

Missing baryons recovered: A measurement of the gas fraction in galaxies and groups with the kinematic Sunyaev-Zel'dovich effect and CMB lensing

Boryana Hadzhiyska^{1,2,3,4,*} Simone Ferraro^{1,2} Gerrit S. Farren^{1,2} Noah Sailer^{2,1} and Rongpu Zhou^{1,2}

¹*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

²*Berkeley Center for Cosmological Physics, Department of Physics, University of California, Berkeley, California 94720, USA*

³*Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA, United Kingdom*

⁴*Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, United Kingdom*



(Received 26 August 2025; accepted 7 November 2025; published 3 December 2025)

We present new constraints on the halo masses and matter density profiles of DESI galaxy groups by cross-correlating samples of luminous red galaxies (LRGs) and bright galaxy survey (BGS) galaxies with the publicly available CMB lensing convergence map from ACT DR6. This provides an independent, lensing-based calibration of halo masses, complementary to methods relying on clustering or dynamics. We derive constraints on the mean halo mass for three DESI-selected samples, finding $\log(M_{\text{halo}}/(M_{\odot}/h)) \approx 13.18, 13.03$ and 13.02 for the main LRG, extended LRG, and BGS samples, respectively. Using a halo model approach, we also compare the projected galaxy-matter density profiles with previously reported gas profiles inferred from measurements of the kinematic Sunyaev-Zel'dovich (kSZ) effect. This work addresses one of the key uncertainties in interpreting kSZ signals—the unknown host halo mass distribution—by providing an independent and consistent mass calibration. The agreement between the gas and total mass profiles at large aperture suggests that sufficiently far from the group center (2–3 virial radii), we recover all the baryons, offering a resolution to the missing baryon problem. We further study the cumulative gas fractions for all galaxies as well as for the most massive galaxy groups in the sample [$\log(M_{\text{halo}}/(M_{\odot}/h)) \approx 13.5$], finding values that are physically sensible and in agreement with previous findings using kSZ and x-ray data: compared to the TNG300 simulation, the observed gas fractions are systematically lower at fixed radius by $\gtrsim 4\sigma$, providing compelling, independent evidence for stronger baryonic feedback in the real Universe. These findings highlight the power of combining CMB lensing with galaxy surveys to probe the interplay between baryons and dark matter in group-sized halos.

DOI: [10.1103/mdhz-fgj8](https://doi.org/10.1103/mdhz-fgj8)

I. INTRODUCTION

The standard cosmological model, Λ cold dark matter (Λ CDM), has proven remarkably successful at describing the evolution of the Universe across a wide range of scales and epochs. Observations of the cosmic microwave background (CMB), baryon acoustic oscillations (BAO), and large-scale structure (LSS) have established a precise framework in which only a few parameters are needed to explain the expansion history and the growth of structure. Nevertheless, some of the most fundamental components of this model, dark energy and dark matter, remain enigmatic. With the advent of high-precision datasets from DESI,

Rubin, Euclid, and Simons Observatory, we are entering an era where new data will be sensitive to subtle departures from Λ CDM predictions. A critical challenge in leveraging this sensitivity is our incomplete understanding of baryonic physics, particularly baryonic feedback processes that redistribute gas within and beyond halos, altering the observable matter distribution on nonlinear scales [1,2]. This manifests itself in a large deficit of observed baryon abundance in late-time galaxies and groups, known as the “missing baryon” problem [3,4].

Baryonic feedback, especially from active galactic nuclei (AGN) and supernovae, expels gas from the centers of halos into their outskirts, suppressing the small-scale matter power spectrum and modifying weak lensing observables [5,6]. This has profound implications for cosmological inference: if not properly modeled, baryonic effects can bias constraints on parameters such as S_8 and the total neutrino mass [7–9]. Yet the modeling of feedback remains highly uncertain, in its amplitude, redshift and scale dependence, with state-of-the-art hydrodynamical simulations yielding

*Contact author: boryanah@alumni.princeton.edu

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.

divergent predictions despite matching many galaxy observables [10]. This degeneracy between baryon physics and cosmology motivates the development of new empirical probes that directly constrain the gas distribution across cosmic time.

One of the most powerful tools in this regard is the kinetic Sunyaev-Zel'dovich (kSZ) effect, a secondary CMB anisotropy generated when CMB photons are Doppler shifted by scattering off free electrons with bulk line-of-sight velocities [11]. Unlike the thermal SZ (tSZ) effect, which depends on electron temperature and is concentrated in hot cluster gas, the kSZ signal traces the momentum of all free electrons, making it sensitive to the total ionized baryon distribution. This is especially valuable in the circumgalactic medium (CGM) and the outskirts of halos ($\sim 1\text{--}4$ virial radii R_{vir}), where most of the “missing baryons” have now enabled significant detections of the kSZ signal on halo mass scales down to $10^{13}M_{\odot}$ [12–15]. These measurements suggest that the gas distribution is significantly more extended than the dark matter, consistent with strong feedback scenarios.

Nonetheless, the interpretation of the kSZ signal is limited by uncertainties in the halo mass of the galaxy sample. Since the kSZ amplitude depends linearly on the optical depth, which in turn is proportional to the host halo mass, the comparison between the matter and gas profiles significantly benefits from external mass information; uncertainty in the halo mass propagates directly into uncertainties on feedback strength [16,17]. Current approaches often rely on stellar mass estimates combined with empirical stellar population or halo occupation distribution (HOD) models calibrated on clustering. However, these indirect methods introduce substantial systematic uncertainties. As a result, attempts to use kSZ to test hydrodynamical simulations or inform cosmological models are substantially limited without more accurate halo mass measurements.

Gravitational lensing offers a natural solution. In particular, CMB lensing provides a clean, unbiased measurement of the projected mass distribution along the line of sight, without the complications of intrinsic alignments, shape noise, or shear calibration biases that affect galaxy-galaxy lensing (GGL). When cross-correlated with galaxy positions, CMB lensing can be used to measure the average halo mass and the matter density profile of a given sample [18,19]. This makes it an ideal counterpart to kSZ measurements: while kSZ traces the baryon density, lensing traces the total matter. Their combination enables a direct estimate of the gas-to-mass ratio (or gas fraction) as a function of scale, a critical test for both feedback models and for resolving the missing baryon problem.

An early comparison of the kSZ signal to small-scale GGL was performed in [17] and further studied in [20] within the halo model framework and in [21] with the use of state-of-the-art hydrodynamical simulations. A combined

analysis of cosmic shear and the kSZ effect was conducted in [22].

In this work, we perform a joint analysis of CMB lensing and kSZ measurements around DESI galaxies. Using high-resolution CMB lensing convergence maps from ACT and the Planck legacy release, we extract the small-scale projected mass profiles of photometrically selected DESI targets in multiple redshift and luminosity bins. We then compare these profiles to the corresponding kSZ measurements from Hadzhiyska *et al.* [13] to assess whether the recovered gas content matches theoretical expectations. This approach allows us to: (1) empirically constrain the average host halo mass of DESI galaxies in a manner that is independent of stellar mass modeling, (2) infer their baryon fraction as a function of scale, and (3) place novel constraints on the amplitude and scale dependence of baryonic feedback. Beyond characterizing the total mass of galaxy groups and comparing it with the gas mass, this analysis has significant implications for the use of DESI galaxies in upcoming lensing and clustering cosmology. The ability to infer their halo masses directly enables new inputs for HOD-based analyses and cross checks with GGL and 2-point correlation function analyses.

This paper is structured as follows. In Sec. II, we describe the galaxy and CMB datasets used in our analysis. Section III outlines the halo model framework used to model the lensing signal. Section IV presents our mass profile measurements, the comparison to kSZ measurements of the gas density profiles, and a comment on the observed stellar-to-halo mass relation. We summarize our findings and discuss implications for cosmology and galaxy formation in Sec. V.

II. DATA

In this section, we summarize the galaxy and CMB observational datasets used in this study.

A. Dark energy spectroscopic instrument

The Dark Energy Spectroscopic Instrument (DESI) is a robotic, fiber-fed, highly multiplexed spectrograph operating on the Mayall 4-meter telescope at Kitt Peak National Observatory [23]. It is capable of obtaining simultaneous spectra for nearly 5000 objects across a $\sim 3^{\circ}$ field of view [24–26], and is currently conducting a five-year dark energy survey covering approximately one-third of the sky [27]. Upon completion, the survey will have collected spectra for roughly 40 million galaxies and quasars [28].

In this work, we use the extended photometric sample of luminous red galaxies (LRGs) [29,30], selected from the DESI Legacy Imaging Surveys, which combine data from three telescopes: Blanco for the Dark Energy Camera Legacy Survey (DECaLS), Mayall for the Mayall z -band Legacy Survey (MzLS), and Bok for the Beijing–Arizona Sky Survey (BASS). We make use of the photometric

redshifts provided in Zhou *et al.* [30] for Data Release 9 (DR9), and the stellar mass estimates from [29].

We additionally include bright galaxy sample (BGS) galaxies, as described in Hahn *et al.* [31], Chen *et al.* [32], selected from the same imaging surveys and designed to trace the nearby galaxy distribution with high completeness.

B. Atacama cosmology telescope

We utilize the publicly available CMB lensing convergence (κ) maps from the Atacama Cosmology Telescope (ACT) Data Release 6 (DR6), described in [33–36]. ACT was a 6-meter telescope located at an altitude of 5,190 meters in the Atacama Desert of northern Chile. It operated from 2007 until its decommissioning in 2022, providing high-resolution measurements of the CMB over a wide range of angular scales.

The DR6 lensing dataset used here leveraged multifrequency observations collected between 2017 and 2021, covering roughly a third of the sky in two frequency bands: f090 (77–112 GHz) and f150 (124–172 GHz). The lensing convergence field was reconstructed using a quadratic estimator (QE) applied to both CMB temperature and polarization maps, and included “profile hardened” estimators [35,37] to mitigate the impact of extragalactic foregrounds.

The resulting κ map is provided as spherical harmonic coefficients ($a_{\ell m}$ ’s), along with a corresponding analysis mask. The reconstruction imposes a low-pass filter at $\ell > 3000$ to exclude low signal-to-noise modes and suppress contamination from foregrounds and instrumental effects.

C. Measurement of convergence profiles

We measure the CMB lensing convergence, κ , profiles around galaxies by stacking cutouts from the ACT DR6 lensing map. For each galaxy, we extract a square region centered on its coordinates using the REPROJECT function from the PIXELL package, employing bilinear interpolation to project the 1-arcmin pixel κ map into a tangential projection with a pixel size of 0.5 arcmin. These cutouts are then stacked and averaged in radial bins to construct the galaxy-convergence cross-correlation signal.

We define the radially averaged convergence profile $\hat{\kappa}(\theta_d)$ at angular separation θ_d as

$$\hat{\kappa}(\theta_d) = \mathcal{N}^{-1}(\theta_d) \int d^2\theta W(\theta, \theta_d) \kappa(\theta), \quad (1)$$

where $W(\theta, \theta_d)$ is a top-hat annular (ring) window centered at angular distance θ_d , and $\mathcal{N}(\theta_d)$ is a normalization factor equal to the area of the annulus (i.e., $\mathcal{N}(\theta_d) = \int d^2\theta W(\theta, \theta_d)$). The integration is performed in 2D pixel space for each stacked image.

We bin the signal into 8 linearly spaced annuli in angular separation with $\Delta\theta = 1.5$ arcmin, covering the range from 0 to 10.5 arcmin. This range leverages the need for high signal-to-noise on small scales with the goal of probing the

outskirts of the halo and surrounding large-scale structure. The resulting κ profile represents the average gravitational lensing signal sourced by the projected mass distribution around the selected galaxy sample.

D. Measurement of the gas profiles

In this work, we utilize previously published measurements of the kinetic Sunyaev-Zel’dovich (kSZ) effect obtained via stacking around DESI LRGs to draw a comparison between the total matter distribution and the gas distribution. In this section, we briefly summarize the procedure for self-consistency.

The methodology follows that described in Refs. [12,13] and is implemented in the publicly available ThumbStack pipeline [38]. It requires two ingredients: an estimate of the line-of-sight peculiar velocity for each galaxy and a measurement of the temperature modulation in the CMB at the location of each galaxy. The temperature signal is extracted from CMB maps using a compensated aperture photometry (CAP) filter, which subtracts the mean temperature in an annulus from that in a central disk to isolate small-scale temperature fluctuations associated with individual halos. The resulting temperature profile $T_{\text{kSZ}}(\theta_d)$ is computed as a function of aperture radius θ_d .

The peculiar velocity field is reconstructed by solving the linearized continuity equation in redshift space under the assumption of linear galaxy bias, using the PYRECON package [39]. This yields an estimate of the 3D velocity field from which only the line-of-sight component is used for the kSZ analysis. The kSZ signal $T_{\text{kSZ}}(\theta_d)$ is then estimated using the velocity-weighted stacking estimator from Ref. [12], which correlates the filtered CMB temperature with the reconstructed velocity field. We adopt a velocity cross-correlation coefficient of $r = 0.3$ for the main DESI LRG sample [40] and $r = 0.25$ for the extended sample, based on values of the photometric redshift errors from [30] and the dependence of r on redshift errors from [41]. The uncertainty on that quantity is about 10%; a more precise determination of r using the latest spectroscopic data is deferred to future work. Outliers are removed and a symmetric velocity distribution imposed to suppress contamination from the thermal SZ effect and the cosmic infrared background, ensuring a foreground-free measurement (see [13] for details). In this work, we make use of the final, stacked kSZ profiles for the main and extended LRGs from [13].

III. METHODS

In this section, we describe the simulation-based model we adopt in our analysis for modeling the gas and dark matter distribution.

A. AbacusSummit

AbacusSummit is a suite of large-volume high-resolution cosmological N -body simulations, which was designed to

meet the cosmological simulation requirements of the DESI survey [42,43]. The simulations were run with the cosmological N -body simulation code Abacus [44,45], optimized for GPU architectures on the Summit supercomputer at the Oak Ridge Leadership Computing Facility.

While the AbacusSummit suite spans a wide range of cosmologies and box sizes, here we make use of the BASE resolution boxes, which contain 6912^3 particles in a $2 \text{ Gpc}/h$ box, with a mass resolution of $M_{\text{part}} = 2.1 \times 10^9 M_{\odot}/h$. As this analysis is performed at fixed cosmology, we employ the fiducial cosmology boxes which have cosmological parameters set to their *Planck* 2018 values: $\Omega_b h^2 = 0.02237$, $\Omega_c h^2 = 0.12$, $h = 0.6736$, $A_s = 2.0830 \times 10^{-9}$, $n_s = 0.9649$, $w_0 = -1$, $w_a = 0$.

As we are interested in the DESI LRG and BGS galaxy samples, we adopt the snapshot outputs at $z = 0.3$ (for BGS), 0.5 (for Bin 1 and 2 of LRGs), 0.8 (for Bin 3 and 4 of LRGs). We also make use of the particle outputs available for these snapshots, which feature both a 3% and a 7% subsample. As we are not shot noise dominated in our cross-correlation measurements between galaxies and matter, we opt to only use the A particle subsample, which has 3% of all particles in the simulation. For full details on the simulation products, see Maksimova *et al.* [42]. For some of the internal tests for validating our pipelines, we also utilize the halo light cone catalogs and weak lensing maps generated on the HUGE AbacusSummit simulation [42].

All halo masses quoted in this work correspond to the mass definition adopted by the AbacusSummit halo finder CompaSO [46], which defines the virial mass using the spherical collapse model and the fitting formulas from Bryan and Norman [47].

B. Halo occupation distribution model

To model the distribution of luminous red galaxies (LRGs) within dark matter halos, we adopt a standard (vanilla) five-parameter halo occupation distribution (HOD) framework [48]. In this model, the mean number of central and satellite galaxies in a halo of mass M is given by

$$\langle N_{\text{cen}}(M) \rangle = \frac{1}{2} \left[1 + \text{erf} \left(\frac{\log M - \log M_{\text{cut}}}{2\sigma_{\log M}} \right) \right], \quad (2)$$

$$\langle N_{\text{sat}}(M) \rangle = \langle N_{\text{cen}}(M) \rangle \left(\frac{M - \kappa M_{\text{cut}}}{M_1} \right)^\alpha, \quad \text{for } M > \kappa M_{\text{cut}}. \quad (3)$$

We adopt the AbacusHOD prescription within the ABACUSUTILS package [49]. The five free parameters are

- (i) M_{cut} : the characteristic halo mass at which a halo has a 50% probability of hosting a central galaxy.
- (ii) $\sigma_{\log M}$: the scatter in $\log M$ describing the smooth transition of the central galaxy occupation function.

- (iii) κM_{cut} : the cutoff mass below which no satellites are hosted.
- (iv) M_1 : the typical halo mass required to host one satellite galaxy.
- (v) α : the power-law slope governing the number of satellites in high-mass halos.

This HOD formalism provides a simple, flexible, and empirically motivated description of galaxy bias and halo occupation that we use to construct the three-dimensional galaxy-matter power spectrum $P^{gm}(k, z)$.

C. Lensing observables

To predict the stacked CMB lensing signal for a given galaxy sample, we Fourier transform the galaxy-matter power spectrum into the real-space continuous convergence profile $\kappa(\theta)$ via

$$\kappa(\theta) = \int \frac{\ell d\ell}{2\pi} J_0(\ell\theta) C_\ell^{gk}, \quad (4)$$

where J_0 is the zeroth-order Bessel function and C_ℓ^{gk} is the angular cross-power spectrum between galaxies and convergence, computed from the simulation-measured cross-power spectrum $P_{gm}(k, z)$. In particular, under the Limber approximation, the cross-correlation angular power spectrum is given by

$$C_\ell^{gk} = \int_0^{\chi_*} \frac{d\chi}{\chi^2} W^k(\chi) W^g(\chi) P^{gm} \left(k = \frac{\ell + 0.5}{\chi}, z(\chi) \right) \quad (5)$$

where χ is the comoving distance, χ_* the comoving distance to the last-scattering surface ($z_* = 1089.3$), and $P^{gm}(k, z)$ the three-dimensional cross-power spectrum. The lensing and galaxy kernels are defined as

$$W^k(\chi) = \frac{3H_0^2 \Omega_m}{2c^2} \frac{\chi}{a(\chi)} \left(1 - \frac{\chi}{\chi_*} \right), \quad (6)$$

$$W^g(\chi) = H(z(\chi)) n_g(z(\chi)), \quad (7)$$

where $H(z)$ is the Hubble parameter, H_0 the Hubble constant, c the speed of light, Ω_m the energy density of matter, and $n_g(z)$ the normalized lens galaxy redshift distribution. Since in building our model, we use the AbacusSummit snapshots at fixed redshift, we can substitute the lens redshift distribution with a Dirac delta function at the lens redshift, z_1 , $n_g(z) = \delta^D(z - z_1)$, arriving at

$$C_\ell^{gk} = \frac{1}{\chi_1^2} W^k(\chi_1) P^{gm} \left(k = \frac{\ell + 0.5}{\chi_1}, z_1 \right). \quad (8)$$

Having converted the cross-power spectrum into an estimate of the convergence profile, $\kappa(\theta)$, we apply a low-pass filter with a cutoff at $L_{\text{max}} = 3000$, consistent with the filtering applied in the ACT DR6 lensing map

reconstruction [34]. We then bin the resulting $\kappa(\theta)$ profile into the same annular bins used in the data measurements to obtain the final binned profiles, $\hat{\kappa}(\theta)$. This ensures that we forward-model the lensing signal in a manner that accounts for the full measurement pipeline, enabling an accurate comparison with observations.

D. Validation

A rigorous test we administered is an end-to-end validation of the modeling pipeline using the AbacusSummit light cone maps and catalogs. We utilized the main and extended LRG mock samples from [30], which have a matching redshift distribution and footprint to the DESI Legacy Survey. We first tested the validity of the approach in Sec. III C by directly comparing the stacked profiles from the light cone simulation (i.e., performing an explicit stack using the 2D κ map and LRG-like catalogs on the DESI and ACT footprints) against the predicted $\kappa(\theta)$ from Eq. (4) and found excellent agreement.

An essential component of our modeling pipeline is the assumption that the reconstructed convergence field from the DR6 lensing map provides an unbiased estimate of the true projected mass distribution. To test this, we create a lensed CMB temperature map by first generating a Gaussian realization of the primary CMB using CAMB [50] with the *Planck* 2018 cosmological parameters. This map is then lensed using the convergence map from [51] via the PIXELL [52] package. To approximate ACT DR6 observations, we add white noise with an amplitude of 15 $\mu\text{K-arcmin}$ and convolved the map with a Gaussian beam of $\text{FWHM} = 1.6$ arcmin. We then perform QE CMB lensing reconstruction on the lensed CMB map with the ACT DR6 analysis pipeline which utilizes the SO-LENPIPE [53] package, keeping the settings unchanged. Finally, we directly compare the reconstructed κ to the true input convergence maps over the same scales used in this analysis. We find excellent agreement between the reconstructed and true κ profiles across all angular scales of interest, confirming that the DR6 QE reconstruction yields unbiased measurements of κ for our galaxy sample and modeling framework. This validates the fidelity of our lensing modeling and removes potential concerns about reconstruction bias in the small-scale regime. While we have shown here that for the masses and noise levels of interest the standard QE is appropriate to recover unbiased masses, we note that for more massive clusters and in the low-noise regime of future experiments, asymmetric estimators such as those in [54,55] or optimal methods [56–59] become advantageous.

To further validate the robustness of our method, we conduct two key tests focused on recovering the mean halo mass—our primary science result, for LRG and BGS host halos. First, we evaluate the accuracy of the mean mass inference on a mock galaxy sample generated from a randomly chosen set of HOD parameters. These parameters

are drawn from within the range of our latin hypercube training set (see Sec. III F, i.e., no extrapolation is required), but were not used during training. For this sample, we compute the true mean halo mass as well as its corresponding $\kappa(\theta)$ profile. We then recover the best-fit parameters using a standard minimization procedure with the emulator. The recovered mean halo mass agrees with the true value to within 0.01 dex in $\log_{10}(M/h^{-1}M_{\odot})$, approximately 2%, demonstrating excellent internal consistency of our pipeline in the regime where the emulator is well trained.

To provide a more stringent and conservative test, we also assess the emulator’s performance on a mock sample with fixed mass ($M = 10^{13.5}h^{-1}M_{\odot}$) and no satellites. This is near the upper edge of our mass range and presents a challenging test case, as the emulator has seen few examples at such high masses and was not explicitly trained on pure halo samples. Moreover, unlike realistic HOD samples, the halo-only sample has no satellites and a sharply peaked mass distribution, making it structurally different from the emulator training set. In this case, we find that the recovered mass differs from the true value by 0.05 dex, corresponding to approximately a 10% discrepancy. Given that the actual galaxy samples of interest are more similar to the HOD-based test case than to the extreme pure-halo sample, we adopt a conservative estimate of 5% as the systematic uncertainty on our mean mass measurements.

We also test for potential contamination of the lensing signal by extragalactic foregrounds using the WebSky simulations [60], which include a range of astrophysical foregrounds (e.g., tSZ, CIB, radio sources) embedded in a simulated CMB sky. As a first test, we use lensing maps reconstructed on foregrounds only (but with the weights and analysis choices appropriate for reconstruction in the presence of primary CMB), similarly to what was performed in [35–37]. We find that the maximum bias to $\kappa(\theta)$ is approximately 1.3% for the BGS sample, and less for the others. In addition, we repeat the analysis with the ACT DR6 CIB-deprojected lensing map, finding 5% differences at worst, noting that part of the difference is likely due to the different noise realization and level in the two lensing maps. Finally, the difference between the full minimum variance reconstruction (which includes polarization data) and temperature-only reconstruction is 3% at worst.

Therefore, we conservatively budget a 5% systematic uncertainty to account for foreground biases. This is in line with the results from the profile hardened QE in [37]: at the highest multipole used here ($L = 3000$), the total foreground bias in cross-correlation with low-redshift galaxies is of $\mathcal{O}(5\%)$, but much lower at lower L . We note that smaller scales, for which the bias rapidly increases, are not used in our analysis. This uncertainty is subdominant to the other sources of uncertainty at play (especially the uncertainty on the kSZ measurements to which the lensing results are compared), and therefore, we will not explore this topic further here. However, we note that these biases could be further mitigated by shear-only estimators [61,62],

gradient cleaning methods [63], polarization-only reconstruction [64], the use of asymmetric, or optimal estimators discussed above. Future experiments will also be able to leverage more extensive multifrequency information to further reduce any remaining contamination [65,66].

E. Gas observables

Here we derive the connection between the two-dimensional gas density profile and the observed gas density signal from the kSZ.

The gas is assumed to be fully ionized with primordial abundances at the scales of interest. Under this assumption, the free electron number density is given by

$$n_e = \frac{\rho_{\text{gas}}}{\mu_e m_p}, \quad (9)$$

where $\mu_e = 2/(X_H + 1)$ is the electron chemical potential, $X_H \approx 0.76$ is the primordial hydrogen mass fraction, and m_p denotes the atomic mass unit. Given a model for the three-dimensional gas density profile $\rho_{\text{gas}}(r)$, the stacked (averaged) kSZ temperature fluctuation imprinted on the CMB is given by

$$\frac{T_{\text{kSZ}}(\theta)}{T_{\text{CMB}}} = \frac{v_{\text{rms}}^{\text{true}}}{c} \frac{\sigma_T}{\mu_e m_p} \int_{\text{LoS}} dl \rho_{\text{gas}} \left(\sqrt{l^2 + d_A(z)^2 \theta^2} \right). \quad (10)$$

Here, $v_{\text{rms}}^{\text{true}}$ is the true (three-dimensional) root-mean-square velocity of the galaxy sample (≈ 300 km/s), σ_T is the Thomson scattering cross section, and $d_A(z)$ is the angular diameter distance to redshift z . The integral computes the line-of-sight projection of the gas density profile, yielding a two-dimensional projected gas density.

To enable direct comparison with the CAP-filtered temperature fluctuation maps observed by ACT, we must account for the instrumental beam response. This is done by convolving the projected gas density field with the ACT beam profile $B(\theta)$ in real space. The resulting beam-convolved temperature fluctuation is given by

$$T_{\text{kSZ}}(\theta) = T_{\text{CMB}} \left(\frac{v_{\text{rms}}^{\text{true}}}{c} \right) \frac{\sigma_T}{\mu_e m_p} [\rho_{\text{gas}}^{2D}(\theta) * B(\theta)]. \quad (11)$$

This expression gives the expected kSZ signal for the average galaxy in our sample, incorporating the effects of line-of-sight projection and instrumental resolution. Note that for ACT DR6, the beam FWHM is 1.6 arcmin. To compare with the data, we apply the same CAP filtering procedure described in Sec. II D, which is necessary to account for the signal attenuation introduced by the filter. This ensures a consistent forward-modeling pipeline for connecting theoretical predictions to observational measurements.

F. Emulator

To accelerate inference of halo occupation parameters from CMB lensing measurements, we developed a Gaussian process (GP) emulator for the lensing convergence signal, $\kappa(\theta)$, over a range of angular separations θ . We train this emulator at the mean redshifts for each of the lens samples considered in this study. The emulator takes as input five HOD parameters

- (i) $\log M_{\text{cut}} \in [12.24, 13.34]$.
- (ii) $\log M_1 \in [13.38, 14.38]$.
- (iii) $\sigma \in [0.01, 1.06]$.
- (iv) $\alpha \in [0.35, 1.85]$.
- (v) $\kappa \in [0.01, 2.65]$.

These bounds were sampled using latin hypercube sampling (LHS) to efficiently cover the five-dimensional parameter space. A total of 1000 samples were drawn.

For each set of HOD parameters, we compute $\kappa(\theta)$ at 9 angular bins linearly spaced between 0 and 15 arcminutes. We train an independent GP model for each θ bin using the `scikit-learn` implementation of the `GaussianProcessRegressor` with an RBF kernel plus a white noise term. Inputs were scaled using a standard normalization prior to training.

To evaluate the emulator's accuracy, we held out a test set and compared predicted values against simulated truth. Across all bins, the emulator achieves excellent performance, with a total RMSE $\sim 10^{-4}$ and R^2 scores close to unity ($R^2 \approx 0.9998$). The per bin performance is similarly high, with the worst offender having an RMSE of 0.0001 and R^2 of 0.9997.

The mean fractional error across bins ranges from 0.0006 to 0.0032, well below typical statistical uncertainties on the CMB lensing signal, validating the emulator's fidelity for use in inference.

G. Likelihood

To constrain the halo occupation parameters, we compare the measured CMB lensing signal, $\mathbf{d}_\kappa$, to the model predictions $\mathbf{m}_\kappa(\theta)$ obtained from the Gaussian process emulator (Sec. III F). The likelihood is assumed to be Gaussian in the difference between the observed and predicted $\kappa(\theta)$ profiles

$$\ln \mathcal{L}[\mathbf{d}_\kappa | \theta] = -\frac{1}{2} (\mathbf{d}_\kappa - \mathbf{m}_\kappa(\theta))^T \mathbf{C}_\kappa^{-1} (\mathbf{d}_\kappa - \mathbf{m}_\kappa(\theta)), \quad (12)$$

where θ denotes the set of five HOD parameters and $\mathbf{C}_\kappa$ is the covariance matrix associated with the measurement.

The covariance matrix $\mathbf{C}_\kappa$ is estimated empirically from the stacked $\kappa(\theta)$ profiles of individual galaxies. Specifically, we compute the covariance between radial bins across the galaxy sample, and normalize it by the number of galaxies N_g

$$\mathbf{C}_\kappa \approx \frac{1}{N_g} \hat{\mathbf{C}}, \quad (13)$$

where $\hat{\mathbf{C}}$ is the unnormalized empirical covariance. This method assumes the independence of individual galaxy profiles and yields an estimate accurate to approximately 10% [67]. We find that the off-diagonal terms of the resulting $\mathbf{C}_\kappa$ are dominated by noise, which can lead to instabilities when inverting the matrix. To regularize the covariance, we perform an eigenvalue decomposition and clip the condition number by enforcing a maximum eigenvalue ratio of 100

$$\mathbf{C}_\kappa \rightarrow \mathbf{U}\mathbf{\Lambda}_{\text{clip}}\mathbf{U}^\top, \quad (14)$$

where $\mathbf{C}_\kappa = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ and the eigenvalues in $\mathbf{\Lambda}$ are clipped such that $\lambda_{\text{max}}/\lambda_{\text{min}} \leq 100$. This procedure preserves the dominant structure in the covariance while mitigating the impact of noisy modes. We check that this changes the χ^2 value by about 10%–15%.

We explore the posterior distribution of the HOD parameters using dynamic nested sampling, as implemented in the DYNesty sampler [68]. This approach efficiently samples multimodal posteriors and computes the Bayesian evidence. The prior on each HOD parameter is uniform over the same range used to train the emulator (see Sec. III F). In practice, this allows us to evaluate the likelihood over the domain of validity of the emulator, ensuring accurate interpolation.

To validate our results, we also perform a direct minimization of the negative log-likelihood to identify the best-fitting point. We find good agreement between the minimizer and the peak of the posterior distribution in most cases, though the nested sampler provides additional information on the parameter posteriors and degeneracies.

IV. RESULTS

In this section, we present our main findings regarding the masses of the DESI samples and compare the gas density profiles against the dark matter profiles derived from the measured CMB lensing convergence profiles.

A. Mass estimates

We begin with Fig. 1, which examines the stacked κ map at the positions of DESI extended LRGs using the ACT DR6 CMB lensing convergence map. This image corresponds to the average projected matter distribution around the selected galaxy sample. The central feature corresponds to the mean lensing signal around DESI galaxy halo hosts and is clearly detected. The convergence map released by ACT features a multipole cutoff of $L < 3000$, which is designed to reduce contamination from small-scale foregrounds and noise. At the same time, it reduces the small-scale power, causing the profile to appear smeared and broadened. This effect is also visible in the surrounding structure of the 2D stacked image, where the correlations introduced by the filtering manifest as coherent features on a few-arcminute scales. These features are a known




FIG. 1. Stacked convergence map centered on the positions of DESI extended LRGs using the ACT DR6 κ map. The central feature corresponds to the mean lensing signal around DESI galaxy halo hosts. The κ map released by ACT features a Fourier-space cutoff at $L < 3000$, which is applied to suppress small-scale noise. It introduces coherent features on few-arcminute scales and smears out the 2D profile.

consequence of the filtering and do not reflect the intrinsic structure of the halos.

In Fig. 2, we find that our halo model provides an excellent fit to the DR6 κ signal for both main and extended LRGs. The BGS sample, which consists of lower-redshift and lower-bias galaxies, shows larger residuals between the data and the model prediction. This is not unexpected: BGS galaxies represent a more diverse population, including lower-mass and star-forming systems for which our simple five-parameter HOD may not be as accurate. In contrast, the LRG samples are dominated by red, massive galaxies, where the standard HOD form is known to perform well. The relatively poor fit for the BGS sample likely reflects two important factors: (a) as the sample spans a wide range of low redshifts, $z = (0, 0.5)$, there is a substantial mixture of angular scales in the averaged κ profiles; additionally, our theoretical model is evaluated at a fixed redshift, $z = 0.3$, and hence lacks the redshift evolution seen in the data profiles; (b) the increasing importance of the two-halo term and the need to account for effects such as central/satellite assembly bias or more complex halo occupation dependencies, which are more pronounced for the diverse BGS objects.

Error on the mass estimate: Nevertheless, since our primary goal is to infer the mean halo mass of these samples, the limitations of the model are acceptable. Systematic errors from HOD model misspecification and




FIG. 2. Lensing convergence κ as a function of aperture angle for the three DESI galaxy samples: main LRGs (red), extended LRGs (blue), and BGS galaxies (orange). Points show the measured κ signal in the DR6 map, and the solid lines show the best-fit models based on the simulation-based framework described in Sec. III C. Error bars correspond to the square root of the diagonal of the covariance matrix described in Sec. III G. The shaded regions around the best-fit curves denote 68% confidence intervals. All model fits are computed by forward-modeling the predicted galaxy-matter correlation into $\kappa(\theta)$, applying an $L < 3000$ multipole cut and binning in the same annular bins as the ones used in the data. We attribute the worst fit in the case of the BGS sample to the rapid redshift evolution of the signal (as a function of angular scale), which is not accounted for in our fixed-snapshot approach.

unmodeled assembly bias and the stellar component are expected to affect the inferred halo mass at the $\sim 5\%$ level, which is comparable to other sources of uncertainty in the analysis. Since we take a $\sim 5\%$ uncertainty on modeling and a further (uncorrelated) $\sim 5\%$ uncertainty due to possible foreground biases, we shall quote a $\sim 5\sqrt{2} \approx 7\%$ systematic uncertainty per sample, noting that it dominates over the statistical uncertainty of the lensing measurement, but is still smaller than the uncertainty on the gas profiles from kSZ and therefore not the limiting factor in our analysis at present.

One may also worry about the effect of baryonic feedback on CMB lensing, which can redistribute gas beyond the virial radius and lead to a mismatch between the true and modeled mass profiles. However, since we restrict our modeling to multipoles $L < 3000$, our analysis focuses on sufficiently large angular scales where baryonic effects are minimized and the full halo mass is still captured well. Baryon effects in CMB lensing have been studied in [69] and shown to be negligible for most of the multipole L range used in our analysis, though they may become sizeable around $L \approx 3000$ for the most extreme models. While smaller than the kSZ uncertainties in this analysis, they motivate a joint and self-consistent analysis of both

TABLE I. Best-fit values and 68% confidence intervals for the five HOD parameters and three derived parameters: comoving number density $\bar{n}$ (in $[\text{Mpc}/h]^{-3}$), satellite fraction f_{sat} , and mean halo mass $\langle M_{\text{halo}} \rangle$ (in $M_{\odot}/h$). Results are shown for each of the three tracer samples: main LRGs, extended LRGs, and BGS. All mass units are in $M_{\odot}/h$. The masses correspond to the virial mass definition from Bryan and Norman [47]. We budget around 7% for the systematic bias on the mean halo mass, as described in the main text.

Param.	Main LRGs	Extended LRGs	BGS _{log M > 10.5}
$\log M_{\text{cut}}$	$12.683^{+0.129}_{-0.063}$	$12.491^{+0.07}_{-0.045}$	$12.133^{+0.102}_{-0.125}$
$\log M_1$	$14.063^{+0.241}_{-0.395}$	$14.196^{+0.144}_{-0.433}$	$13.83^{+0.276}_{-0.67}$
$\sigma_{\log M}$	$0.133^{+0.179}_{-0.083}$	$0.108^{+0.126}_{-0.06}$	$0.107^{+0.194}_{-0.074}$
α	$0.848^{+0.611}_{-0.362}$	$0.642^{+0.725}_{-0.222}$	$1.219^{+0.501}_{-0.77}$
κ	$1.245^{+0.924}_{-0.816}$	$0.941^{+1.075}_{-0.737}$	$1.069^{+0.992}_{-0.79}$
$\bar{n} \times 1000$	$0.745^{+0.223}_{-0.133}$	$1.254^{+0.223}_{-0.198}$	$3.608^{+1.615}_{-0.723}$
f_{sat}	$0.08^{+0.104}_{-0.057}$	$0.098^{+0.091}_{-0.078}$	$0.108^{+0.195}_{-0.032}$
$\log \tilde{M}_h$	$13.179^{+0.029}_{-0.021}$	$13.025^{+0.022}_{-0.016}$	$13.022^{+0.091}_{-0.061}$

kSZ and CMB lensing (for example, similar to what was performed in [20,21]), which is left to upcoming work.

Table I summarizes the best-fit halo occupation distribution (HOD) parameters and derived quantities for the main LRG, extended LRG, and BGS samples. As expected, the main LRGs occupy the most massive halos on average, followed by the extended LRGs and then BGS, which reflects their target selection and redshift distributions. The satellite fraction and number density are highest for BGS, which is consistent with it being a low-redshift, high-density sample. On the other hand, they are lower for the extended LRGs, and lowest for the main LRGs.

The five HOD parameters (threshold mass M_{cut} , scatter $\sigma_{\log M}$, satellite mass scale M_1 , slope α , and satellite cutoff κ) all follow the trend of increasing with tracer mass and decreasing redshift. Main LRGs, being the most massive and sparsest sample, require higher occupation thresholds and steeper satellite scaling to match the clustering and number densities. Overall, the results across samples are internally consistent and reflect our expectations from both theory and observational selection.

Figure 3 shows the correlation matrix of the measured $\kappa(\theta)$ signal for the extended sample, combining all four redshift bins. We normalize the covariance matrix to obtain the correlation coefficients between the angular bins. We note that the correlation matrix is qualitatively identical across the different samples. A clear blocklike pattern emerges with alternating positive and negative correlated patches. This structure is characteristic of lensing measurements derived from filtered maps with an $\ell < 3000$ cut, which correlates structure on angular scales of order a few arcminutes, inducing coherent fluctuations across adjacent angular bins in the profile.




FIG. 3. Correlation matrix of the lensing signal $\kappa(\theta)$ for the extended sample, combining all four redshift bins. The pattern is representative of all samples and redshift bins, so we show this case as a typical example. We see a characteristic block structure with alternating positive and negative correlations, which arises due to the $\ell < 3000$ multipole cut applied to the lensing maps, as it induces correlations on angular scales of a few arcminutes.

B. Comparison with kSZ

In this section, we compare the convergence profiles $\kappa(\theta)$ measured in this work against previous measurements of the kSZ profiles. To examine the connection between the kSZ effect and gravitational lensing convergence (κ), we compare the amplitudes of the two signals for the same galaxy sample in Fig. 4. Specifically, we focus on the main LRG sample and consider the kSZ amplitude at the third radial bin from [13], and the κ amplitude at the first radial bin from this work, corresponding to the highest SNR bins in our end-to-end simulation tests [70]. The four points in the plot correspond to different stellar mass selections: the full sample (all galaxies), and three stellar mass bins: $10^{11.0} < M_*/M_\odot < 10^{11.25}$, $10^{11.25} < M_*/M_\odot < 10^{11.5}$, and $M_* > 10^{11.5} M_\odot$. The dashed black line, fitted to the points with a linear model, serves to guide the eye.

We find that the points lie approximately on a straight line. This linear relationship provides strong empirical support for the idea that the kSZ signal traces the total gas mass within an aperture, while the κ signal traces the total matter (halo) mass. The proportionality between the two signals is consistent with the approximate expectation¹ that $M_{\text{gas}} \propto M_{\text{halo}}$, further validating both observables as tracers of halo-scale baryon and matter content.

Next, we compare the radial profiles of the total matter and gas density. To this end, we take the best-fit HOD

¹This statement neglects the mass dependence of feedback, which is limited here given our limited range of masses, but would likely be detectable if we had access to a larger mass range.




FIG. 4. Scatter plot comparing the kSZ amplitude at the third radial bin (from [13]) to the κ amplitude at the first radial bin (from this work) for the main LRG sample. Each point represents a different stellar mass selection: All galaxies (black point), and the stellar mass bins $10^{11.0} < M_* < 10^{11.25} M_\odot$, $10^{11.25} < M_* < 10^{11.5} M_\odot$, and $M_* > 10^{11.5} M_\odot$. The dashed black line is fitted to the points with a simple linear model and serves to guide the eye. The approximately linear trend indicates that the kSZ and κ signals scale proportionally with gas mass and total halo mass, respectively, supporting the interpretation that $M_{\text{gas}} \propto M_{\text{halo}}$. The “max SNR” label always refers to the third radial bin for kSZ and the first radial bin for the $\kappa(\theta)$ profile.

values for the main and extended LRG samples (see Table I) and compute the galaxy-matter cross-power spectrum. We then convert it to a kSZ-like CAP measurement by adopting Eq. (10) but replacing the gas density profile, $\rho_{\text{gas}}^{2D}(\theta)$, with the matter density profile, $\rho_m^{2D}(\theta)$, which is in turn obtained from

$$\rho_m^{2D}(\theta) = \int \frac{\ell d\ell}{2\pi} J_0(\ell\theta) C_\ell^{gm}, \quad (15)$$

where C_ℓ^{gm} can be linked to the three-dimensional cross-power spectrum, $P^{gm}(k, z)$, analogously to Eq. (8).

Figure 5 presents a direct comparison between the gas density profile around LRGs inferred from kSZ measurements and the total matter density profile derived from gravitational lensing. The two panels show results for the main (top) and extended (bottom) LRG samples. The lensing-based profiles [derived from $\kappa(\theta)$] are converted into the same units as the kSZ measurement as described above.

What makes this comparison especially powerful is that, unlike earlier works interpreting the kSZ signal, we no longer need to marginalize over the halo mass or rescale the theoretical prediction. Our lensing measurement provides an independent and accurate mass calibration, enabling a largely parameter-free comparison. At large




FIG. 5. Comparison between the gas density profile measured from the kSZ effect [13] and the total matter (dark matter-dominated) profile inferred from our $\kappa(\theta)$ measurements, for the main (top panel) and extended (bottom panel) LRG samples. The κ -based profile has been translated into the same units as the kSZ profile, assuming that gas traces the total matter distribution and that there is no baryonic feedback. The profiles are plotted as a function of aperture radius (in comoving Mpc/h). The shaded regions around the κ -derived curves denote 68% confidence intervals, and the error bars on the kSZ points indicate the error on the kSZ measurements. The agreement on large scales indicates that the baryon fraction is recovered sufficiently far from the center, indicating no evidence of large amounts of unbound gas.

apertures (2–3 Mpc/h, i.e., a few virial radii), we find excellent agreement between the two profiles, supporting the expectation that on large scales, gas and dark matter trace each other well. Thus, the entire baryonic content is recovered sufficiently far from the group center. We note that for the extended LRG sample, the gas appears to lie systematically lower than the matter, suggesting that a fraction of the gas surrounding these galaxies has been

pushed out even further out and is unbound. Intuitively, this makes sense, as the mean mass of the extended sample is slightly lower and thus the gravitational pull the gas feels toward the group center is slightly weaker.

However, at smaller apertures, we observe a significant discrepancy between the lensing-inferred and kSZ-inferred profiles at the level of $\sim 6.5\sigma$, indicative of large baryonic feedback processes. These processes push gas beyond the virial radius, causing the gas profile to continue rising at larger scales, whereas the dark matter profile flattens out around $\sim 1.5 R_{\text{vir}}$. This behavior aligns with previous findings [13–15, 17, 71] of large baryonic feedback and underscores the importance of feedback in redistributing baryons.

The only remaining source of systematic uncertainty comes from the velocity reconstruction used in the kSZ analysis, which affects the amplitude of the recovered gas profiles at the 10%–15% level (as quantified in [13, 15]). Nonetheless, the comparison clearly demonstrates the complementarity of weak lensing and kSZ as probes of the matter and gas distribution, and marks an important step forward in constraining feedback models and cosmology robustly and self-consistently.

We note that the conversion from the measured $\kappa(\theta)$ profile to the corresponding CAP total matter density profile, $\rho_{\text{matter}}(\theta)$, relies on an intermediate step involving simulations. However, this step is largely model-independent: it primarily serves to extend the range of the ℓ -modes in the cross-spectrum $C_{\ell}^{\kappa g}$ —which $\kappa(\theta)$ effectively probes, to smaller scales not directly accessible in the data. Since $C_{\ell}^{\kappa g}$ is a smooth and featureless function, this extrapolation could equally well be performed using a spline fit or similar functional extension without reference to a specific simulation. We verify in simulations that such approaches yield consistent results with our method, confirming that our inferred matter profiles are robust to the details of this modeling choice.

Figure 6 shows the inferred baryon fraction $f_{\text{gas}}(< R)/(\Omega_b/\Omega_m)$ as a function of aperture radius for both the main and extended LRG samples, where Ω_b/Ω_m is the cosmic baryon fraction. This ratio is computed by dividing the gas density profile derived from the kSZ signal [13] by the total matter profile inferred from our $\kappa(\theta)$ measurements, and correcting for the CMB beam suppression using a simulation-calibrated transfer function (note that this correction is about 3%). Notably, this transfer function is only weakly sensitive to baryonic feedback effects, especially at large radii.

On large scales ($R \gtrsim 5$ arcmin), the main sample recovers a baryon fraction consistent with unity, indicating that essentially all the expected baryons are accounted for in these regions. For the extended sample, we observe a slight shortfall ($f_{\text{gas}}/(\Omega_b/\Omega_m) < 1$), which could stem from statistical fluctuations, a slight mischaracterization of the correlation coefficient r in the estimator due to the larger




FIG. 6. Cumulative baryon fraction, $f_{\text{gas}}(< R)/(\Omega_b/\Omega_m)$, as a function of aperture radius for the main (top) and extended (bottom) LRG samples. This is computed by dividing the kSZ-inferred gas profile by the κ -based total matter profile, after correcting for the beam. A value of unity (horizontal dashed line) indicates the cosmic baryon fraction. Vertical dashed lines show the halo virial radii (0.6 and 0.53 Mpc/h). The black shaded band shows results from TNG300 for a mainlike sample, using the same method. The simulation overpredicts the gas fraction by about 4σ , reinforcing the case for stronger baryonic feedback in the real Universe. At the virial radius, our inferred gas fraction is in excellent agreement with the empirical relation derived by Popesso *et al* [72] using eROSITA data, which predicts $f_{\text{gas}}(R_{200m})/(\Omega_b/\Omega_m) \approx 0.34$ for main-LRG-like halos.

photometric noise of that sample, or the presence of gas that is no longer bound to the halo. These statements can be rephrased in terms of the “closure radius” R_c , introduced by [73], which is the radius within which all baryons associated with haloes are found. Our results suggest that for the samples of galaxies studied here, $R_c \gtrsim 3R_{\text{vir}}$.

Most strikingly, at small radii ($R \lesssim 1$ arcmin), we find $f_{\text{gas}}/(\Omega_b/\Omega_m) \sim 0.3$, indicating that a substantial fraction of baryons have been removed from the inner halo regions—strong evidence for the effects of feedback mechanisms such as AGN-driven outflows. This result provides direct observational support for baryon depletion in massive galaxy halos and highlights the complementarity of kSZ and κ measurements in tracing baryon dynamics.

We compare our gas fraction at the virial radius with the empirical relation from Popesso *et al.* [72], based on eROSITA data for low-redshift ($z < 0.2$) systems, and find excellent agreement: $f_{\text{gas}}^{\text{X-ray}}(R_{200m})/(\Omega_b/\Omega_m) \approx 0.34$ compared with $f_{\text{gas}}^{\text{kSZ}}(R_{200m})/(\Omega_b/\Omega_m) \approx 0.3$ and well within the uncertainty. Their analysis, like ours, indicates that baryonic feedback must be significantly stronger in the real Universe than in most state-of-the-art hydrodynamical simulations, which they find to overpredict the gas content, particularly in the $10^{13.5} - 10^{14.5} M_\odot$ mass range. While the redshift range differs from ours ($z \sim 0.7$), this consistency suggests that feedback capable of expelling gas from halos without overly suppressing star formation is a necessary ingredient across cosmic time. Our results are also consistent with the self-similarity of the gas fractions with redshift highlighted by Fig. 5 in [74]. Similar conclusions were also recently obtained by [75] by using a “baryonification” prescription to relate f_{gas} measured with eROSITA data to kSZ profiles from BOSS galaxies. We leave a more direct comparison to x-ray data for future work.

To facilitate a robust comparison with simulations, we also compute the gas fraction in the TNG300 simulation for a main-LRG-like sample, shown as a black shaded region. This is done by taking the ratio of the kSZ and dark matter profiles in the simulation, using the same beam-corrected method applied to the data. We find that the TNG300 gas fraction is consistently higher by $\gtrsim 4\sigma$ than the observed one at all radii, only approaching unity at large distances. This underscores the fact that baryonic feedback in TNG300 appears to be weaker than in the real Universe. Because this method compares relative rather than absolute quantities, it is also less sensitive to uncertainties in normalization and mean halo mass, making it a more direct probe of feedback strength.

To better understand the effects of baryonic feedback on the gas distribution around LRG-like halos, we analyze the compensated aperture photometry (CAP) profiles in kSZ units for two hydrodynamical simulations: Illustris and TNG300. For both simulations, we construct stellar-mass-selected galaxy samples using a number density cut such that the mean halo mass matches that of the observed sample, $\log_{10}(\bar{M}_{\text{halo}}/M_\odot h^{-1}) = 13.22$. The effective number density cut that we apply to each is 0.9 and $1.1 \times 10^{-4} [\text{Mpc}/h]^{-3}$, respectively. This ensures a fair comparison with the data. We note that both of these simulations, and especially Illustris, lack the most massive




FIG. 7. CAP profiles in kSZ units for gas (solid lines) and dark matter (dashed lines) from the Illustris and TNG300 simulations. Stellar-mass-selected samples are constructed such that the mean halo mass is $\log(\bar{M}_{\text{halo}}/M_{\odot}h^{-1}) = 13.22$ in both cases, matching the main LRG sample in the data (see Table I). The gas profiles reflect the strength of baryonic feedback: Illustris (red) exhibits a significantly steeper profile and reduced amplitude at large apertures due to strong AGN feedback that pushes out, heats up, and unbinds gas from halos. TNG300 (blue) shows a shallower gas profile and no evidence for unbound gas. The dark matter profiles from both simulations are nearly identical, validating the halo matching and convergence of the N -body solvers. No artificial rescaling has been applied to the simulation outputs.

galaxy groups due to their smaller volume, which artificially upweights the less massive galaxy groups when computing the average profiles.

In Fig. 7, we show the CAP profiles of gas (solid lines) and dark matter (dashed lines) for both simulations. As expected, baryonic feedback redistributes gas within halos, steepening the gas profiles relative to the dark matter. TNG300, which includes more moderate feedback prescriptions, shows only a modest steepening. Illustris, by contrast, exhibits much steeper gas profiles, indicating overly strong feedback that ejects gas out to large radii or beyond the halo’s virial boundary. In addition, Illustris shows a suppressed gas profile amplitude at large apertures due to unbound gas escaping group-sized halos. This feature is not prominently seen in the data (except perhaps marginally in the extended sample [13]), though the observed kSZ profile depends on the velocity–reconstruction cross-correlation coefficient r , which affects the amplitude.

Importantly, the dark matter profiles in both simulations agree remarkably well, demonstrating that the underlying mass distributions and stellar mass selections are consistent. These profiles are shown without any normalization or rescaling applied. Qualitatively, they agree extremely well with the data curves shown in Fig. 5.




FIG. 8. Top: lensing signal $\kappa(\theta)$ for the main sample, split into cumulative stellar mass bins: all galaxies (solid), $\log(M_*/M_{\odot}) > 11$ (dashed), > 11.25 (dash-dotted), and > 11.5 (dotted). The solid lines correspond to the best-fit model predictions, and the error bars are obtained from the diagonal of the covariance matrix. Bottom: same, but for the extended sample. As expected, the signal amplitude increases with stellar mass threshold due to the corresponding increase in mean halo mass. The extended sample shows systematically lower amplitudes, reflecting its lower average halo mass.

C. Mass evolution

In Fig. 8, we show the measured convergence profiles $\kappa(\theta)$ for the main and extended samples, split into three cumulative stellar mass bins, along with the full sample. As expected, the amplitude of the signal increases with stellar mass threshold. This behavior reflects the well-established correlation between stellar mass and halo mass: more massive galaxies typically reside in more massive halos, which produce stronger lensing signals.

Since all samples lie at similar redshifts, the observed differences in signal amplitude are attributable to variations in halo mass rather than redshift evolution. Compared to the main sample, the extended sample exhibits systematically lower signal amplitudes, consistent with its lower average

halo mass. The difference between the full sample and the $\log(M_*/M_\odot) > 11$ bin is also more pronounced for the extended sample. This is because the stellar mass completeness limit for extended is slightly lower, and the full sample includes a larger fraction of lower-mass galaxies, which reduces the mean halo mass and thus the overall signal amplitude.

Next, we study the stellar-to-halo mass ratio (SHMR) for the main and extended LRG samples when we split the galaxies into 4 bins. We first quote the mean halo masses per bin, and we comment on the SHMR, comparing it against the hydrodynamical simulation TNG300 [76].

We summarize the mean halo masses for the main LRG and extended LRG samples below, computed using our lensing-based inference method. The quoted uncertainties represent the 68% confidence intervals, and all halo masses are reported in $\log(M_{\text{halo}}/M_\odot h^{-1})$.

Main LRG:

- (i) All galaxies: $13.179^{+0.029}_{-0.021}$.
- (ii) $M_* > 10^{11.0} M_\odot$: $13.190^{+0.028}_{-0.020}$.
- (iii) $M_* > 10^{11.25} M_\odot$: $13.282^{+0.026}_{-0.019}$.
- (iv) $M_* > 10^{11.5} M_\odot$: $13.531^{+0.022}_{-0.019}$.

Extended LRG:

- (i) All galaxies: $13.025^{+0.022}_{-0.016}$.
- (ii) $M_* > 10^{11.0} M_\odot$: $13.096^{+0.020}_{-0.015}$.
- (iii) $M_* > 10^{11.25} M_\odot$: $13.258^{+0.017}_{-0.014}$.
- (iv) $M_* > 10^{11.5} M_\odot$: $13.524^{+0.015}_{-0.014}$.

We observe a clear trend in both samples: the mean halo mass increases monotonically with the stellar mass threshold. This is consistent with expectations from galaxy-halo coevolution models and validates that our lensing measurements robustly capture the stellar-to-halo mass relation across the LRG population.

We next examine the stellar-halo mass relation for the extended LRG samples across four stellar mass bins in Fig. 9 and compare it with simulations. We note that we add 0.3 dex of Gaussian scatter in $\log M_*$ to the simulated galaxies to mimic observational uncertainties in the mass determination of the DECaLS catalogs [77]. The stellar mass bins correspond to: (1) all galaxies, (2) $\log(M_*/M_\odot) > 11.0$, (3) > 11.25 , and (4) > 11.5 . For each bin, we show the mean stellar mass and mean halo mass inferred from the DESI DR6 CMB lensing measurements, and compare them against galaxies in the TNG300 simulation. For the main LRG sample, the mean stellar masses across the four bins are: $2.29 \times 10^{11} M_\odot$, $2.36 \times 10^{11} M_\odot$, $2.84 \times 10^{11} M_\odot$, and $4.24 \times 10^{11} M_\odot$, respectively. For the extended LRG sample, the mean stellar masses are: $1.18 \times 10^{11} M_\odot$, $1.93 \times 10^{11} M_\odot$, $2.67 \times 10^{11} M_\odot$, and $4.24 \times 10^{11} M_\odot$, respectively.

Across all bins, we find that the mean halo masses inferred from the DESI data are in good agreement with those of the matched galaxy samples in the TNG300 simulation, validating our lensing-based mass estimates. The stellar mass trend is also well reproduced, with




FIG. 9. Stellar-halo mass relation for the extended LRG samples from data and simulations. We show the mean stellar mass versus mean halo mass in four stellar mass-selected bins: all galaxies, $\log(M_*/M_\odot) > 11.0$, > 11.25 , and > 11.5 . Results are shown both for DESI data (circles) and the matched galaxies in the TNG300 simulation with added stellar mass scatter (crosses). The TNG300 results include a 0.3 dex Gaussian scatter in $\log M_*$ to emulate observational uncertainties. The data points have error bars in the x direction from the halo masses inferred in this work and error bars in the y direction from the systematic bias of 0.1 dex on the stellar mass estimates reported in Appendix C of Ref. [29].

increasing halo mass corresponding to increasing stellar mass, as expected.

However, we observe that the TNG300 stellar masses are systematically higher than the DESI estimates in each bin. This offset is consistent with known findings in the literature: for instance, [78] noted that TNG predicts slightly elevated stellar masses compared to data, which may point to insufficient quenching of star formation in massive galaxies in the simulation. Since our stellar mass estimates for DESI are derived from photometry-trained machine learning models [77], some additional systematics may contribute to the offset, though the qualitative agreement remains good. We leave for future work the exploration of using more accurate stellar mass estimates (e.g., inferred from the full galaxy spectra as opposed to photometry).

In Fig. 10, we focus on the most massive stellar mass bin of the main sample ($\log_{10}(M_*/M_\odot) > 11.5$), for which a direct comparison with [13] is possible. The top panel shows the comparison between the gas profile inferred from the kSZ signal and the total matter profile derived from $\kappa(\theta)$ measurements. The bottom panel presents the corresponding cumulative gas fraction profile, $f_{\text{gas}}(< R)/(\Omega_b/\Omega_m)$, corrected for the effect of the ACT beam using a transfer function calibrated on simulations. We restrict our detailed modeling to the highest bin, as stellar mass-selected subsamples with both lower and upper bounds lead to




FIG. 10. Top panel: comparison between the gas density profile inferred from the kSZ signal [13] and the total matter profile inferred from our $\kappa(\theta)$ measurements for the highest stellar mass bin of the main LRG sample ($\log_{10}(M_*/M_\odot) > 11.5$). The κ -derived profile is rescaled into kSZ units as outlined in Sec. III. Bottom panel: cumulative baryon fraction, $f_{\text{gas}}(< R)/(\Omega_b/\Omega_m)$, computed from the ratio of the kSZ and κ profiles and corrected for beam effects using a transfer function derived from simulations. The horizontal dashed line marks the cosmic baryon fraction, and the vertical dashed line denotes the halo virial radius for this subsample ($R_{\text{vir}} = 0.83 \text{ Mpc}/h$ comoving). At the virial radius, our inferred gas fraction matches the empirical relation from Popesso *et al.* [72], based on eROSITA observations very closely, which predicts $f_{\text{gas}}(R_{200m})/(\Omega_b/\Omega_m) \approx 0.43$ for halos similar to those in the main-LRG sample.

nonstandard HOD shapes that are challenging to model robustly via the vanilla HOD formalism typically adopted for LRGs.

We find that, as with the full main sample (cf. Fig. 6), the gas fraction asymptotes to unity at large radii, indicating that the baryons are fully accounted for beyond a few times

the virial radius. Notably, the gas fraction in this high-mass bin is slightly elevated relative to the full sample, which is consistent with expectations: more massive halos have deeper gravitational potentials and are therefore less susceptible to baryon removal via feedback processes such as AGN activity. We also see indications that the baryon fraction reaches unit at smaller apertures compared with the full main sample, which is consistent with the picture that less gas is expelled from high-mass halos. However, in this regime the error bars are also largest and thus follow-up with a larger cluster sample would be beneficial.

D. Redshift evolution

Figure 11 displays the redshift evolution of the lensing convergence profile, $\kappa(\theta)$, for the main and extended LRG samples, split into four tomographic bins. The overall




FIG. 11. Redshift evolution of the lensing signal $\kappa(\theta)$ for the main (top panel) and extended (bottom panel) LRG samples, shown in four tomographic bins. We observe a mild increase in the signal amplitude with redshift, particularly near $\theta \sim 5$ arcmin where the lensing kernel is most sensitive. The curves correspond to the best-fit model predictions, and error bars are derived from the diagonal of the covariance matrices.

TABLE II. Best-fit HOD and derived parameters for the main LRG sample, shown across four redshift bins. The redshift bins correspond to: Bin 1: $0.4 < z < 0.54$, Bin 2: $0.54 < z < 0.713$, Bin 3: $0.713 < z < 0.86$, Bin 4: $0.86 < z < 1.024$. All mass units are in $M_\odot/h$, and comoving number densities are in $[\text{Mpc}/h]^{-3}$. The masses correspond to the virial mass definition from Bryan and Norman [47]. We budget around 5% for the systematic bias on the mean halo mass (see Section III D).

Param.	z_1	z_2	z_3	z_4
$\log M_{\text{cut}}$	$12.61^{+0.39}_{-0.15}$	$12.63^{+0.23}_{-0.09}$	$12.73^{+0.27}_{-0.1}$	$12.63^{+0.32}_{-0.11}$
$\log M_1$	$13.91^{+0.33}_{-0.35}$	$14.02^{+0.26}_{-0.41}$	$13.98^{+0.29}_{-0.37}$	$13.97^{+0.29}_{-0.37}$
$\sigma_{\log M}$	$0.3^{+0.29}_{-0.21}$	$0.18^{+0.25}_{-0.13}$	$0.21^{+0.23}_{-0.14}$	$0.23^{+0.25}_{-0.15}$
α	$1.0^{+0.56}_{-0.48}$	$0.84^{+0.6}_{-0.39}$	$0.96^{+0.57}_{-0.44}$	$0.96^{+0.56}_{-0.45}$
κ	$1.3^{+0.92}_{-0.88}$	$1.3^{+0.91}_{-0.9}$	$1.33^{+0.89}_{-0.89}$	$1.28^{+0.91}_{-0.85}$
$\bar{n} \times 1000$	$1.3^{+0.6}_{-0.41}$	$1.09^{+0.37}_{-0.25}$	$0.69^{+0.26}_{-0.19}$	$0.91^{+0.34}_{-0.27}$
f_{sat}	$0.08^{+0.09}_{-0.05}$	$0.09^{+0.1}_{-0.06}$	$0.08^{+0.09}_{-0.05}$	$0.07^{+0.09}_{-0.05}$
$\log \bar{M}_h$	$13.19^{+0.09}_{-0.06}$	$13.22^{+0.05}_{-0.04}$	$13.2^{+0.05}_{-0.03}$	$13.12^{+0.06}_{-0.04}$

trend shows a mild increase in amplitude with redshift, particularly around $\theta \sim 5$ arcmin. This monotonic behavior is consistent with the expectations from the lensing kernel, which peaks at higher values for more distant and more massive lenses. The observed redshift dependence complements the corresponding mass evolution summarized in Table I, confirming that the higher-redshift bins are, on average, associated with more massive halos. The agreement between data and model remains good across all bins, and no strong deviations from the fiducial HOD form are required to explain the redshift dependence.

In Table II, we show the HOD fits and derived quantities for the main LRG sample split into four redshift bins. Across the redshift range, the derived quantities (number density, satellite fraction, and mean halo mass) remain relatively stable, with no statistically significant trends. The stability in these parameters is expected, as the main LRG selection targets a stable population of massive red galaxies across a relatively narrow redshift range. Minor fluctuations in the HOD parameters likely reflect small differences in completeness or selection efficiency, but the overall HOD structure remains consistent with a massive, primarily central-dominated population.

TABLE III. Same as Table II, but for the extended LRG sample. The redshift bins are identical to those used in Table II.

Param.	z_1	z_2	z_3	z_4
$\log M_{\text{cut}}$	$12.49^{+0.37}_{-0.1}$	$12.44^{+0.17}_{-0.07}$	$12.52^{+0.25}_{-0.08}$	$12.38^{+0.33}_{-0.07}$
$\log M_1$	$13.96^{+0.29}_{-0.37}$	$14.01^{+0.27}_{-0.49}$	$14.05^{+0.24}_{-0.39}$	$14.1^{+0.21}_{-0.39}$
$\sigma_{\log M}$	$0.24^{+0.31}_{-0.17}$	$0.14^{+0.23}_{-0.1}$	$0.17^{+0.25}_{-0.11}$	$0.2^{+0.27}_{-0.13}$
α	$0.92^{+0.58}_{-0.44}$	$0.76^{+0.66}_{-0.34}$	$0.89^{+0.61}_{-0.4}$	$0.82^{+0.63}_{-0.35}$
κ	$1.28^{+0.94}_{-0.89}$	$1.23^{+0.94}_{-0.85}$	$1.23^{+0.94}_{-0.83}$	$1.1^{+1.04}_{-0.79}$
$\bar{n} \times 1000$	$1.58^{+0.54}_{-0.43}$	$1.66^{+0.54}_{-0.34}$	$1.16^{+0.33}_{-0.28}$	$1.72^{+0.35}_{-0.47}$
f_{sat}	$0.07^{+0.1}_{-0.05}$	$0.08^{+0.14}_{-0.06}$	$0.06^{+0.09}_{-0.04}$	$0.06^{+0.09}_{-0.04}$
$\log \bar{M}_h$	$13.11^{+0.07}_{-0.04}$	$13.09^{+0.04}_{-0.03}$	$13.05^{+0.05}_{-0.03}$	$12.92^{+0.06}_{-0.03}$

Table III presents the redshift evolution of the HOD fits for the extended LRG sample. We observe a mild decrease in the mean halo mass $\langle M_{\text{halo}} \rangle$ with increasing redshift, consistent with the expectation that halos hosting galaxies of fixed stellar mass are less massive at earlier cosmic times. The satellite fraction and number density also exhibit modest variations across bins, reflecting changes in sample completeness and evolution in the underlying halo population. Nonetheless, the general HOD shape remains stable, supporting the robustness of our modeling framework across redshift.

V. SUMMARY AND CONCLUSIONS

In this work, we measure the gravitational lensing convergence profile, $\kappa(\theta)$, around photometric galaxies in DESI using the CMB lensing maps from ACT DR6 and compare it with existing kSZ measurements. We focus on three DESI tracer samples: the main luminous red galaxies (LRGs), the extended LRGs, and the bright galaxy sample (BGS), and interpret the measurements through a simulation-based halo model, allowing us to infer key halo properties such as mean halo mass and satellite fraction. Our joint kSZ and CMB lensing analysis allows us to test whether we recover the full baryonic content predicted by cosmological models and hydrodynamical simulations. Moreover, because CMB lensing provides a direct estimate of halo mass, our results offer new insights into the stellar-to-halo mass relation (SHMR) and the mass evolution of DESI-selected galaxies. Unlike traditional GGL, our approach is free from many of the noise and systematics

issues that limit optical lensing, making it a powerful tool for future cosmological analyses.

We develop a simulation-based modeling framework, which we show to be fast, accurate and flexible. Overall, we find that our model provides an excellent fit to the $\kappa(\theta)$ measurements for both main and extended LRGs. The fit for the BGS sample is slightly worse likely due to the more complex halo occupation distribution (HOD) of the BGS, which includes a broader mix of galaxy types, including star-forming galaxies that are not as well described by the vanilla HOD model used here [e.g., [79,80]]. We speculate that incorporating assembly bias and allowing for central and satellite galaxy biasing could further improve the fit, though this is beyond the scope of this work.

From the fits, we extract constraints on the five standard HOD parameters and three derived quantities (mean halo mass, number density, and satellite fraction). We find that the mean halo masses are highest for the main sample, followed by the extended and BGS samples, in line with expectations based on galaxy formation and evolution models. Satellite fractions and number densities follow the opposite trend, being highest for the BGS and lowest for the main sample, again consistent with expectations.

A central result of this work is the comparison of our $\kappa(\theta)$ -derived dark matter density profiles with gas density profiles inferred from the kSZ effect (from Ref. [13]). By converting the κ measurements into CAP density profiles in the same units as T_{kSZ} , we perform a nearly assumption-free comparison between the gas and dark matter density. We find excellent agreement between the two profiles at large radii ($\sim 2\text{--}3 \text{ Mpc}/h$), consistent with the notion that gas and dark matter follow each other on large scales. However, at small apertures, the kSZ-inferred profiles lie below the lensing-based profiles, revealing the impact of baryonic feedback, which pushes gas out of the inner regions of halos. Both for the main LRGs as well as the extended LRGs, the baryon and matter curves are significantly discrepant. Note that this analysis is an improvement over Ref. [13] since the amplitude no longer needs adjusting. Using the SNR definition from Ref. [13], this discrepancy between gas and dark matter translates to $\text{SNR} = \sqrt{\Delta\chi^2} \approx 8.2$ (c.f. $\text{SNR} \approx 6$ when fitting for a free amplitude parameter).

Additionally, we study the gas fractions inferred by the ratio of the kSZ and CMB lensing profiles. We find a $\gtrsim 4\sigma$

discrepancy with the gas fractions of LRG-like galaxies in the hydrodynamical simulation TNG300 when comparing with our gas fraction measurements using kSZ and CMB lensing. Interestingly, the low gas fractions we find of about 0.3 for the LRG sample are in agreement with group-size gas fraction measurements using x-ray data. Upon studying the gas fraction of our highest-mass sample, we find that it is higher, as expected theoretically, and consistent with the gas fraction inferred from x-ray analyses of halos of similar mass. In accordance with previous findings, this suggests that large baryonic feedback is necessary to match the observed gas distribution.

We explore the redshift evolution of the halo properties by performing the same analysis in four redshift bins for the main and extended LRG samples. We observe mild redshift trends in the halo masses, particularly for the extended sample, where the inferred halo mass decreases at higher redshift. These trends are consistent with expectations from the evolving lensing kernel and halo mass function.

Finally, we compare the SHMR from the data against the hydrodynamical simulation TNG300. We find good agreement in the halo mass estimates across the four stellar mass bins (all galaxies, $\log(M_*/M_\odot) > 11, 11.25, 11.5$). However, TNG300 tends to overpredict the stellar masses compared to observations, hinting that its galaxies may not be sufficiently quenched. This finding aligns with previous research (e.g., [78]) and highlights the value of combining CMB and galaxy survey data to investigate galaxy formation physics.

Our results demonstrate that joint analyses of κ and kSZ can put tight constraints on the baryonic content and feedback processes by providing estimates of the halo mass. Future work leveraging high-resolution CMB lensing and temperature maps from upcoming experiments such as Simons Observatory and CMB-S4 as well as galaxy survey data from DESI, *Euclid* and the Vera Rubin Observatory (LSST) will allow even more precise tests of baryonic and dark matter physics on small scales.

ACKNOWLEDGMENTS

We thank Ant3n Baleato Lizancos, Esra Bulbul, Matt Johnson, Henry Liu, Luisa Lucie-Smith, Mathew Madhavacheril, Ian McCarthy, Daisuke Nagai, Hiranya Peiris, Andrew Pontzen, Frank Qu, Bernardita Ried Guachalla, Emmanuel Schaan, and Uroš Seljak for useful

discussions, as well as our anonymous referee. B. H. thanks the Miller Institute for supporting her postdoctoral research. S.F., G.S.F., N.S., and R.Z. are supported by Lawrence Berkeley National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

DATA AVAILABILITY

The data that support the findings of this article are openly available [29].

APPENDIX: 2D VISUALIZATION OF κ AND KSZ STACKS

In this appendix, we present 2D stacked images of the CMB lensing convergence κ and the kinematic Sunyaev-Zel'dovich (kSZ) effect for the main LRG sample to provide a visual comparison of the spatial distributions of dark matter and gas.

The top panel of Fig. 12 shows the 2D κ stacks on ACT DR6 data for the main LRG sample, arranged as 2×2 subfigures, each one corresponding to a different redshift bin. To account for the small-scale mode removal ($L < 3000$) applied in the ACT lensing quadratic estimator, we apply a correction derived from the *AbacusSummit* simulations. Specifically, we compute a transfer function by taking the ratio of simulated κ maps with and without the $L > 3000$ cutoff (both convolved with the ACT 1.6' FWHM beam), matched to halos of mass $M_{\text{vir}} \sim 10^{13.2} M_{\odot}/h$ (i.e., the average halo mass for the main sample). This transfer function is then applied multiplicatively to the observed κ stacks to approximately restore large-scale power. The most visible smearing that we observe is thus due to the ACT beam. We check in simulations that this procedure works well; however, small residual effects may remain. We emphasize that these visualizations are intended for qualitative illustration only and are not used for any quantitative analysis.

The bottom panel shows the corresponding 2D kSZ stacks for the same sample, reproduced from Hadzhiyska *et al.* [13]. The kSZ signal is subdominant relative to the total CMB anisotropies, so for visualization purposes only, we apply a high-pass filter governed by the ratio of the kSZ power spectrum to the total power. Inevitably, that suppresses some of the large-scale gas density signal in the 2D kSZ stacks. Nonetheless, it remains evident that the kSZ profiles appear more extended and diffuse compared to the κ stacks.

This visual comparison highlights the core physical result of this work: baryonic feedback processes displace gas from the central regions of halos, leading to a more extended gas distribution relative to the underlying dark matter. The difference in spatial extent between the κ and kSZ stacks is consistent with this interpretation.




FIG. 12. Top: 2D stacked images of ACT DR6 CMB lensing convergence κ for the main LRG sample. A transfer function correction, derived from *AbacusSummit* simulations, has been applied to approximately restore the shape of the small-scale structure removed by the $L > 3000$ cutoff in the ACT lensing pipeline, using halos of mass matching that of the main sample, $M_{\text{vir}} \sim 10^{13.2} M_{\odot}/h$. This correction is used here for visualization only. Bottom: corresponding 2D stacks of the kSZ signal for the same galaxy sample, reproduced from Hadzhiyska *et al.* [13]. Due to the faintness of the kSZ effect, the stacks are effectively high-pass filtered by the kSZ-to-CMB power spectrum ratio, suppressing large-scale modes. The broader appearance of the kSZ signal compared to κ reflects the more extended nature of the gas, consistent with expectations from baryonic feedback processes.

- [1] M. P. van Daalen, J. Schaye, C. M. Booth, and C. Dalla Vecchia, The effects of galaxy formation on the matter power spectrum: A challenge for precision cosmology, *Mon. Not. R. Astron. Soc.* **415**, 3649 (2011).
- [2] 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).
- [3] M. Fukugita and P. J. E. Peebles, The cosmic energy inventory, *Astrophys. J.* **616**, 643 (2004).
- [4] R. Cen and J. P. Ostriker, Where are the baryons? II. Feedback effects, *Astrophys. J.* **650**, 560 (2006).
- [5] A. Schneider, R. Teyssier, J. Stadel, N. E. Chisari, A. M. C. Le Brun, A. Amara, and A. Refregier, Quantifying baryon effects on the matter power spectrum and the weak lensing shear correlation, *J. Cosmol. Astropart. Phys.* **03** (2019) 020.
- [6] L. Bigwood *et al.*, Weak lensing combined with the kinetic Sunyaev Zel'dovich effect: A study of baryonic feedback, [arXiv:2404.06098](https://arxiv.org/abs/2404.06098).
- [7] A. Schneider, R. Teyssier, J. Stadel, N. E. Chisari, A. M. C. Le Brun, A. Amara, and A. Refregier, Quantifying baryon effects on the matter power spectrum and the weak lensing shear correlation, *J. Cosmol. Astropart. Phys.* **03** (2019) 020.
- [8] A. Amon and G. Efstathiou, A non-linear solution to the S_8 tension?, *Mon. Not. R. Astron. Soc.* **516**, 5355 (2022).
- [9] W. Elbers, C. S. Frenk, A. Jenkins, B. Li, J. C. Helly, R. Kugel, M. Schaller, J. Schaye, J. Braspennig, J. Kwan, I. G. McCarthy, J. Salcido, M. P. van Daalen, B. Vandenbroucke, and S. Pascoli, The FLAMINGO project: The coupling between baryonic feedback and cosmology in light of the S_8 tension, *Mon. Not. R. Astron. Soc.* **537**, 2160 (2025).
- [10] H.-J. Huang, T. Eifler, R. Mandelbaum, and S. Dodelson, Modelling baryonic physics in future weak lensing surveys, *Mon. Not. R. Astron. Soc.* **488**, 1652 (2019).
- [11] T. Mroczkowski, D. Nagai, K. Basu, J. Chluba, J. Sayers, R. Adam, E. Churazov, A. Crites, L. Di Mascio, D. Eckert, J. Macias-Perez, F. Mayet, L. Perotto, E. Pointecouteau, C. Romero, F. Ruppin, 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).
- [12] 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).
- [13] 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).
- [14] 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).
- [15] B. Ried Guachalla *et al.*, Backlighting extended gas halos around luminous red galaxies: Kinematic Sunyaev-Zel'dovich effect from DESI Y1 x ACT, *Phys. Rev. D* **112**, 103512 (2025).
- [16] E. Moser, S. Amodeo, N. Battaglia, M. A. Alvarez, S. Ferraro, and E. Schaan, The impacts of modeling choices on the inference of circumgalactic medium properties from Sunyaev-Zeldovich observations, *Astrophys. J.* **919**, 2 (2021).
- [17] S. Amodeo *et al.*, Atacama Cosmology Telescope: Modeling the gas thermodynamics in BOSS CMASS galaxies from kinematic and thermal Sunyaev-Zel'dovich measurements, *Phys. Rev. D* **103**, 063514 (2021).
- [18] E. J. Baxter *et al.*, A measurement of gravitational lensing of the cosmic microwave background by galaxy clusters using data from the south pole telescope, *Astrophys. J.* **806**, 247 (2015).
- [19] M. S. Madhavacheril *et al.*, The Atacama Cosmology Telescope: Weighing distant clusters with the most ancient light, *Astrophys. J. Lett.* **903**, L13 (2020).
- [20] 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).
- [21] 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).
- [22] 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).
- [23] B. Abareschi *et al.* (DESI Collaboration), Overview of the instrumentation for the dark energy spectroscopic instrument, *Astron. J.* **164**, 207 (2022).
- [24] A. Aghamousa *et al.* (DESI Collaboration), The DESI experiment part II: Instrument design, [arXiv:1611.00037](https://arxiv.org/abs/1611.00037).
- [25] J. H. Silber *et al.* (DESI Collaboration), The robotic multi-object focal plane system of the dark energy spectroscopic instrument (DESI), *Astron. J.* **165**, 9 (2023).
- [26] T. N. Miller *et al.*, The optical corrector for the dark energy spectroscopic instrument, *Astron. J.* **168**, 95 (2024).
- [27] M. Levi, C. Bebek, T. Beers, R. Blum, R. Cahn, D. Eisenstein, B. Flaugher, K. Honscheid, R. Kron, O. Lahav, P. McDonald, N. Roe, and D. Schlegel (representing the DESI Collaboration), The DESI experiment, a white-paper for snowmass 2013, [arXiv:1308.0847](https://arxiv.org/abs/1308.0847).
- [28] A. Aghamousa *et al.* (DESI Collaboration), The DESI experiment part I: Science, targeting, and survey design, [arXiv:1611.00036](https://arxiv.org/abs/1611.00036).
- [29] R. Zhou *et al.*, Target selection and validation of DESI luminous red galaxies, *Astron. J.* **165**, 58 (2023).
- [30] R. Zhou *et al.*, DESI luminous red galaxy samples for cross-correlations, *J. Cosmol. Astropart. Phys.* **11** (2023) 097.
- [31] C. Hahn *et al.*, The DESI bright galaxy survey: Final target selection, design, and validation, *Astron. J.* **165**, 253 (2023).

- [32] S. Chen *et al.*, Analysis of DESI $\times$ DES using the Lagrangian effective theory of LSS, *Phys. Rev. D* **110**, 103518 (2024).
- [33] M. S. Madhavacheril *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: DR6 gravitational lensing map and cosmological parameters, *Astrophys. J.* **962**, 113 (2024).
- [34] F. J. Qu *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: A measurement of the DR6 CMB lensing power spectrum and its implications for structure growth, *Astrophys. J.* **962**, 112 (2024).
- [35] N. MacCrann *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: Mitigating the impact of extragalactic foregrounds for the DR6 cosmic microwave background lensing analysis, *Astrophys. J.* **966**, 138 (2024).
- [36] G. S. Farren *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: Cosmology from cross-correlations of unWISE Galaxies and ACT DR6 CMB lensing, *Astrophys. J.* **966**, 157 (2024).
- [37] N. Sailer, E. Schaan, and S. Ferraro, Lower bias, lower noise CMB lensing with foreground-hardened estimators, *Phys. Rev. D* **102**, 063517 (2020).
- [38] <https://github.com/EmmanuelSchaan/ThumbStack>
- [39] M. White, Reconstruction within the Zeldovich approximation, *Mon. Not. R. Astron. Soc.* **450**, 3822 (2015).
- [40] B. Hadzhiyska, S. Ferraro, B. Ried Guachalla, and E. Schaan, Velocity reconstruction in the era of DESI and Rubin/LSST. II. Realistic samples on the light cone, *Phys. Rev. D* **109**, 103534 (2024).
- [41] B. Ried Guachalla, E. Schaan, B. Hadzhiyska, and S. Ferraro, Velocity reconstruction in the era of DESI and Rubin/LSST. I. Exploring spectroscopic, photometric, and hybrid samples, *Phys. Rev. D* **109**, 103533 (2024).
- [42] 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).
- [43] B. Hadzhiyska, L. H. Garrison, D. Eisenstein, and S. Bose, The halo light-cone catalogues of AbacusSummit, *Mon. Not. R. Astron. Soc.* **509**, 2194 (2022).
- [44] L. H. Garrison, D. J. Eisenstein, and P. A. Pinto, A high-fidelity realization of the Euclid code comparison N -body simulation with Abacus, *Mon. Not. R. Astron. Soc.* **485**, 3370 (2019).
- [45] L. H. Garrison, D. J. Eisenstein, D. Ferrer, N. A. Maksimova, and P. A. Pinto, The Abacus cosmological N -body code, *Mon. Not. R. Astron. Soc.* **508**, 575 (2021).
- [46] B. Hadzhiyska, D. Eisenstein, S. Bose, L. H. Garrison, and N. Maksimova, COMPASO: A new halo finder for competitive assignment to spherical overdensities, *Mon. Not. R. Astron. Soc.* **509**, 501 (2022).
- [47] G. L. Bryan and M. L. Norman, Statistical properties of x-ray clusters: Analytic and numerical comparisons, *Astrophys. J.* **495**, 80 (1998).
- [48] Z. Zheng, A. A. Berlind, D. H. Weinberg, A. J. Benson, C. M. Baugh, S. Cole, R. Dave, C. S. Frenk, N. Katz, and C. G. Lacey, Theoretical models of the halo occupation distribution: Separating central and satellite galaxies, *Astrophys. J.* **633**, 791 (2005).
- [49] <https://github.com/abacusorg/abacusutils>
- [50] <https://camb.readthedocs.io/en/latest/>
- [51] B. Hadzhiyska *et al.*, Synthetic light-cone catalogues of modern redshift and weak lensing surveys with AbacusSummit, *Mon. Not. R. Astron. Soc.* **525**, 4367 (2023).
- [52] <https://github.com/simonsobs/pixell>
- [53] <https://github.com/simonsobs/so-lenspipe>
- [54] W. Hu, S. DeDeo, and C. Vale, Cluster mass estimators from CMB temperature and polarization lensing, *New J. Phys.* **9**, 441 (2007).
- [55] M. S. Madhavacheril *et al.* (ACT Collaboration), The Atacama Cosmology Telescope: Weighing distant clusters with the most ancient light, *Astrophys. J. Lett.* **903**, L13 (2020).
- [56] E. J. Baxter *et al.*, A measurement of gravitational lensing of the cosmic microwave background by galaxy clusters using data from the south pole telescope, *Astrophys. J.* **806**, 247 (2015).
- [57] S. Raghunathan, S. Patil, E. J. Baxter, F. Bianchini, L. E. Bleem, T. M. Crawford, G. P. Holder, A. Manzotti, and C. L. Reichardt, Measuring galaxy cluster masses with CMB lensing using a Maximum Likelihood estimator: Statistical and systematic error budgets for future experiments, *J. Cosmol. Astropart. Phys.* **08** (2017) 030.
- [58] B. Horowitz, S. Ferraro, and B. D. Sherwin, Reconstructing small scale lenses from the cosmic microwave background temperature fluctuations, *Mon. Not. R. Astron. Soc.* **485**, 3919 (2019).
- [59] K. Levy, S. Raghunathan, and K. Basu, A foreground-immune CMB-cluster lensing estimator, *J. Cosmol. Astropart. Phys.* **08** (2023) 020.
- [60] G. Stein, M. A. Alvarez, J. R. Bond, A. van Engelen, and N. Battaglia, The websky extragalactic CMB simulations, *J. Cosmol. Astropart. Phys.* **10** (2020) 012.
- [61] E. Schaan and S. Ferraro, Foreground-immune cosmic microwave background lensing with shear-only reconstruction, *Phys. Rev. Lett.* **122**, 181301 (2019).
- [62] F. J. Qu, A. Challinor, and B. D. Sherwin, CMB lensing with shear-only reconstruction on the full sky, *Phys. Rev. D* **108**, 063518 (2023).
- [63] M. S. Madhavacheril and J. C. Hill, Mitigating foreground biases in CMB lensing reconstruction using cleaned gradients, *Phys. Rev. D* **98**, 023534 (2018).
- [64] N. Sailer, S. Ferraro, and E. Schaan, Foreground-immune CMB lensing reconstruction with polarization, *Phys. Rev. D* **107**, 023504 (2023).
- [65] O. Darwish, B. D. Sherwin, N. Sailer, E. Schaan, and S. Ferraro, Optimizing foreground mitigation for CMB lensing with combined multifrequency and geometric methods, *Phys. Rev. D* **107**, 043519 (2023).
- [66] N. Sailer, E. Schaan, S. Ferraro, O. Darwish, and B. Sherwin, Optimal multifrequency weighting for CMB lensing, *Phys. Rev. D* **104**, 123514 (2021).
- [67] W. R. Coulton *et al.*, The Atacama Cosmology Telescope: Detection of patchy screening of the cosmic microwave background, *arXiv:2401.13033*.
- [68] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [69] E. Chung, S. Foreman, and A. van Engelen, Baryonic effects on CMB lensing and neutrino mass constraints,

- Phys. Rev. D **101**, 063534 (2020); Phys. Rev. D **102**, 109903(E) (2020).
- [70] To avoid potential bias from upward fluctuations, we do not select the bin with the highest SNR in the data, but instead use simulations. We also note that the first radial bin in the κ measurements correlates best with the one-halo term amplitude, as it is sensitive to smaller scales.
- [71] 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).
- [72] P. Popesso *et al.*, The hot gas mass fraction in halos. From Milky Way-like groups to massive clusters, [arXiv:2411.16555](https://arxiv.org/abs/2411.16555).
- [73] M. Ayromlou, D. Nelson, and A. Pillepich, Feedback reshapes the baryon distribution within haloes, in halo outskirts, and beyond: The closure radius from dwarfs to massive clusters, *Mon. Not. R. Astron. Soc.* **524**, 5391 (2023).
- [74] L. Lucie-Smith, H. V. Peiris, A. Pontzen, A. Halder, J. Schaye, M. Schaller, J. Helly, R. J. McGibbon, and W. Elbers, Cosmological feedback from a halo assembly perspective, *Phys. Rev. D* **112**, 063541 (2025).
- [75] M. Kovač, A. Nicola, J. Bucko, A. Schneider, R. Reischke, S. K. Giri, R. Teyssier, M. Schaller, and J. Schaye, Baryonification II: Constraining feedback with x-ray and kinematic Sunyaev-Zel'dovich observations, *J. Cosmol. Astropart. Phys.* **11** (2025) 046.
- [76] D. Nelson, V. Springel, A. Pillepich, V. Rodriguez-Gomez, P. Torrey, S. Genel, M. Vogelsberger, R. Pakmor, F. Marinacci, R. Weinberger, L. Kelley, M. Lovell, B. Diemer, and L. Hernquist, The IllustrisTNG simulations: Public data release, *Comput. Astrophys. Cosmol.* **6**, 2 (2019).
- [77] 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).
- [78] A. Pillepich *et al.*, First results from the IllustrisTNG simulations: The stellar mass content of groups and clusters of galaxies, *Mon. Not. R. Astron. Soc.* **475**, 648 (2018).
- [79] B. Hadzhiyska, S. Tacchella, S. Bose, and D. J. Eisenstein, The galaxy-halo connection of emission-line galaxies in IllustrisTNG, *Mon. Not. R. Astron. Soc.* **502**, 3599 (2021).
- [80] B. Hadzhiyska, L. Hernquist, D. Eisenstein, A. M. Delgado, S. Bose, R. Kannan, R. Pakmor, V. Springel, S. Contreras, M. Barrera, F. Ferlito, C. Hernández-Aguayo, S. D. M. White, and C. Frenk, The MillenniumTNG Project: Refining the one-halo model of red and blue galaxies at different redshifts, *Mon. Not. R. Astron. Soc.* **524**, 2524 (2023).