found in Sec. III B grow significantly when adding noise to the simulated kSZ map. To further verify that the differences observed are due to pixelization effects, we implement an alternative stacking method, free from such effects, which we label “harmonic-space stacking.” In this case, the CAP-filtered signal at the position of a given source is calculated by multiplying the harmonic transform of the CMB map by the SHT of the CAP filter window function and then evaluating this CAP-convolved map at the exact source position via an irregular-grid inverse SHT (as done in Sec. III A to avoid pixelization effects). The result of this method is shown as dark blue points in Fig. 12 and agree with the C_ℓ -derived signal at very high precision.

As in Sec. III B, the errors shown in the right panel of Fig. 12 were obtained via bootstrapping the individual profiles for both real- and harmonic-space stacks and by converting the Gaussian power spectrum covariance using Eq. (38) for the C_ℓ -derived stack. Interestingly, this time the analytical Gaussian errors are larger than their empirical (bootstrap) counterparts. This partly accounts for the slightly lower SNR at which the C_ℓ -derived stack (equivalently, the kSZ-galaxy power spectrum) is detected in this work, compared to the previous results of [65] who detected the LRG-kSZ correlation at the $\sim 10\sigma$ level. We illustrate this point using our hybrid harmonic stacking estimator, which matches both the C_ℓ approach in terms of stack amplitude and the real-space stacking in terms of covariance: the harmonic stack, together with its bootstrap covariance, is detected at SNR = 11.3, but its SNR drops back to the same SNR value as the power spectrum’s (SNR = 8.7) when using instead the Gaussian power spectrum covariance estimate. We leave a more thorough study of the kSZ power spectrum covariance, and the appropriateness of bootstrap uncertainties, for future work.

In the standard stacking measurements, a subset of the galaxies are further excluded to avoid those that lie too close to point sources or to the mask edges, in order to mitigate biases to the measurements caused by survey geometry effects. One of the advantages of the pseudo- C_ℓ estimator is its immunity to these effects. However, in general, limiting the sample to the overlap between the two fields being correlated may lead to an enhancement in the measured cross-correlation, since the map fluctuations far away from the overlapping region could act as a source of noise in harmonic space. It is therefore interesting to consider the impact in our measurements of removing the galaxies in the LRG sample that lie outside of the ACT footprint. Repeating our analysis using the full LRG sample shown in Fig. 11, we find that the total SNR drops to

¹¹<https://github.com/cosmodesi/pyrecon>.

¹²Note that, as we use the same reconstructed velocities in both measurements, we assume $r_v = 1$ in this conversion.




FIG. 12. Left: measurement of the kSZ-galaxy power spectrum from ACT DR6 and the DESI Y1 LRG sample (turquoise data points) with a theoretical prediction for the baryonic model of [68] with the parameters listed in Eq. (39) and a correlation coefficient of $r_v = 0.64$. Right: comparison of the corresponding stacked radial profiles and theory line with other stacking methods. In turquoise, we show the stack recovered from the coupled kSZ power spectrum; in dark blue, the “harmonic stack” estimate as described in the main text and, in orange, the standard stacking estimator. Bootstrap errors are shown for the standard and harmonic stacking estimators, and Gaussian errors [Eq. (38)] for the C_ℓ -derived profile.

$\text{SNR} = 8.5$. This small change is caused by a few-percent reduction in the amplitude of the measured spectrum, which in turn leads to smaller errors (as the Gaussian covariance estimate depends on the power spectra themselves). This change in amplitude could be a simple statistical fluctuation or a sign of potential large-scale systematics in the galaxy distribution, leading to a spatial modulation in the density or reconstructed galaxy velocities.

In Fig. 12, we plot theoretical predictions for the kSZ power spectrum (left panel) and stacks (right panel), based on the framework described in [47]. For simplicity, we neglect the subdominant contributions presented in [47], such as the cross term and the connected trispectrum. The prediction is based on a baryonic correction model [69], specifically that described in [68], in which the electron density profile is described by five parameters: $\log_{10} M_c$ sets the mass scale at which half of the gas content has been ejected by AGN outflows. β characterizes the mass dependence of the ejected fraction. η_b quantifies the radius out to which the ejected gas is expelled. The polytropic index Γ determines the density profile of the bound gas. Finally, A_* quantifies the fraction of the gas converted into stars (see [68] for further details). The amplitude of the theoretical prediction is also modified by the correlation coefficient r_v between the reconstructed and true velocities. We use $r_v = 0.64$ for the LRG sample used here [65]. The theoretical prediction shown in Fig. 12 corresponds to a model with parameters

$$\{\log_{10} M_c, \eta_b, \beta, A_*, \Gamma\} = \{14.9, 0.6, 0.6, 0.03, 1.16\}, \quad (39)$$

roughly compatible with the model found in [68] to describe the correlation between weak lensing, thermal SZ, and x-ray data. Although we have not attempted a robust analysis of our measurements within this model, the theory prediction shows a remarkable level of agreement with the kSZ measurements presented here. Moreover, Fig. 12 also allows us to assess the importance of secondary contributions to the kSZ-galaxy power spectrum, such as the 2-halo term (shown as a dotted line in the figure). Where this term becomes relevant ($1000 \lesssim \ell \lesssim 3500$), the error on the measurement is roughly comparable to the size of its contribution, suggesting that the 2-halo term may not be neglected in future modeling of the kSZ-galaxy power spectrum (as pointed out in [47]). We leave a more thorough theoretical exploration of the kSZ power spectrum measurements in combination with other gas probes for future work.

IV. CONCLUSIONS

This paper presents a new angular power spectrum approach to measuring the kSZ effect on projected fields; we show that this approach is mathematically equivalent to the more traditional kSZ “stacking estimator,” while offering a number of advantages. This work is part of a wider effort toward simplifying the inclusion of baryonic probes into cosmological analyses, from early work showing equivalence between real-space kSZ estimators and the ggT bispectrum [37] to current efforts to express the kSZ effect as a well-understood power spectrum measurement [40,45,46].

The estimator presented here consists in taking the cross-power spectrum of a galaxy momentum density field π_g ,

constructed from individual galaxy positions and velocities, with the CMB temperature map T . By reusing the existing infrastructure for power spectrum measurements, including catalog-based methods to bypass pixelization effects and aliasing, the same information can be extracted from the data at a much lower computational cost. The resulting measurements are robust to the effects of sky masks (e.g. mode coupling) and come with well-tested covariance estimates. As such, the kSZ effect can be seamlessly incorporated into large-scale structure analyses that use the same power spectrum framework, and the cross-covariances between probes can be obtained easily and reliably.

We validate our implementation of this estimator using a toy model in which we simulate fields and catalogs from known power spectra and compare their cross-correlation to the estimator outputs. Having validated the estimator, we apply it to both simulated and real data. We find that the C_ℓ approach, once converted to a stacked profile, reproduces the stacking approach, recovering all information present in the latter. Although we find small differences between both at small scales, we can trace these to the impact of finite pixel size, interpolation, and flat-sky projection, effects that the power spectrum estimator avoids altogether. We showed also that the power spectrum estimator offers a more scalable computational performance than stacking for large catalogs.

Finally, we applied this estimator to real data for the first time, specifically CMB maps from ACT DR6 and the DESI Y1 spectroscopic LRG sample. We show that the kSZ-galaxy power spectrum can be detected at high significance ($\text{SNR} \sim 9$) and that the result is compatible with stacking measurements at the estimator level once finite-resolution effects are taken into account. A tentative and qualitative comparison of this measurement against theoretical predictions using a simple parametrization of the gas distribution in halos (using the formalism described in [47]) shows a reasonable agreement with the physical models favored by the analysis of other independent gas probes (specifically tSZ and x-ray data in [15]).

In summary, the estimator presented here is ready for deployment on real data and allows for the inclusion of kSZ cross-correlation in multitracer cosmological analyses in a seamless and robust manner. To ensure a reliable accounting and propagation of all statistical uncertainties, a more

detailed study of the kSZ covariance (in particular contributions beyond those readily included in the standard pseudo- C_ℓ Gaussian covariance framework) is necessary, a problem that we have not addressed here. The work presented here could also be generalized to reproduce the information in other stacking estimators in a pseudo- C_ℓ framework. This would be particularly interesting in the context of “directional” stacking analyses carried out in e.g. the search for signals aligned with cosmic filaments [70], the moving-lens effect [71,72], and general dipoles correlated with the transverse velocity field [73]. This will be the subject of future work.

Note added. Recently, we learned of similar work being developed by Qu *et al.* [45] and Hadzhiyska *et al.* [46].

ACKNOWLEDGMENTS

We thank Emmanuel Schaan and the other authors of these works for useful discussions. We also thank Adrien La Posta and Andrina Nicola for useful comments. L. H. is supported by a Hintze studentship, which is funded through the Hintze Family Charitable Foundation. K. W. acknowledges support from the Science and Technology Facilities Council (STFC) under Grant No. ST/X006344/1. D. A. acknowledges support from STFC and the Beecroft Trust. A. W. is supported by a Science and Technology Facilities Council (STFC) studentship. We made extensive use of computational resources at the University of Oxford Department of Physics, funded by the John Fell Oxford University Press Research Fund.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available because they are owned by a third party and the terms of use prevent public distribution. The data are available from the authors upon reasonable request.

APPENDIX: THE POWER SPECTRUM OF $(1 + \delta)f$

Let $\pi(\hat{\mathbf{n}}) = (1 + \delta(\hat{\mathbf{n}}))f(\hat{\mathbf{n}})$, with $\delta(\hat{\mathbf{n}})$ the overdensity field, and $f(\hat{\mathbf{n}})$ another field (e.g. the radial velocity, in the case of a momentum density field). The harmonic coefficients of π are then

$$\pi_{\ell m} = f_{\ell m} + \sum_{\ell' m', \ell'' m''} \mathcal{G}_{mm'm''}^{\ell \ell' \ell''} \delta_{\ell' m'} f_{\ell'' m''}, \quad \text{where } \mathcal{G}_{mm'm''}^{\ell \ell' \ell''} \equiv \int d\hat{\mathbf{n}} Y_{\ell m}^*(\hat{\mathbf{n}}) Y_{\ell' m'}(\hat{\mathbf{n}}) Y_{\ell'' m''}(\hat{\mathbf{n}}). \quad (\text{A1})$$

The power spectrum of π is then

$$C_\ell^{\pi\pi} = C_\ell^{ff} + \sum_{\{\ell m\}_{1,2,3,4}} \mathcal{G}_{mm_1 m_2}^{\ell \ell_1 \ell_2} \mathcal{G}_{mm_3 m_4}^{\ell \ell_3 \ell_4} \langle \delta_{\ell_1 m_1} f_{\ell_2 m_2} \delta_{\ell_3 m_3}^* f_{\ell_4 m_4}^* \rangle \quad (\text{A2})$$

$$= C_\ell^{ff} + \sum_{\ell_1 \ell_2} \left[\sum_{m_1 m_2} (\mathcal{G}_{mm_1 m_2}^{\ell \ell_1 \ell_2})^2 \right] (C_{\ell_1}^{\delta\delta} C_{\ell_2}^{ff} + C_{\ell_1}^{\delta f} C_{\ell_2}^{\delta f}) \quad (\text{A3})$$

where, in the second line, we have assumed that all fields are Gaussian. Finally, expressing $\mathcal{G}_{mm'm''}^{\ell\ell'\ell''}$ in terms of Wigner 3j symbols, and using their orthogonality properties, we obtain the final result

$$C_\ell^{\pi\pi} = C_\ell^{ff} + \sum_{\ell_1\ell_2} \frac{(2\ell_1+1)(2\ell_2+1)}{4\pi} \begin{pmatrix} \ell & \ell_1 & \ell_2 \\ 0 & 0 & 0 \end{pmatrix}^2 (C_{\ell_1}^{\delta\delta} C_{\ell_2}^{ff} + C_{\ell_1}^{\delta f} C_{\ell_2}^{\delta f}). \quad (\text{A4})$$

-
- [1] A. Schneider, N. Stoira, A. Refregier, A. J. Weiss, M. Knabenhans, J. Stadel, and R. Teyssier, Baryonic effects for weak lensing. Part I. Power spectrum and covariance matrix, *J. Cosmol. Astropart. Phys.* **04** (2020) 019.
- [2] G. Aricò, R. E. Angulo, S. Contreras, L. Ondaro-Mallea, M. Pellejero-Ibañez, and M. Zennaro, The BACCO simulation project: A baryonification emulator with neural networks, *Mon. Not. R. Astron. Soc.* **506**, 4070 (2021).
- [3] L. Bigwood *et al.*, Weak lensing combined with the kinetic Sunyaev-Zel'dovich effect: A study of baryonic feedback, *Mon. Not. R. Astron. Soc.* **534**, 655 (2024).
- [4] A. Wayland, D. Alonso, and M. Zennaro, Calibrating baryonic effects in cosmic shear with external data in the LSST era, *Mon. Not. R. Astron. Soc.* **543**, 1518 (2025).
- [5] A. Schneider, M. Kovač, J. Bucko, A. Nicola, R. Reischke, S. K. Giri, R. Teyssier, T. Tröster, A. Refregier, M. Schaller, and J. Schaye, Baryonification: An alternative to hydrodynamical simulations for cosmological studies, [arXiv:2507.07892](https://arxiv.org/abs/2507.07892).
- [6] N. E. Chisari, A. J. Mead, S. Joudaki, P. G. Ferreira, A. Schneider, J. Mohr, T. Tröster, D. Alonso, I. G. McCarthy, S. Martin-Alvarez, J. Devriendt, A. Slyz, and M. P. van Daalen, Modelling baryonic feedback for survey cosmology, *Open J. Astrophys.* **2**, 4 (2019).
- [7] N. Kokron, J. DeRose, S.-F. Chen, M. White, and R. H. Wechsler, The cosmology dependence of galaxy clustering and lensing from a hybrid N-body-perturbation theory model, *Mon. Not. R. Astron. Soc.* **505**, 1422 (2021).
- [8] M. Zennaro, R. E. Angulo, M. Pellejero-Ibañez, J. Stücker, S. Contreras, and G. Aricò, The BACCO simulation project: biased tracers in real space, *Mon. Not. R. Astron. Soc.* **524**, 2407 (2023).
- [9] A. Nicola, B. Hadzhiyska, N. Findlay, C. García-García, D. Alonso, A. Slosar, Z. Guo, N. Kokron, R. Angulo, A. Aviles, J. Blazek, J. Dunkley, B. Jain, M. Pellejero, J. Sullivan, C. W. Walter, and M. Zennaro (LSST Dark Energy Science Collaboration), Galaxy bias in the era of LSST: Perturbative bias expansions, *J. Cosmol. Astropart. Phys.* **02** (2024) 015.
- [10] T. Tröster, A. J. Mead, C. Heymans, Z. Yan, D. Alonso, M. Asgari, M. Bilicki, A. Dvornik, H. Hildebrandt, B. Joachimi, A. Kannawadi, K. Kuijken, P. Schneider, H. Y. Shan, L. van Waerbeke, and A. H. Wright, Joint constraints on cosmology and the impact of baryon feedback: Combining KiDS-1000 lensing with the thermal Sunyaev-Zeldovich effect from Planck and ACT, *Astron. Astrophys.* **660**, A27 (2022).
- [11] X. Fang, E. Krause, T. Eifler, S. Ferraro, K. Benabed, R. S. Pranjali, E. Ayçoberry, Y. Dubois, and V. Miranda, Cosmology from weak lensing, galaxy clustering, CMB lensing, and tSZ-I. 10×2 pt modelling methodology, *Mon. Not. R. Astron. Soc.* **527**, 9581 (2024).
- [12] T. Ferreira, D. Alonso, C. García-García, and N. E. Chisari, X-ray-cosmic-shear cross-correlations: First detection and constraints on baryonic effects, *Phys. Rev. Lett.* **133**, 051001 (2024).
- [13] S. Pandey, J. Salcido, C.-H. To, J. C. Hill, D. Anbajagane, E. J. Baxter, and I. G. McCarthy, Accurate connected modeling of gas thermodynamics and matter distribution, *Phys. Rev. D* **111**, 043529 (2025).
- [14] T. Eifler, X. Fang, E. Krause, C. M. Hirata, J. Xu, K. Benabed, S. Ferraro, V. Miranda, S. Pranjali R., E. Ayçoberry, and Y. Dubois, Cosmology from weak lensing, galaxy clustering, CMB lensing and tSZ: II. Optimizing Roman survey design for CMB cross-correlation science, [arXiv:2411.04088](https://arxiv.org/abs/2411.04088).
- [15] A. La Posta, D. Alonso, N. E. Chisari, T. Ferreira, and C. García-García, Insights on gas thermodynamics from the combination of x-ray and thermal Sunyaev-Zel'dovich data cross correlated with cosmic shear, *Phys. Rev. D* **112**, 043525 (2025).
- [16] J. Sunseri, A. Amon, J. Dunkley, N. Battaglia, S. Ferraro, B. Hadzhiyska, B. Ried Guachalla, and E. Schaan, Disentangling the halo: Joint model for measurements of the kinetic Sunyaev-Zeldovich effect and galaxy-galaxy lensing, [arXiv:2505.20413](https://arxiv.org/abs/2505.20413).
- [17] J. Siegel, A. Amon, I. G. McCarthy, L. Bigwood, M. Yamamoto, E. Bulbul, J. E. Greene, J. McCullough, M. Schaller, and J. Schaye, Joint X-ray, kinetic Sunyaev-Zeldovich, and weak lensing measurements: Toward a consensus picture of efficient gas expulsion from groups and clusters, [arXiv:2509.10455](https://arxiv.org/abs/2509.10455).
- [18] J. Xu, T. Eifler, E. Krause, V. Miranda, J. Salcido, and I. McCarthy, Constraining baryonic feedback and cosmology from DES Y3 and Planck PR4 6×2 pt data. I. Λ CDM models, [arXiv:2510.25596](https://arxiv.org/abs/2510.25596).
- [19] M. Birkinshaw, The Sunyaev-Zel'dovich effect, *Phys. Rep.* **310**, 97 (1999).
- [20] T. Mroczkowski, D. Nagai, K. Basu, J. Chluba, J. Sayers, R. Adam, E. Churazov, A. Crites, L. Di Mascolo, D. Eckert, J. Macias-Perez, F. Mayet, L. Perotto, E. Pointecouteau, C. Romero, F. Ruppen, E. Scannapieco, and J. ZuHone, Astrophysics with the spatially and spectrally resolved Sunyaev-Zeldovich effects. A millimetre/submillimetre

- probe of the warm and hot universe, *Space Sci. Rev.* **215**, 17 (2019).
- [21] B. Hadzhiyska *et al.*, Evidence for large baryonic feedback at low and intermediate redshifts from kinematic Sunyaev-Zel'dovich observations with ACT and DESI photometric galaxies, *Phys. Rev. D* **112**, 083509 (2025).
- [22] I. G. McCarthy, A. Amon, J. Schaye, E. Schaan, R. E. Angulo, J. Salcido, M. Schaller, L. Bigwood, W. Elbers, R. Kugel, J. C. Helly, V. J. Forouhar Moreno, C. S. Frenk, R. J. McGibbon, L. Ondaro-Mallea, and M. P. van Daalen, FLAMINGO: combining kinetic SZ effect and galaxy-galaxy lensing measurements to gauge the impact of feedback on large-scale structure, *Mon. Not. R. Astron. Soc.* **540**, 143 (2025).
- [23] B. Ried Guachalla *et al.*, Backlighting extended gas halos around luminous red galaxies: Kinematic Sunyaev-Zel'dovich effect from DESI Y1 and ACT data, *Phys. Rev. D* **112**, 103512 (2025).
- [24] B. Hadzhiyska, N. Sailer, and S. Ferraro, Mapping the gas density with the kinematic Sunyaev-Zel'dovich and patchy screening effects: A self-consistent comparison, *Phys. Rev. D* **112**, 103532 (2025).
- [25] L. Bigwood, M. Yamamoto, J. Siegel, A. Amon, I. G. McCarthy, R. Dave, J. Salcido, M. Schaller, J. Schaye, and T. Yang, The kinetic Sunyaev Zeldovich effect as a benchmark for AGN feedback models in hydrodynamical simulations: Insights from DESI + ACT, [arXiv:2510.15822](#).
- [26] M. Li, R. E. Angulo, S. D. M. White, and J. Jasche, Matched filter optimization of kSZ measurements with a reconstructed cosmological flow field, *Mon. Not. R. Astron. Soc.* **443**, 2311 (2014).
- [27] E. Schaan *et al.* (ACTPol Collaboration), Evidence for the kinematic Sunyaev-Zel'dovich effect with the Atacama Cosmology Telescope and velocity reconstruction from the Baryon Oscillation Spectroscopic Survey, *Phys. Rev. D* **93**, 082002 (2016).
- [28] E. Schaan *et al.* (Atacama Cosmology Telescope Collaboration), Atacama Cosmology Telescope: Combined kinematic and thermal Sunyaev-Zel'dovich measurements from BOSS CMASS and LOWZ halos, *Phys. Rev. D* **103**, 063513 (2021).
- [29] M. Mallaby-Kay *et al.*, Kinematic Sunyaev-Zel'dovich effect with ACT, DES, and BOSS: A novel hybrid estimator, *Phys. Rev. D* **108**, 023516 (2023).
- [30] J. Siegel, L. Bigwood, A. Amon, J. McCullough, M. Yamamoto, I. G. McCarthy, M. Schaller, A. Schneider, and J. Schaye, The suppression of the matter power spectrum: Strong feedback from X-ray gas mass fractions, kSZ effect profiles, and galaxy-galaxy lensing, [arXiv:2512.02954](#).
- [31] M. Oguri and M. Takada, Combining cluster observables and stacked weak lensing to probe dark energy: Self-calibration of systematic uncertainties, *Phys. Rev. D* **83**, 023008 (2011).
- [32] M. Shirasaki, E. T. Lau, and D. Nagai, Probing cosmology and cluster astrophysics with multiwavelength surveys—I. Correlation statistics, *Mon. Not. R. Astron. Soc.* **491**, 235 (2020).
- [33] A. Nicola, J. Dunkley, and D. N. Spergel, Joint cosmology and mass calibration from thermal Sunyaev-Zel'dovich cluster counts and cosmic shear, *Phys. Rev. D* **102**, 083505 (2020).
- [34] M. Jarvis, G. Bernstein, and B. Jain, The skewness of the aperture mass statistic, *Mon. Not. R. Astron. Soc.* **352**, 338 (2004).
- [35] D. Alonso, CUTE solutions for two-point correlation functions from large cosmological datasets, [arXiv:1210.1833](#).
- [36] M. Sinha and L. H. Garrison, CORRFUNC—a suite of blazing fast correlation functions on the CPU, *Mon. Not. R. Astron. Soc.* **491**, 3022 (2020).
- [37] K. M. Smith, M. S. Madhavacheril, M. Münchmeyer, S. Ferraro, U. Giri, and M. C. Johnson, KSZ tomography and the bispectrum, [arXiv:1810.13423](#).
- [38] N. Hand *et al.*, Evidence of galaxy cluster motions with the kinematic Sunyaev-Zel'dovich effect, *Phys. Rev. Lett.* **109**, 041101 (2012).
- [39] R. Patki, N. Battaglia, and J. C. Hill, Novel bispectrum estimator of the kinematic Sunyaev-Zel'dovich effect using projected fields, *Phys. Rev. D* **112**, 063524 (2025).
- [40] J. C. Hill, S. Ferraro, N. Battaglia, J. Liu, and D. N. Spergel, Kinematic Sunyaev-Zel'dovich effect with projected fields: A novel probe of the baryon distribution with Planck, WMAP, and WISE data, *Phys. Rev. Lett.* **117**, 051301 (2016).
- [41] S. Ho, S. Dedeo, and D. Spergel, Finding the missing baryons using CMB as a backlight, [arXiv:0903.2845](#).
- [42] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, MASTER of the cosmic microwave background anisotropy power spectrum: A fast method for statistical analysis of large and complex cosmic microwave background data sets, *Astrophys. J.* **567**, 2 (2002).
- [43] D. Alonso, J. Sanchez, and A. Slosar (LSST Dark Energy Science Collaboration), A unified pseudo- C_ℓ framework, *Mon. Not. R. Astron. Soc.* **484**, 4127 (2019).
- [44] K. Wolz, D. Alonso, and A. Nicola, Catalog-based pseudo- C_ℓ s, *J. Cosmol. Astropart. Phys.* **01** (2025) 028.
- [45] F. J. Qu *et al.*, Precision Kinematic Sunyaev-Zel'dovich Measurements Across Halo Mass and Redshift with DESI DR2 and ACT DR6: Part I. Luminous Red Galaxies, [arXiv:2604.19744](#).
- [46] B. Hadzhiyska *et al.*, Precision kinematic Sunyaev-Zel'dovich measurements across halo mass and redshift with DESI DR2 and ACT DR6: Part II. Bright galaxy survey and emission-line galaxies, [arXiv:2604.19745](#).
- [47] A. Wayland, D. Alonso, and A. La Posta, Detailed theoretical modelling of the kinetic Sunyaev-Zel'dovich stacking power spectrum, [arXiv:2509.18732](#).
- [48] A. Laguë, M. S. Madhavacheril, J. Borrow, K. M. Smith, X. Chen, M. Schaller, and J. Schaye, Inferring the impacts of baryonic feedback from kinetic Sunyaev-Zeldovich cross-correlations, [arXiv:2511.20595](#).
- [49] C. García-García, D. Alonso, and E. Bellini, Disconnected pseudo- C_ℓ covariances for projected large-scale structure data, *J. Cosmol. Astropart. Phys.* **11** (2019) 043.
- [50] A. Nicola, C. García-García, D. Alonso, J. Dunkley, P. G. Ferreira, A. Slosar, and D. N. Spergel, Cosmic shear power spectra in practice, *J. Cosmol. Astropart. Phys.* **03** (2021) 067.

- [51] A. Baleato Lizancos and M. White, Harmonic analysis of discrete tracers of large-scale structure, *J. Cosmol. Astropart. Phys.* **05** (2024) 010.
- [52] M. Reinecke, S. Belkner, and J. Carron, Improved cosmic microwave background (de-)lensing using general spherical harmonic transforms, *Astron. Astrophys.* **678**, A165 (2023).
- [53] N. Tessore *et al.* (Euclid Collaboration), Euclid preparation: LIX. Angular power spectra from discrete observations, *Astron. Astrophys.* **694**, A141 (2025).
- [54] K. Storey-Fisher, D. W. Hogg, H.-W. Rix, A.-C. Eilers, G. Fabbian, M. R. Blanton, and D. Alonso, Quiaia, the gaia-unWISE quasar catalog: An all-sky spectroscopic quasar sample, *Astrophys. J.* **964**, 69 (2024).
- [55] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere, *Astrophys. J.* **622**, 759 (2005).
- [56] N. A. Maksimova, L. H. Garrison, D. J. Eisenstein, B. Hadzhiyska, S. Bose, and T. P. Satterthwaite, AbacusSummit: A massive set of high-accuracy, high-resolution N-body simulations, *Mon. Not. R. Astron. Soc.* **508**, 4017 (2021).
- [57] B. Hadzhiyska, L. H. Garrison, D. Eisenstein, and S. Bose, The halo light cone catalogues of \ABACUSSUMMIT, *Mon. Not. R. Astron. Soc.* **509**, 2194 (2021).
- [58] R. Zhou, J. A. Newman, Y.-Y. Mao, A. Meisner, J. Moustakas, A. D. Myers, A. Prakash, A. R. Zentner, D. Brooks, Y. Duan, M. Landriau, M. E. Levi, F. Prada, and G. Tarle, The clustering of DESI-like luminous red galaxies using photometric redshifts, *Mon. Not. R. Astron. Soc.* **501**, 3309 (2021).
- [59] R. H. Liu, B. Hadzhiyska, S. Ferraro, S. Bose, and C. Hernández-Aguayo, Fast baryonic field painting for Sunyaev-Zel'dovich analyses: Transfer function vs. hybrid effective field theory, [arXiv:2504.11794](https://arxiv.org/abs/2504.11794).
- [60] S. Naess *et al.* (The Atacama Cosmology Telescope Collaboration), The Atacama Cosmology Telescope: DR6 maps, *J. Cosmol. Astropart. Phys.* **11** (2025) 061.
- [61] R. Zhou *et al.*, Target selection and validation of DESI luminous red galaxies, *Astron. J.* **165**, 58 (2023).
- [62] L. Harscouet, D. Alonso, A. Nicola, and A. Slosar, Constraints from CMB lensing tomography with projected bispectra, [arXiv:2507.07968](https://arxiv.org/abs/2507.07968).
- [63] P. Ade *et al.*, The Simons Observatory: Science goals and forecasts, *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [64] W. Coulton *et al.*, Atacama Cosmology Telescope: High-resolution component-separated maps across one third of the sky, *Phys. Rev. D* **109**, 063530 (2024).
- [65] B. Ried Guachalla *et al.*, Backlighting extended gas halos around luminous red galaxies: Kinematic Sunyaev-Zel'dovich effect from DESI Y1 and ACT data, *Phys. Rev. D* **112**, 103512 (2025).
- [66] E. Paillas *et al.*, Optimal reconstruction of baryon acoustic oscillations for DESI 2024, *J. Cosmol. Astropart. Phys.* **01** (2025) 142.
- [67] X. Chen *et al.*, Extensive analysis of reconstruction algorithms for DESI 2024 baryon acoustic oscillations, [arXiv:2411.19738](https://arxiv.org/abs/2411.19738).
- [68] A. La Posta, D. Alonso, N. E. Chisari, T. Ferreira, and C. García-García, Insights on gas thermodynamics from the combination of x-ray and thermal Sunyaev-Zel'dovich data cross correlated with cosmic shear, *Phys. Rev. D* **112**, 043525 (2025).
- [69] A. Schneider and R. Teyssier, A new method to quantify the effects of baryons on the matter power spectrum, *J. Cosmol. Astropart. Phys.* **12** (2015) 049.
- [70] B. Hadzhiyska, S. Ferraro, and R. Zhou, Tracing cosmic gas in filaments and halos: Low-redshift insights from the kinematic Sunyaev-Zel'dovich effect, *Phys. Rev. D* **111**, 023534 (2025).
- [71] S. C. Hotinli, E. Pierpaoli, S. Ferraro, and K. Smith, Transverse velocities and matter gradient correlations: A new signal and a new challenge to moving-lens analyses, *Phys. Rev. D* **108**, 083508 (2023).
- [72] A. Beheshti, E. Schaan, and A. Kosowsky, Moving lens effect: Simulations, forecasts, and foreground mitigation, *Phys. Rev. D* **111**, 043510 (2025).
- [73] Y.-C. Cai, J. A. Peacock, A. de Graaff, and S. Alam, Detection of cosmological dipoles aligned with transverse peculiar velocities, *Mon. Not. R. Astron. Soc.* **541**, 2093 (2025).