

The Atacama Cosmology Telescope: Reionization kSZ trispectrum methodology and limits

Niall MacCrann,^{1,2}★ Frank J. Qu,^{1,2} Toshiya Namikawa,³ Boris Bolliet,^{1,2,4} Hongbo Cai,⁵ Erminia Calabrese,⁶ Steve K. Choi,⁷ William Coulton,^{1,2} Omar Darwish,⁸ Simone Ferraro,⁹ Yilun Guan,¹⁰ J. Colin Hill,¹¹ Matt Hilton,^{12,13} Renée Hložek,^{14,15} Darby Kramer,¹⁶ Mathew S. Madhavacheril,¹⁷ Kavilan Moodley,^{18,19} Neelima Sehgal,²⁰ Blake D. Sherwin,^{1,2} Cristóbal Sifón,²¹ Suzanne T. Staggs,²² Hy Trac,²³ Alexander Van Engelen¹⁶ and Eve M. Vavagiakis²⁴

Affiliations are listed at the end of the paper

Accepted 2024 July 11. Received 2024 July 11; in original form 2024 May 30

ABSTRACT

Patchy reionization generates kinematic Sunyaev–Zel’dovich (kSZ) anisotropies in the cosmic microwave background (CMB). Large-scale velocity perturbations along the line of sight modulate the small-scale kSZ power spectrum, leading to a trispectrum (or four-point function) in the CMB that depends on the physics of reionization. We investigate the challenges in detecting this trispectrum and use tools developed for CMB lensing, such as realization-dependent bias subtraction and cross-correlation based estimators, to counter uncertainties in the instrumental noise and assumed CMB power spectrum. We also find that both lensing and extragalactic foregrounds can impart larger trispectrum contributions than the reionization kSZ signal. We present a range of mitigation methods for both of these sources of contamination, validated on microwave-sky simulations. We use ACT DR6 and *Planck* data to calculate an upper limit on the reionization kSZ trispectrum from a measurement dominated by foregrounds. The upper limit is about 50 times the signal predicted from recent simulations.

Key words: cosmology: observations – cosmic background radiation – dark ages, reionization, first stars – large-scale structure of Universe.

1 INTRODUCTION

The epoch of reionization (EoR) is the period during which the Universe’s first stars and galaxies formed, and the resulting radiation began the process of ionizing the Universe, eventually leading to a Universe populated with the low-density, ionized hydrogen we observe today. The EoR is a challenging period to observe, as the high-redshift ionizing sources are extremely faint, although recent data from the *JWST* presents a dramatic increase in available data on the formation of these first galaxies (e.g. Eisenstein et al. 2023; Finkelstein et al. 2023a,b). While current and upcoming 21 cm experiments, such as the Hydrogen Epoch of Reionization Array (HERA)¹ and the Square Kilometre Array (SKA)², aim to map out the neutral hydrogen abundance during reionization, they will need to address foregrounds that are orders of magnitude above the signal.

Reionization also leaves imprints on the cosmic microwave background (CMB). As well as a large-scale polarization signature (e.g. Kogut 2003), there is a contribution to the kinematic Sunyaev–Zeldovich (kSZ) signal from reionization, commonly referred to as the ‘patchy’ kSZ effect (e.g. Santos et al. 2003, McQuinn et al. 2005, Trac, Bode & Ostriker 2011, Shaw, Rudd & Nagai 2012, Battaglia

et al. 2013). During reionization, there are large inhomogeneities in the free electron density as ‘bubbles’ surrounding ionizing sources ionize before other regions. These bubbles will have a range of radial peculiar velocities with respect to the observer, leading to a scattering and Doppler shifting of the CMB photons that corresponds to a boosting (for bubbles moving towards us) or reduction (for bubbles moving away) in the observed CMB temperature. This generates a contribution to the kSZ power spectrum of the same order of magnitude as the low-redshift effect, which arises due to radially moving inhomogeneities in the (now fully ionized) electron density field.

The kSZ power spectrum dominates over the blackbody primary CMB signal at small scales, $l \gtrsim 4500$. Data from ground-based, high-resolution, low-noise, CMB experiments such as the South Pole Telescope (SPT) and the Atacama Cosmology Telescope (ACT; which we use here) already have greater sensitivity to this small-scale regime than *Planck*. While there have been recent attempts to disentangle this reionization kSZ signal from the different contributions to the small-scale CMB temperature (Planck Collaboration XLVII 2016; Choi et al. 2020; Reichardt et al. 2021; Gorce, Douspis & Salvati 2022) using only the power spectrum, the presence of foregrounds (including the late-time kSZ) has thus far posed significant challenges. These challenges are also discussed in Calabrese et al. (2014) and Ade et al. (2019) who forecast the potential of upcoming experiments

* E-mail: nm746@cam.ac.uk

¹<https://reionization.org/>

²<https://www.skao.int/en>

like Simons Observatory³ to constrain reionization using the kSZ power spectrum.

In this work, we focus on the non-Gaussian signature imparted on the CMB due to the kSZ effect during reionization. As laid out in Smith & Ferraro (2017) and Ferraro & Smith (2018), this non-Gaussianity should be present due to the modulation of the kSZ signal by large-scale radial velocity modes. For a given line of sight $\hat{n}$, the small-scale kSZ power spectrum is amplified depending on the squared radial velocity field, projected along the line of sight, in direction $\hat{n}$, with respect to the sky-averaged projected squared radial velocity. Given the large coherence length of the velocity perturbations, even after the line-of-sight projection there should be significant variation of the (squared) radial velocity across the sky, leading to a position-dependent kSZ power spectrum, i.e. a trispectrum. Both the power spectrum and the kSZ trispectrum include contributions from moving ionized overdensities in the more recent universe ($z \lesssim 3$) as well as from reionization. The trispectrum gives us additional information with which to disentangle the contributions from the two kSZ epochs (Alvarez et al. 2021).

In an idealized forecast,⁴ we forecast a signal-to-noise ratio for the patchy kSZ trispectrum of roughly unity, for the current ACT data set used in this work, when assuming the signal is described by the Alvarez (2016) simulation. As we will describe, complexities in the measurement mean the achieved signal-to-noise ratio is less than this, however, there is significant theoretical uncertainty in the predicted signal, and an upper limit could still rule out more extreme theoretical models.

While this paper was in collaboration review, Raghunathan et al. (2024) presented a first upper limit on the kSZ trispectrum using a combination of SPT and *Herschel*-SPIRE data, finding that the contribution from foregrounds dominates over the kSZ signal. In this work, we also find that foregrounds present a major challenge for measuring the kSZ trispectrum, dominating over the predicted signal for current ACT + *Planck* data. With upcoming, lower noise data however, we expect that residual foregrounds can be greatly reduced (because, e.g. fainter sources and clusters can be detected and mitigated against), and with the mitigation methods described in Section 3, the kSZ trispectrum may become an informative probe of reionization physics.

Estimating an unbiased trispectrum from CMB data which has complex, anisotropic noise, and non-Gaussian foregrounds, has various challenges. In Section 2, we describe our trispectrum estimator and bias corrections. In Section 3, we describe the biases due to lensing and foregrounds, and our mitigation strategies. In Section 4, we estimate foreground biases from microwave-sky simulations. In Section 5, we describe the ACT DR6 + *Planck* data and simulations used, and in Section 6, we describe our measurement. In Section 7, we conclude and discuss future improvements.

2 THEORETICAL BACKGROUND AND TRISPECTRUM ESTIMATOR

In Section 2.1, we briefly outline and motivate the trispectrum estimator we use, and refer the reader to Smith & Ferraro (2017) and Ferraro & Smith (2018) for further details. We also discuss the sensitivity of the trispectrum statistic to the properties of reionization in Section 2.2, and discuss the bias corrections required in Section 2.3.

³<https://simonsobservatory.org/>

⁴Not using the cross-correlation-based estimator (see Section 2.3) and not including foregrounds in the noise (see Section 3.2).

2.1 The estimator

The kSZ effect imparts a fractional temperature fluctuation on the CMB:

$$\Delta\Theta^{\text{kSZ}}(\hat{n}) = -\frac{1}{c} \int d\chi g(\chi) [1 + \delta_e(\chi\hat{n})] v_r(\chi\hat{n}), \quad (1)$$

where χ is the comoving distance, $g(\chi) = e^{-\tau} \frac{d\tau}{d\chi}$ is the visibility function, τ is the optical depth, $\delta_e(\chi\hat{n})$ is the free electron overdensity, and $v_r(\chi\hat{n}) = \hat{n} \cdot \mathbf{v}(\chi\hat{n})$ is the line-of-sight component of the peculiar velocity. As we will describe, in this work we are interested in the small-scale (high- l) temperature fluctuation generated by the kSZ, which depends on small-scale electron density perturbations, hence we drop the constant electron density term in equation (1). In the flat-sky limit, in the Fourier domain, the multiplication becomes a convolution:

$$\Delta\Theta(\mathbf{l}) = -\frac{1}{c} \int d\chi g(\chi) \int \frac{d^2\mathbf{l}'}{(2\pi)^2} \delta_e(\mathbf{l}' - \mathbf{l}, \chi) v_r(\mathbf{l}', \chi), \quad (2)$$

where $v_r(\mathbf{l}, \chi) = \int d\Omega e^{i\mathbf{l} \cdot \hat{n}} \hat{n} \cdot \mathbf{v}(\chi\hat{n})$ and $d\Omega$ is the area element on the sky.

Ignoring for now the integral over the line-of-sight, it is useful to consider the case where the radial velocity field has only a single large scale velocity mode V_L , i.e. $v_r(\mathbf{l}', r) = \delta_D(\mathbf{l}' - \mathbf{L})V_L$. Ignoring other contributions to the temperature, we would then have

$$\Delta\Theta(\mathbf{l}) \propto \int d^2\mathbf{l}' \delta_e(\mathbf{l} - \mathbf{l}') \delta_D(\mathbf{l}' - \mathbf{L}) V_L, \quad (3)$$

$$\propto V_L \delta_e(\mathbf{l} - \mathbf{L}), \quad (4)$$

where $\delta_D(\mathbf{l})$ denotes the Dirac delta function. The product of two temperature modes is then

$$\Delta\Theta(\mathbf{l}_1) \Delta\Theta(\mathbf{l}_2) \propto V_L^2 \delta_e(\mathbf{l}_1 - \mathbf{L}) \delta_e(\mathbf{l}_2 - \mathbf{L}), \quad (5)$$

and has expectation, for fixed V_L

$$\langle \Delta\Theta(\mathbf{l}_1) \Delta\Theta(\mathbf{l}_2) \rangle_{\text{fixed } V_L} \propto V_L^2 \delta_D(\mathbf{l}_1 + \mathbf{l}_2 - 2\mathbf{L}) C_{\delta_e}^{\text{ee}}[\frac{\mathbf{l}_1 - \mathbf{l}_2}{2}], \quad (6)$$

where $C_{\delta_e}^{\text{ee}}$ is the electron density power spectrum. We see then that a product of two temperature modes is an estimator for the radial velocity squared V_L^2 , i.e. the kSZ effect enables us to probe the Universe's velocity field, a rich source of cosmological information, using the CMB.

The physical picture here is that the small-scale power in the CMB temperature due to the kSZ is modulated up or down for a given line of sight $\hat{n}$, depending on the size of $\langle V_L^2 \rangle(\hat{n})$ for that region (where angle brackets imply an averaging along the line of sight). Note that in reality the kSZ effect has a broad redshift kernel, so the quadratic estimator is sensitive to the radial velocity squared projected over this broad redshift kernel. Given the large coherence length of the velocity field, there is still significant variation across the sky in this projected radial velocity squared, and this modulation generates a position-dependent power spectrum. Note that the trispectrum (i.e. four-point function) is the lowest order non-zero statistic of the kSZ temperature perturbations beyond the power spectrum, since the bispectrum (i.e. three-point function) is zero by symmetry.⁵

⁵One can argue this in the following way: Consider the correlation of three radial velocity modes $\langle v_r(k_1) v_r(k_2) v_r(k_3) \rangle$. The configuration with $v_r \rightarrow -v_r$ is equally likely, in which case the three-point function is $-\langle v_r(k_1) v_r(k_2) v_r(k_3) \rangle$. In the ensemble average over possible values of v_r , these two configurations will always cancel, and $\langle v_r(k_1) v_r(k_2) v_r(k_3) \rangle$, averaged over realizations of the v_r field, is zero.

Motivated by this, Smith & Ferraro (2017) formed a quadratic estimator for this non-Gaussianity induced by the kSZ at reionization

$$K_L = N_L^{KK} \int \frac{d^2l}{(2\pi)^2} W_l W_{L-l} T_l T_{L-l}, \quad (7)$$

with the filter function,

$$W_l = \sqrt{C_l^{\text{pkSZ}} / C_l^{\text{tot}}}, \quad (8)$$

chosen to up-weight the kSZ contribution to the temperature modes. Here, C_l^{pkSZ} is the contribution to the temperature power spectrum due to the kSZ at reionization (i.e. the ‘patchy’ term), and C_l^{tot} is the total observed temperature power spectrum (i.e. including noise). Throughout this work, when constructing the filter W_l , we set C_l^{pkSZ} equal to the angular power spectrum measured from the Alvarez (2016) simulations, and use as C_l^{tot} a smoothed measurement of the power spectrum of the (beam-deconvolved) data temperature map (see Section 5 for more details). In the squeezed limit $L \ll l$, K is just a measure of the local amplitude of the (l -weighted) temperature power spectrum, integrated over scales l . One can think of $K(\hat{n})$ then as some measure of the local temperature power spectrum around the direction $\hat{n}$. The power spectrum of K , C_L^{KK} , is a measure of how much it varies across the sky. Note that throughout, we will use L to denote the large scale on which we measure K , while l labels the (small scale) CMB temperature fluctuations which enter our quadratic estimator for K_L .

N_L^{KK} is a normalization factor which we define following Sailer, Schaan & Ferraro (2020):

$$(N_L^{KK})^{-1} = \int \frac{d^2l}{(2\pi)^2} W_l^2 W_{L-l}^2 C_l^{\text{tot}} C_{L-l}^{\text{tot}}. \quad (9)$$

We note that this normalization is fairly arbitrary – so long as measurement and theoretical prediction are normalized consistently, the exact choice is not crucial. None the less, we discuss alternative approaches, including that taken by Raghunathan et al. (2024), in Appendix C.

As a baseline theory model, we measure C_L^{KK} from the Alvarez (2016) simulations, using the same W_l as in the ACT DR6 + Planck data measurement (see Section 6). This model is shown as the black solid line in Fig. 2, where it is scaled by L^2 , hence one can see it has strong support at low L i.e. large angular scales. We use this model as a template to which we compare the data measurement and foreground biases, often reporting a fitted amplitude parameter that simply re-scales this template, which we call A_{A16} i.e.

$$A_{A16} = \frac{\sum_L C_L^{KK} C_L^{KK, A16} / \sigma_L^2}{\sum_L (C_L^{KK, A16} / \sigma_L)^2}, \quad (10)$$

where C_L^{KK} is the data measurement, $C_L^{KK, A16}$ is the Alvarez (2016) model, and σ_L is the data measurement uncertainty.

We note that there is considerable theoretical uncertainty in the C_L^{KK} signal and that the Abundance Matching Box for the Epoch of Reionization (AMBER) simulations (Trac et al. 2022; Chen et al. 2023) provide an updated treatment relative to Alvarez (2016). They also include a range of reionization scenarios. We consider these more extensive simulations in Appendix A.

2.2 Sensitivity to properties of reionization

Smith & Ferraro (2017) and Ferraro & Smith (2018) motivate the following theoretical model for C_L^{KK} :

$$C_L^{KK} \propto \int dz \frac{H(z)}{\chi^2(z)} \left[\frac{dC_l^{\text{pkSZ}}}{dz} \right]^2 P_\eta [L/\chi(z), z], \quad (11)$$

where $\chi(z)$ is the comoving distance to redshift z and

$$\eta(\hat{n}, z) = v_r^2(\hat{n}, z) / \langle v_r^2(z, \hat{n}) \rangle_{\hat{n}} \quad (12)$$

is the squared radial velocity contrast, i.e. the squared radial velocity divided by its sky average. The dependence on the details of reionization enters largely through $\left[\frac{dC_l^{\text{pkSZ}}}{dz} \right]^2$, which is the derivative of the reionization contribution to the kSZ power spectrum as a function of redshift. Thus C_L^{KK} may be increased by physics that increases C_l^{pkSZ} , such as in models where more of the reionization occurs at higher redshift, since the physical electron density is higher (scaling as $(1+z)^3$). The trispectrum signal additionally depends on $P_\eta [L/\chi(z), z]$, the power spectrum of the squared radial velocity contrast. This largely sets the scale-dependence of the C_L^{KK} , with earlier reionization corresponding to the signal peaking on smaller (angular) scales.

As discussed in Smith & Ferraro (2017), Ferraro & Smith (2018), Alvarez et al. (2021), and Raghunathan et al. (2024) the kSZ trispectrum signal is mainly sensitive to the midpoint, z_{mid} and width, Δz of reionization. In particular, when more of the reionization occurs at higher redshift (which could result from larger z_{mid} or Δz), the kSZ power spectrum tends to be increased since ionized bubbles of a given size have higher mean density, and thus higher optical depth. Since the trispectrum signal arises from the modulation of this power spectrum, it also increases. For the trispectrum signal we note that there should be a competing effect whereby larger Δz means the modulating radial velocity field is averaged over a longer line-of-sight and therefore should have smaller fluctuations. Hence, we do expect a somewhat different sensitivity to changes in the reionization parameters for the power spectrum and the trispectrum, allowing for breaking of the $\tau - \Delta z$ degeneracy demonstrated in Alvarez et al. (2021).

The complementary information provided by the trispectrum is likely even more important when realistic foregrounds are present, making interpretation of either statistic alone more difficult. In Appendix A, we study the response of the reionization kSZ power spectrum and trispectrum to changes in the reionization parameters.

2.3 N^0 bias, mean-field, and the cross-correlation estimator

There is a bias to C_L^{KK} arising from the disconnected trispectrum, i.e. the trispectrum one would measure for a Gaussian field with the same power spectrum. We follow the CMB lensing convention here and denote this bias as N_L^0 .

As in the case of CMB lensing, this correction depends on the total power spectrum of the input temperature map (i.e. including CMB, foregrounds, and instrumental noise), and inaccuracy in the assumed total power spectrum will lead to bias in the calculated N^0 correction, thus biasing the inferred C_L^{KK} . Hence, a calculation where the estimator is simply applied to many Gaussian simulations with representative power, and averaged (an approach we will refer to as MCN^0), is susceptible to biases in that assumed representative power spectrum. As we will see, the N^0 bias can be orders of magnitude larger than the kSZ signal we are trying to recover here, so even a small fractional bias in the N^0 can cause biases in C_L^{KK} .

In this work, we propose to follow the approach taken in CMB lensing and use a realization-dependent N^0 (RDN^0 ; Hanson et al. 2011; Namikawa, Hanson & Takahashi 2013).⁶ This method removes

⁶While the MCN^0 method described above can be denoted as an averaging of the tripsectrum estimator over pairs of random, Gaussian, simulated

biases in the N^0 that depend linearly on the fractional difference, $\Delta C_l/C_l^{\text{tot}}$, between the true power spectrum and that assumed for the simulations. With this method remaining biases are then $\mathcal{O}[(\Delta C_l/C_l^{\text{tot}})^2]$ or higher, thus a sub-percent accuracy N^0 can be achieved even with few percent uncertainty in the total power spectrum. We show in Section 6 that this method accounts for small errors in the assumed temperature power spectrum. We believe this is especially important given the high l s considered for this measurement ($l \gtrsim 3000$), where highly uncertain foregrounds and the kSZ signal itself dominate over the better understood primary CMB power spectrum.

There is another bias to K_L , and hence C_L^{KK} , arising from statistical anisotropy of the noise or survey mask, which can also generate an observed position-dependent power spectrum or trispectrum. We again follow the CMB literature by calling this the *mean-field*, $\bar{K}_L$. It is calculated as the mean of the K_L estimator applied to simulations which have the same statistical anisotropy due to noise and mask as the data, as well as the same underlying power spectrum.

We also follow the lead of CMB lensing approaches by using the cross-correlation-only trispectrum estimator of Madhavacheril et al. (2021), as described in Section 6, which makes both the N^0 bias and the mean-field independent of the instrumental noise (which would otherwise be difficult to model accurately). We note that the cross-correlation estimator results in a reduction of the signal-to-noise ratio (S/N) of the measurement, since it removes the autocorrelations of the independent data splits. This reduction is significant in the high l regime used here, where noise that is uncorrelated between data splits becomes dominant over signal. Madhavacheril et al. (2021) show that using a greater number of splits reduces the S/N cost of the cross-correlation estimator. In this work, we use four splits, which results in roughly a factor of 2 to 3 increase of our measurement uncertainties. As we will demonstrate in Section 6, given that our upper-limit is dominated by foreground contamination rather than statistical uncertainty, we think the extra robustness provided by the cross-correlation estimator is worth the S/N cost.

The N^0 and mean-field biases are thus accounted for in our estimator for C_L^{KK} :

$$C_L^{KK} = [\langle (\hat{K}_{AB} - \bar{K}_{AB})_L (\hat{K}_{CD}^* - \bar{K}_{CD}^*)_L \rangle - \text{RDN}_L^{0,ABCD}] / (w_4 T_L^{KK}). \quad (15)$$

In this expression the angled brackets, $\langle XY \rangle$ denote the angular power spectrum between fields X and Y . $\hat{K}_L^{AB}$ is the ‘raw’ (i.e. before mean-field correction) data measurement of the K_L field in equation (7), computed from two temperature maps labelled A and B . These two temperature maps may be the same or different (for example if different foreground cleaning measures have been applied to each, see Section 3.2 for discussion). $\bar{K}_L^{AB}$ is mean-field and $\text{RDN}_L^{0,ABCD}$ is the realization-dependent N^0 correction, both described above.

temperature maps S and S' :

$$\text{MCN}_L^0 \sim 2 \langle \langle K_L(S, S') K_L(S, S') \rangle \rangle_{S, S'}, \quad (13)$$

the RDN^0 additionally uses combinations of data and simulations to compute the Gaussian bias:

$$\text{RDN}_L^0 \sim 2 \langle \langle K_L(D, S) K_L(D, S) \rangle - \langle K_L(S, S') K_L(S, S') \rangle \rangle_{S, S'}. \quad (14)$$

In these expressions, $K_L(X, Y)$ denote the K_L estimate from two temperature maps X and Y , the inner angle brackets is simply the angular power spectrum of two K_L estimates, and $\langle \rangle_{S, S'}$ denotes averaging over pairs of random simulations.

Finally, we also include multiplicative factors w_4^7 and T_L^{KK} (detailed in Section 5.1) to account for the survey mask and survey transfer function. These final corrections mean our estimator is for the full-sky C_L^{KK} .

It is this estimator that we apply to ACT DR6 and *Planck* in Section 6, where we also demonstrate the importance of using the cross-correlation-based estimator, and RDN^0 , for our measurement. First though, we turn to the significant contamination of the C_L^{KK} signal by extragalactic foregrounds and CMB lensing.

3 NEW METHODS FOR MITIGATION OF LENSING AND FOREGROUND BIASES

Our trispectrum statistic, C_L^{KK} is susceptible to contamination from other sources of non-Gaussianity in the observed CMB, in particular, gravitational lensing and extragalactic foregrounds. In Section 4, we estimate biases to C_L^{KK} using microwave sky simulations. In this section, we introduce and motivate methods to mitigate against these biases: lensing-hardening and frequency-based methods. Note also that the use of RDN^0 , as proposed in the previous section, will also mitigate the impact of foreground uncertainty via reducing the sensitivity to the assumed foreground contribution to the assumed power spectrum.

3.1 Lensing and bias-hardening

Lensing of the CMB by a lensing potential field ϕ_L generates a mode-coupling

$$\langle \Theta(l) \Theta(L-l) \rangle = f_{l,L}^\phi \phi_L, \quad (16)$$

where

$$f_{l,L}^\phi = C_l l \cdot L + C_{|L-l|} l \cdot (L-l), \quad (17)$$

and C_l is the CMB power spectrum without noise.

Lensing therefore produces a bias to our K estimate, K_L^ϕ , which has expectation

$$K_L^\phi = N_L^K \phi_L \int \frac{d^2 l}{(2\pi)^2} W_l W_{L-l} f_{l,L}^\phi. \quad (18)$$

As we will show below, this signal is larger than the expected signal from the kSZ, but can be mitigated via bias-hardening, a method developed for CMB lensing (Namikawa et al. 2013; Osborne, Hanson & Doré 2014) that aims to isolate the lensing potential from other sources of mode-coupling in the reconstruction, such as Poisson distributed point-sources or anisotropy due to the mask. When the functional form of mode-coupling is given, e.g. equation (17) for lensing, one can write down a response of some other quadratic estimator to that source of mode-coupling. Our case is most similar to that investigated in Sailer et al. (2020), which used a bias-hardened estimator to remove the contamination to lensing due to Poisson-distributed extended sources. Noting the similarity of our K estimator with the extended source estimator of Sailer et al. (2020), here we can instead use bias-hardening to remove the contamination due to lensing from our K estimator.

⁷ $w_4 = \sum_i a_i (m^i)^4 / \sum_i a_i$ where the sum is over pixels i , a_i is the area of pixel i , and m^i is the value of the (apodized) mask (between 0 and 1) for pixel i .




Figure 1. Expected fractional foreground biases from the WEBSKY (blue lines) and S10 (orange lines) simulations, for our baseline CMB temperature l range $3000 < l < 4000$. Note the y-axis is linear in the range $[-100, 100]$, and log-scaled otherwise. Solid lines use a minimum variance ILC to combine frequency channels. For dashed lines, one of the temperature maps in the trispectrum has tSZ and CIB spectra de-projected (see Section 3.2). In the legend, we quote ΔA_{A16} which is the bias in the amplitude fitted to the C_L^{KK} prediction from the Alvarez (2016) simulations. We also show the bias due to lensing multiplied by 10, without lensing mitigation (black solid line), when lensing-hardening (‘LH’) is used, (dashed black line, see Section 3.1) and when an additional N^1 correction is used (dotted black line).

In general a bias-hardened estimator for a field x_L in the presence of a contaminant y_L is

$$x_L = \frac{\hat{x}_L - N_L^x R_L^{xy} \hat{y}_L}{1 - N_L^x N_L^y (R_L^{xy})^2}, \quad (19)$$

where $\hat{x}_L$ is the non-hardened estimator,

$$R_L^{xy} = \int \frac{d^2 l}{(2\pi)^2} \frac{f_{l,L}^y f_{l,L}^x}{2 C_l^{\text{tot}} C_{|L-l|}^{\text{tot}}}, \quad (20)$$

and N_L^x is the estimator normalization, given by $1/R_L^{xx}$. The functions $f_{l,L}^x$ and $f_{l,L}^y$ set the mode-coupling generated by the fields x and y , e.g. the function for lensing is given above in equation (17).

The noise on the bias-hardened estimator is

$$N_L^{B,H,x} = N_L^x / (1 - N_L^x N_L^y (R_L^{xy})^2). \quad (21)$$

For the case of lensing, the contaminant field is ϕ_L , with $f_{l,L}^\phi$ given by equation (17). Our ‘lensing-hardened’ estimator is then

$$K_L^{\text{LH}} = \frac{\hat{K}_L - N_L^K R_L^{K\phi} \hat{\phi}_L}{1 - N_L^K N_L^\phi (R_L^{K\phi})^2}. \quad (22)$$

We have implicitly assumed here that the mode-coupling generated by the K field is equivalent to that generated by Poisson-

distributed extended sources with profile $\sqrt{C_l^{\text{pkSZ}}}$, i.e. setting $f_{l,L}^K = \sqrt{C_l^{\text{pkSZ}} C_{|L-l|}^{\text{pkSZ}}}$. While this is an approximation, we only assume this in order to construct the lensing-hardened estimator, and test the accuracy of this below.

In Fig. 1, we show the lensing bias measured as the average of our C_L^{KK} estimator applied to simulations of the lensed CMB. We use the same filtering and l -ranges as for our baseline data measurement (described in Section 6). We find that lensing-hardening alone reduces the lensing bias to C_L^{KK} by roughly an order of magnitude. We see also that the lensing-hardened estimator can be further improved by subtracting a simulation-based N^1 correction. This term, identified for the CMB lensing power spectrum estimator by Kesden, Cooray & Kamionkowski (2003), is required to correct the $\langle \hat{\phi} \hat{\phi} \rangle$ terms appearing in C_L^{KK} when the lensing-hardened estimator of equation (22) is used.

Bias-hardening can also be used to remove the contribution from Poisson-distributed point-sources (this was one of its original motivations in the CMB-lensing context). We might expect this to significantly help with contamination from the cosmic infrared background (CIB) and radio galaxies. However, one would also expect a significant noise cost, since the form of mode-coupling

induced is similar to our signal of interest (which appear roughly like Poisson-distributed extended sources).

3.2 Extragalactic foregrounds and frequency-cleaning methods

Non-Gaussian signatures in the CMB are also generated by astrophysical sources of microwave radiation, such as radio sources and the CIB and the late-time kSZ and thermal SZ (tSZ) effects, which arise from inverse-Compton scattering of CMB photons. These will in general impart a trispectrum. As well as contributions from the trispectrum of each individual foreground, various other terms will arise due to correlations between the foreground fields (e.g. the tSZ and CIB fields are quite correlated).

Various strategies can be employed to reduce the contamination of the CMB temperature by foregrounds. Finding (e.g. via matched filter methods) and masking (or modelling, or inpainting) point-sources and clusters reduces the contribution from radio sources, the CIB and the tSZ. In addition, one can use the frequency spectral dependence of extragalactic foregrounds to ‘clean’ temperature maps. In internal linear combination (ILC) methods (see e.g. Remazeilles, Delabrouille & Cardoso 2011), linear combinations of observed frequencies are constructed which preserve the CMB signal, but null (or ‘deproject’) one or more assumed foreground frequency spectra. We note that the contribution from the low-redshift kSZ can not be deprojected, since it has the same frequency spectral dependence as the reionization kSZ that we are aiming to measure. Deprojecting foreground spectra will result in a noisier ILC map compared to the ILC where the only constraint is that the variance is minimized. For our case this noise cost is amplified by the fact that we are measuring a trispectrum of the temperature map (e.g. a factor of 2 increase in noise at the map level corresponds to a factor of 8 increase in noise in the trispectrum).

However, here we are interested in the kSZ signal at reionization, which we can assume is to a good approximation uncorrelated with the lower redshift large-scale structure generating the foreground contamination. Unlike for the CMB lensing case then, to remove the foreground contribution from the trispectrum, we need not use a foreground-cleaned map in all four legs of the C_L^{KK} estimator. Denoting our K estimator from two temperature maps Θ_A and Θ_B as $K(\Theta_A, \Theta_B)$, then the following C_L^{KK} estimators are all unbiased by foregrounds:

$$\langle K(\Theta_{\text{clean}}, \Theta_{\text{clean}}), K(\Theta_{\text{clean}}, \Theta) \rangle, \quad (23)$$

$$\langle K(\Theta_{\text{clean}}, \Theta), K(\Theta_{\text{clean}}, \Theta) \rangle, \quad (24)$$

$$\text{and } \langle K(\Theta_{\text{clean}}, \Theta), K(\Theta, \Theta) \rangle, \quad (25)$$

where Θ_{clean} is a foreground-cleaned map, and Θ is the minimum-variance ILC map (see e.g. Madhavacheril & Hill 2018; Darwish et al. 2023 for similar approaches, sometimes referred to as ‘gradient-cleaning’, in CMB lensing and Kusiak et al. 2021 for a similar approach for the projected fields kSZ estimator).

These estimators will however have different noise properties, since the statistical uncertainty on the trispectrum depends on the power spectrum of (and level of correlation between) the input maps. In the realistic case that foreground cleaning does not work perfectly (e.g. since the CIB is not perfectly described by a single l -dependent frequency spectrum), they will also have different levels of residual foreground bias.

We will explore the performance of various of these estimators on simulations in Section 4.

4 FOREGROUND CONTAMINATION ESTIMATES

In order to predict the levels of extragalactic foreground contamination, we use the ACT DR6-like versions of the WEBSKY (Stein et al. 2020) and Sehgal et al. (2010; S10 henceforth) simulations produced in MacCrann et al. (2023) – see section 4 of that work for a description of the simulation processing. The simulations include tSZ, late-time kSZ, radio sources, and the CIB. After adding simulated CMB and realistic ACT DR6-like noise, we run the NEMO⁸ source and cluster finding code, and use the outputs to subtract models for $S/N > 4$ point sources, and the tSZ contamination from $S/N > 5$ clusters (the same thresholds applied in the ACT DR6 data processing). The predicted contamination to C_L^{KK} , which we will refer to as ΔC_L^{KK} , is simply estimated by measuring the C_L^{KK} of the total (i.e. summed over all astrophysical components) residual foreground map, after applying the same sky mask we use when performing the data measurement (see Section 5). We correct for the mean-field (necessary since we apply a mask to the foreground simulations in order to make realistic the source and cluster finding) and N^0 .⁹

As described in Section 3.2, in each leg of our trispectrum estimator, we can use a (possibly different) linear combination of the temperature maps from our range of frequency channels (channels centred at 97 and 149 GHz from ACT, and the 217, 353, and 545 GHz channels from *Planck*). We test linear combinations that minimize the variance in harmonic space, and also the case where frequency spectra for tSZ and CIB are additionally deprojected. Specifically, we generate temperature maps with the following configurations:

(i) Minimum variance ILC, ‘MV ILC’ henceforth.

(ii) tSZ and CIB-deprojected ILC, ‘tSZ + CIB-nulled ILC’ henceforth. We assume a CIB frequency spectrum:

$$f_{\text{CIB}}(\nu) \propto \frac{\nu^{3+\beta}}{e^{h\nu/k_B T_{\text{CIB}}} - 1} \left(\frac{dB(\nu, T)}{dT} \Big|_{T_{\text{CMB}}} \right)^{-1}, \quad (27)$$

where h is Planck’s constant, k_B is the Boltzmann constant, $B(\nu, T)$ is the Planck function, and T_{CMB} is the CMB temperature. We set $\beta = 1.2$ and $T_{\text{CIB}} = 24\text{K}$ (following Ade et al. 2019). We note this is only an approximation to the true (l -dependent) CIB frequency spectrum, so we should not expect CIB to be perfectly removed.

We also investigated CIB-deprojected ILC (without tSZ-deprojection), as well as deprojecting additional components using the ‘moments method’ (Chluba, Hill & Abitbol 2017; Rotti & Chluba 2021), but find that they do not reduce biases further (we think that given the relatively high noise in the high-frequency *Planck* channels at the high l we are using, there is insufficient information to usefully constrain the additional components introduced in the moments method).

We show in Fig. 1 the fractional bias to C_L^{KK} predicted from the WEBSKY (blue lines) and S10 (orange lines) simulations for several estimator variations. The bias, ΔC_L^{KK} is divided by the Alvarez (2016) C_L^{KK} prediction, hence can be interpreted as a fractional bias if one assumes the Alvarez (2016) prediction is correct. In the legend,

⁸<https://nemo-sz.readthedocs.io/en/latest/>

⁹For the bias-hardened case, this is given by

$$N_{\text{fg}, \text{BH}}^0 = \frac{N_{\text{fg}}^0 + (R^{xy})^2 (N^x)^2 N^y - 2R^{xy} N^y (A^x)^2 R_{\text{fg}}^{xy}}{[1 - N^x N^y (R^{xy})^2]^2}, \quad (26)$$

where N_{fg}^0 is the N^0 for the foreground-only map without bias hardening, and all quantities are functions of L .

we also include the predicted bias to A_{A16} , denoted ΔA_{A16} . We show the bias for the following cases, which we find to perform best:

- (i) We use a minimum-variance ILC in all trispectrum legs, labelled ‘ILC’.
- (ii) The same as (i) but replacing the temperature map in just one leg with a tSZ and CIB-nulled version, labelled ‘one leg (tSZ and CIB)-nulled’.

We also checked the case where we deproject tSZ and CIB from more than one trispectrum leg, but this generally results in larger foreground biases,¹⁰ as well as higher noise. While the ‘ILC’ case shows very significant biases for both simulations, this estimator does have the advantage that the foreground biases are very likely positive,¹¹ so a data measurement can be interpreted simply as an upper-limit on the reionization kSZ trispectrum. In contrast, the case where one trispectrum leg is ‘cleaned’ may generate positive or negative foreground biases (since the residual foreground contribution in the cleaned leg may be correlated or anticorrelated with the foreground contributions in the other legs). While there is still significant uncertainty in the theoretical prediction of the foreground bias (as evidenced by the difference in the WEBSKY and S10 predictions), we conclude that the strict positive upper-limit makes the ILC case more interpretable, and we use that as our baseline case.

We note here that for our current data, the biases due to extragalactic foregrounds are significantly larger than those due to lensing (especially when lensing-hardening is used) – note that the lensing biases have been multiplied by 10 in Fig. 1. There will also exist biases to C_L^{KK} arising from the correlation between lensing and extragalactic foregrounds, i.e. trispectrum terms like

$$\langle K(\Theta_{\text{cmb}}, \Theta_{\text{cmb}})K(\Theta_{\text{fg}}, \Theta_{\text{fg}}) \rangle, \quad (28)$$

$$\text{and } \langle K(\Theta_{\text{cmb}}, \Theta_{\text{fg}})K(\Theta_{\text{cmb}}, \Theta_{\text{fg}}) \rangle, \quad (29)$$

which depend on the bispectrum $\langle \phi \Theta_{\text{fg}} \Theta_{\text{fg}} \rangle$, where Θ_{fg} is the temperature perturbation due to foregrounds. However, given the dominance of the foreground-only biases for our current data, we do not consider those cross-terms here.

We investigate in Appendix B the dependence of the fractional foreground bias on the maximum CMB temperature multipole, l used, for the ILC case. For both WEBSKY and S10, the foreground biases increase when increasing l_{max} , and we decide to use $l_{\text{max}} = 4000$ as our baseline choice. We use $l_{\text{min}} = 3000$ throughout. We conclude that extragalactic foregrounds are a major systematic for the kSZ trispectrum, and will dominate the limits on the kSZ reionization trispectrum in this work.

5 ACT DR6 + Planck DATA AND SIMULATIONS USED

We use CMB data from the ACT. ACT was a telescope located in the Atacama Desert in northern Chile and observed the millimetre wavelength sky from 2007 until 2022. For this analysis we use night-time temperature data collected from 2017 to 2022 in the CMB-dominated bands $F090$ (77–112 GHz) and $F150$ (124–172 GHz).

¹⁰This can occur when residual, i.e. post-deprojection foreground contributions are amplified by the large weights required to null the assumed frequency spectra in the ILC.

¹¹This is true for Poisson distributed foreground sources, which is likely a good approximation for the high CMB temperature l s we use here. We note also that in both WEBSKY and S10 simulations the foreground bias for the ILC case is strongly positive.

These maps have a pixel scale of 0.5 arcmin resolution, with a beam full width at half-maximum of 2.2 and 1.5 arcmin for $F090$ and $F150$, respectively. These observations at $F1090$ ($F150$) were made using two (three) dichroic detector modules, known as polarization arrays (PAs; Thornton et al. 2016). For each frequency channel and PA, four independent data split maps were constructed by using different subsets of the time-ordered data. For this analysis, we obtain four independent data splits for a given frequency channel by combining the data for given split from each PA. We do this by performing an inverse-variance coadd in harmonic space (following Qu et al. 2024, see section 5.6 of that work for details) allowing us to use the cross-correlation estimator described in Section 2.3.

We find approximate equivalent white noise levels of $12.4 \mu\text{K}$ arcmin and $16.3 \mu\text{K}$ arcmin for $F090$ and $F150$, respectively, for the coadd of the four independent splits. We subtract and inpaint a catalogue of 1779 objects that include especially bright sources and extended sources with $\text{SNR} > 10$. We further subtract tSZ model images corresponding to clusters detected with the NEMO software. Refer to MacCrann et al. (2023) for further details regarding point-source and cluster template subtraction. We also include the *Planck*NPIPE (Planck Collaboration LVII 2020) 217, 353, and 545 GHz channels, from which we model and subtract $S/N > 5$ point sources.

For all frequencies (ACT and *Planck*) we mask Fourier modes with $|\ell_x| < 90$ and $|\ell_y| < 50$ to remove contamination by ground, magnetic, and other types of pick-up in the data due to the scanning of the ACT telescope. We apply a 60 per cent galactic sky mask, with a 3° apodization, following Qu et al. (2024).

For *Planck* we do not have four independent data splits, so we just use the same *Planck* data in each leg of the trispectrum estimator. This means our mean-field and N^0 is independent of ACT instrumental noise, but not of *Planck* instrumental noise. However, the *Planck* noise is easier to simulate (e.g. as there is no atmospheric noise and the scan strategy is simpler). Furthermore, for our baseline result where we use a minimum-variance ILC, the *Planck* channels contribute very little weight to the ILC anyway given their higher map noise and larger beam.

We construct the covariance matrix for harmonic ILC weights using heavily-smoothed¹² auto and cross-power spectra measured from each ACT and *Planck* frequency. For the ACT channels, we measure these power spectra from the average of the four independent data splits. Having constructed the ILC maps, we again measure heavily smoothed, beam-deconvolved angular power spectra, which are used as C_l^{tot} in the filter function $W_l = \sqrt{C_l^{\text{pkSZ}}/C_l^{\text{tot}}}$.

We did not find significant improvement from using the Needlet ILC method presented in Coulton et al. (2024), perhaps because in the high l regime considered here, the foregrounds are fairly isotropic. Hence, we decided to stick with a harmonic ILC approach given the lesser computational cost.

5.1 Simulations for bias corrections and covariance

To accurately calculate the $\text{RD}N^0$ and mean-field bias corrections we require simulations with the same power spectrum as the data. For the $\text{RD}N^0$ calculation the simulations need not contain a non-Gaussian kSZ component, while for the mean-field calculation Gaussian signal simulations should also suffice, since we are interested only in

¹²We use a Savitzky–Golay filter (Savitzky & Golay 1964) of window length 301, order 2, as implemented at https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.savgol_filter.html

the statistical anisotropy generated by the noise and the mask. We generate 200 Gaussian signal simulations, ensuring these have the correct power spectrum via the following procedure:

(i) Measure auto and cross-power spectra for all frequency (or ILC) maps, apply a w_2 correction to account for the mask,¹³ and apply a smoothing filter.¹⁴ Call this $C_{l,ij}^{\text{data}}$ for frequencies i and j . We checked our results are not sensitive to using an increased smoothing window length.

(ii) Generate Gaussian random fields with covariance $C_{l,ij}^{\text{data}}$. Apply the survey mask and k-space masking. Measure the resulting power spectra $C_{l,ij}^{\text{cal}}$, again applying the w_2 correction.

(iii) Define

$$T_i = C_{l,ij}^{\text{cal}} / C_{l,ij}^{\text{data}}. \quad (30)$$

(iv) Generate Gaussian random fields with power $C_{l,ij} = C_{l,ij}^{\text{data}} / \sqrt{T_i T_j}$. Apply survey mask and k-space masking as on the data.

We then add these signal simulations to noise-only simulations generated following Atkins et al. (2023) for the ACT channels, and using those provided by Planck Collaboration LVII (2020) for the *Planck* channels. We use these simulations to calculate the mean-field and $\text{RD}N^0$ biases, from which we can calculate our corrected C_L^{KK} according to equation (15).

To compute the covariance matrix and T_L^{KK} (see equation 15), we add a randomly rotated version of the Alvarez (2016) reionization kSZ simulation to each realization. The covariance matrix is then computed from the application of our C_L^{KK} estimator (equation 15) to each of these simulation realizations. We do not recalculate $\text{RD}N^0$ for every simulation realization, instead replacing it with $\text{MC}N^0$ (see Section 2.3). This likely results in a small overestimation of the off-diagonal covariance, however, we will show that our upper limit on the kSZ trispectrum is dominated by the foreground bias rather than statistical uncertainty, so we believe it is reasonable to avert the very high computational cost of including the $\text{RD}N^0$ correction for every realization. We note also that the covariance matrix simulations do not then include non-Gaussian foregrounds, which may result in an underestimate of the covariance. Again, given that the upper limit on the kSZ trispectrum is dominated by foreground bias rather than statistical uncertainty, we do not believe this potential underestimate significantly affects our conclusions.

Finally, we use these simulations to compute a multiplicative correction to C_L^{KK} that arises largely from the Fourier mode masking described in Section 5. For each simulation realization we compute the ratio of our mean-field, N^0 and w_4 -corrected C_L^{KK} estimator to the input, reionization kSZ-only C_L^{KK} . Averaged over realizations, we infer a close to scale-independent multiplicative bias $T_L^{KK} = 0.94$.

6 RESULTS

6.1 ACT DR6 + *Planck* measurements of C_L^{KK}

We show in Fig. 2 the measurement from ACT DR6 + *Planck*, combined via a harmonic space ILC, as described in Section 5. We calculate the best-fitting rescaling amplitude of the Alvarez

¹³We divide all power spectra by $w_2 = \sum_i a_i (m^i)^2 / \sum_i a_i$ where the sum is over pixels i , a_i is the area of pixel i , and m^i is the value of the (apodized) mask (between 0 and 1) for pixel i .

¹⁴We use a Savitzky–Golay filter of window length 101, order 2.

(2016) theory prediction, which is reported in the legend. Following Smith & Ferraro (2017) and Ferraro & Smith (2018), we use $l_{\min} = 3000$. For our baseline measurement, we use $l_{\max} = 4000$. We find consistent results when we use $l_{\max} = 5000$, but the reduction in uncertainty is modest (see left panel of Fig. 3), and as shown in Section 4 the expected foreground biases are smaller for $l_{\max} = 4000$.

The measurement is generally well above the Alvarez (2016) theory prediction, as indicated by the amplitude fit $A_{A16} = 56.4 \pm 19.9$ (see Section 2.1 for the definition of A_{A16}). It is quite consistent with the simulation predictions of the contamination from extragalactic foregrounds (dashed and dotted lines from the WEBSKY and S10 simulations, respectively). The plotted errorbars are the diagonal of a covariance matrix calculated by running the C_L^{KK} estimator on the simulations described in Section 5.1.

In the right panel of Fig. 3, we show cases where we use foreground-cleaned temperature in one of the four trispectrum legs. Orange points show the case where frequency spectra for tSZ and CIB are nulled in one of the trispectrum legs. The C_L^{KK} uncertainty roughly doubles for this case, however, in our scenario, where the upper limit is dominated by the foreground biases, this noise cost is not particularly relevant. The result from the baseline (ILC-only) case has the advantage that the foreground bias is very likely positive, hence the measurement can be straightforwardly interpreted as an upper limit. In contrast, the case where we deproject tSZ and CIB from one leg of the estimator is no longer an autocorrelation, so the residual foreground bias need not be positive.

We do not perform a model fit due to the significant uncertainties in the foreground template, as can be inferred from the difference in the predictions from the WEBSKY and S10 simulations. We discuss improvements to data and methodology designed to enable robust modelling of the signal in Section 7.

When we apply lensing-hardening, we obtain measurements consistent with the baseline case ($A_{A16} = 77 \pm 33$). In the current regime where other foregrounds are dominant, lensing-hardening does not yet improve our upper limits. We find also that point-source hardening results in a very large increase in uncertainty (a factor of 20 with respect to the baseline case). We conclude that our current data do not yet warrant the use of point-source hardening, and we do not present measurements using them here.

6.2 Robustness of N^0 and mean-field subtraction

We discuss here the importance for our measurement of using the cross-correlation based estimator (as discussed in Section 2.3), and a realization-dependent N^0 calculation. The left-hand panel of Fig. 4 demonstrates the improvement when using the cross-correlation based estimator. Both the mean-field and N^0 are much smaller, and given the lack of dependence on the observational noise, presumably more accurate, leading to a tighter upper limit when using the cross-correlation based estimator (blue data points). We also show a set of green points where we do use the cross-correlation-based estimator, but the simulations used for the mean-field and N^0 bias calculations are not matched to data. Instead they are generated assuming a power spectrum that is the sum of a primary CMB theory prediction, a foreground power spectrum contribution estimated from WEBSKY, and a reionization kSZ power spectrum contribution estimated from the Alvarez (2016) simulation (labelled ‘CMB + WEBSKY + A16 sims’). We see a larger amplitude for both of these cases, which we consider to be less robust than our baseline analysis.

In the right-hand panel of Fig. 4, we focus on the N^0 for these same three cases. The solid lines show $\text{RD}N^0$ and dashed lines $\text{MC}N^0$.




Figure 2. Blue points show C_L^{KK} from ACT and *Planck* data (combined via a harmonic space ILC). Negative points are indicated by open symbols. The dashed and solid lines show the corresponding predictions of the foreground bias from WEBSKY and S10, respectively, and do not include the reionization *k*SZ component that we are most interested in here. The black solid line shows the prediction of the *k*SZ 4-point signal from the Alvarez (2016) simulations. Thus, we believe the measurement is dominated by foregrounds, as discussed in Sections 4 and 6, so should be interpreted only as an upper limit on the reionization signal.

For our baseline case, where the simulations are tuned to match the data power spectrum, we see good agreement between MCN^0 and RDN^0 . For the case where the cross-correlation estimator is not used, we also see good agreement, but given the correction is much larger in this case, and two or three orders of magnitude above the signal, one would need to validate the correction extremely carefully. This motivates our choice to use the cross-correlation-only estimator in this work. Again we also show the ‘CMB + WEBSKY + A16 sims’ case, which reveals differences between RDN^0 and MCN^0 . If one used the naive MCN^0 to correct the signal in this case, one would significantly bias the C_L^{KK} signal low. The use of RDN^0 in this case does lead to better agreement with the baseline case which we consider most robust.

7 CONCLUSIONS

We presented a limit on the *k*SZ trispectrum from wide-field CMB data which combines observations from ACT DR6 and *Planck*. Estimating this trispectrum from this data set has presented various challenges, especially the levels of foreground and CMB lensing contamination, and uncertainties in the mean-field and N^0 corrections

that are potentially large compared to the signal of interest. Similarly to Raghunathan et al. (2024), we find that the contamination due to foregrounds dominates over the reionization signal.

We address these challenges by presenting a range of mitigation schemes. For lensing and extragalactic foregrounds, we demonstrated the performance of bias-hardening for the first time in this context. Bias hardening against CMB lensing can very effectively remove the bias due to CMB lensing, and is likely to be useful beyond the *k*SZ trispectrum measurement, for example in the ‘projected fields’ method of *k*SZ measurement (Dore, Hennawi & Spergel 2004; Ferraro et al. 2016; Hill et al. 2016).

We also proposed new trispectrum estimators where only a single leg has foreground frequency spectra nulled, which have significantly better performance than nulling all trispectrum legs. For the mean-field and N^0 , also inspired by CMB lensing approaches, we demonstrate the utility of cross-correlation estimators and realization-dependent N^0 , as well as the importance of having simulations whose power spectrum match the data.

With these approaches in hand, we measure a trispectrum from ACT DR6 + *Planck* data that has an amplitude roughly 50 times the expected reionization signal from the Alvarez (2016) simula-




Figure 3. Measurements of C_L^{KK} from ACT and *Planck* data (combined via a harmonic space ILC), highlighting alternative choices for analysis. Negative data points are indicated by open symbols. *Left panel:* The blue data points show our baseline measurement which uses CMB temperature $l_{\max} = 4000$, while the orange data points use $l_{\max} = 5000$. *Right panel:* The blue data points show our baseline measurement which uses a minimum variance harmonic-space ILC in all trispectrum legs, while for the orange points, one of the four trispectrum legs has tSZ and CIB frequency spectra nulled (see Section 3.2). Note that we have switched to a linear scale here given the expectation that the tSZ and CIB-nulled measurement can be negative, and find that removing the L^2 scaling allows for a clearer view of the data points (and also results in a very slightly different amplitude fit for the baseline case). As elsewhere, the black solid line shows the prediction of the kSZ 4-point signal from the Alvarez (2016) simulations and the dashed and solid lines show the corresponding predictions of the foreground bias from WEBSKY and S10, respectively.




Figure 4. Demonstration of the importance of the cross-correlation based estimator of Madhavacheril et al. (2021), and using simulations with the correct power spectrum, for our measurement. *Left panel:* The orange data points, which do not restrict to cross-correlations in the trispectrum estimation, result in a worse upper-limit compared to the baseline case (blue data points). Green data points do use the cross-correlation-based estimator, but do not use a set simulations that have power spectrum tuned to the data (see Section 6.2 for details). Negative data points are indicated by open symbols. *Right panel:* N^0 estimates for the same three cases. RDN^0 is shown in solid lines, and MCN^0 in dashed lines. For the baseline case in blue, there is good agreement between RDN^0 and MCN^0 , arising from the good agreement in the power spectrum between simulations and data. For the cases where the cross-correlation estimator is not used the N^0 is much higher, making this estimator more susceptible to small errors in the N^0 correction. For the ‘CMB + WEBSKY + A16’ case, we do see significant differences between MCN^0 and RDN^0 .

tions. This signal is roughly consistent with the expectation from extragalactic foregrounds, as we have demonstrated using two independent foreground simulations (from Stein et al. 2020 and

Sehgal et al. 2010). Thus, we are not yet in a position to constrain realistic reionization scenarios.

As this work was being finalized, Raghunathan et al. (2024) presented an upper-limit on the reionization kSZ trispectrum from SPT + *Herschel*-SPIRE data that is ~ 5 to 10 times tighter than this work. The authors argued that their inclusion of high-resolution, high-frequency data from *Herschel*-SPIRE mitigated CIB contamination. Even more importantly, since the SPT maps have roughly 5 times lower noise than the ACT maps used here (though over much smaller area), the authors were able to apply much stricter flux thresholds for point-sources and clusters which are the main contributors to foreground contamination. In this paper, we have emphasized methodologies for robust first detections of the small signatures of reionization on the CMB trispectrum in the future, and demonstrated them on the ACT DR6 + *Planck* data. We anticipate that techniques described here will be essential: bias-hardening, RDN^0 and cross-correlation-based estimators.

We believe the constraint presented here can be improved by incorporating further ACT data (in particular data taken during daytime if we can demonstrate adequate systematics control), and pursuing more aggressive source and cluster finding. An analysis that more optimally takes into account the noise variations in the ACT data (both when performing the ILC, and in the filtering of the temperature modes entering the trispectrum estimator) will also be a focus of future work. Beyond ACT, upcoming data from Simons Observatory and CMB-S4, which will have an order of magnitude lower noise, and additional frequency channels, have the potential to make the kSZ trispectrum a highly informative probe of reionization (Alvarez et al. 2021).

ACKNOWLEDGEMENTS

Support for ACT was through the U.S. National Science Foundation through awards AST-0408698, AST-0965625, and AST-1440226 for the ACT project, as well as awards PHY-0355328, PHY-0855887, and PHY-1214379. Funding was also provided by Princeton University, the University of Pennsylvania, and a Canada Foundation for Innovation (CFI) award to UBC. ACT operated in the Parque Astronómico Atacama in northern Chile under the auspices of the Agencia Nacional de Investigación y Desarrollo (ANID). The development of multichroic detectors and lenses was supported by NASA grants NNX13AE56G and NNX14AB58G. Detector research at NIST was supported by the NIST Innovations in Measurement Science programme. Computing for ACT was performed using the Princeton Research Computing resources at Princeton University, the National Energy Research Scientific Computing Center (NERSC), and the Niagara supercomputer at the SciNet HPC Consortium. SciNet is funded by the CFI under the auspices of Compute Canada, the Government of Ontario, the Ontario Research Fund–Research Excellence, and the University of Toronto. We thank the Republic of Chile for hosting ACT in the northern Atacama, and the local indigenous Licanantay communities whom we follow in observing and learning from the night sky.

CS acknowledges support from the Agencia Nacional de Investigación y Desarrollo (ANID) through Basal project FB210003. EC acknowledges support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 849169). OD acknowledges support from an SNSF Eccellenza Professorial Fellowship (No. 186879). JCH acknowledges support from NSF grant AST-2108536, DOE grant DE-SC0011941, NASA grants 21-ATP21-0129 and 22-ADAP22-0145, the Sloan Foundation, and the Simons Foundation. RH acknowledges support from CIFAR, the Azrieli and Alfred P. Sloan foundations, and the NSERC Discovery Grants

programme. NM acknowledges the support of the Royal Society, grant URFR1221806.

DATA AVAILABILITY

We will make the baseline analysis data points and covariance matrices available upon publication of this work.

REFERENCES

- Ade P. et al., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 056
 Alvarez M. A., 2016, *ApJ*, 824, 118
 Alvarez M. A., Ferraro S., Hill J. C., Hložek R., Ikape M., 2021, *Phys. Rev. D*, 103, 063518
 Atkins Z. et al., 2023, *J. Cosmol. Astropart. Phys.*, 2023, 073
 Battaglia N., Natarajan A., Trac H., Cen R., Loeb A., 2013, *ApJ*, 776, 83
 Calabrese E. et al., 2014, *J. Cosmol. Astropart. Phys.*, 2014, 010
 Chen N., Trac H., Mukherjee S., Cen R., 2023, *ApJ*, 943, 138
 Chluba J., Hill J. C., Abitbol M. H., 2017, *MNRAS*, 472, 1195
 Choi S. K. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 045
 Coulton W. et al., 2024, *Phys. Rev. D*, 109, 063530
 Darwish O. et al., 2023, *Physical Review D*
 Dore O., Hennawi J. F., Spergel D. N., 2004, *ApJ*, 606, 46
 Eisenstein D. J. et al., 2023, preprint (arXiv:2306.02465)
 Ferraro S., Hill J. C., Battaglia N., Liu J., Spergel D. N., 2016, *Phys. Rev. D*, 94, 123526
 Ferraro S., Smith K. M., 2018, *Phys. Rev. D*, 98, 123519
 Finkelstein S. L. et al., 2023a, *ApJ*, 969, L2
 Finkelstein S. L. et al., 2023b, *ApJ*, 946, L13
 Gorce A., Douspis M., Salvati L., 2022, *A&A*, 662, A122
 Hanson D., Challinor A., Efstathiou G., Bielewicz P., 2011, *Phys. Rev. D*, 83, 043005
 Hill J. C., Ferraro S., Battaglia N., Liu J., Spergel D. N., 2016, *Phys. Rev. Lett.*, 117, 051301
 Kesden M., Cooray A., Kamionkowski M., 2003, *Phys. Rev. D*, 67, 123507
 Kogut A., 2003, *New A Rev.*, 47, 977
 Kusiak A., Bolliet B., Ferraro S., Hill J. C., Krolewski A., 2021, *Phys. Rev. D*, 104, 043518
 MacCrann N. et al., 2023, *ApJ*, 966, 138
 Madhavacheril M. et al., 2021, *Journal of Cosmology and Astroparticle Physics*, 028
 Madhavacheril M. S., Hill J. C., 2018, *Phys. Rev. D*, 98, 023534
 McQuinn M., Furlanetto S. R., Hernquist L., Zahn O., Zaldarriaga M., 2005, *ApJ*, 630, 643
 Namikawa T., Hanson D., Takahashi R., 2013, *MNRAS*, 431, 609
 Osborne S. J., Hanson D., Doré O., 2014, *J. Cosmol. Astropart. Phys.*, 2014, 024
 Planck Collaboration LVII, 2020, *A&A*, 643, A42
 Planck Collaboration XLVII, 2016, *A&A*, 596, A108
 Qu F. J. et al., 2024, *ApJ*, 962, 112
 Raghunathan S. et al., 2024, preprint (arXiv:2403.02337)
 Reichardt C. L. et al., 2021, *ApJ*, 908, 199
 Remazeilles M., Delabrouille J., Cardoso J.-F., 2011, *MNRAS*, 410, 2481
 Rotti A., Chluba J., 2021, *MNRAS*, 500, 976
 Sailer N., Schaan E., Ferraro S., 2020, *Phys. Rev. D*, 102, 063517
 Santos M. G., Cooray A., Haiman Z., Knox L., Ma C.-P., 2003, *ApJ*, 598, 756
 Savitzky A., Golay M. J. E., 1964, *Anal. Chem.*, 36, 1627
 Sehgal N., Bode P., Das S., Hernandez-Monteagudo C., Huppenberger K., Lin Y.-T., Ostriker J. P., Trac H., 2010, *ApJ*, 709, 920
 Shaw L. D., Rudd D. H., Nagai D., 2012, *ApJ*, 756, 15
 Smith K. M., Ferraro S., 2017, *Phys. Rev. Lett.*, 119, 021301
 Stein G., Alvarez M. A., Bond J. R., Engelen A. v., Battaglia N., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 012
 Thornton R. J. et al., 2016, *ApJS*, 227, 21




Figure A1. Predictions of the kSZ power spectrum (left panel), and trispectrum (right panel) from reionization, using the AMBER simulations (Trac et al. 2022; Chen et al. 2023). We show the fractional difference with respect to a fiducial AMBER simulation (the third line in the legend). We also show the prediction based on the Alvarez (2016) simulations (blue dashed lines). For the trispectrum, below $L = 10$ the fractional differences become very noisy thus we only show the range $L > 10$ for clarity.

Trac H., Bode P., Ostriker J. P., 2011, *ApJ*, 727, 94

Trac H., Chen N., Holst I., Alvarez M. A., Cen R., 2022, *ApJ*, 927, 186

APPENDIX A: KSZ IN THE AMBER SIMULATIONS

Fig. A1 shows the kSZ power spectrum (left panel) and trispectrum (right panel) measured from the AMBER simulations (Trac et al. 2022; Chen et al. 2023). In both cases the predictions are normalized by the measurement from the fiducial AMBER simulation.¹⁵ The AMBER simulations vary several reionization parameters:

- (i) Δz : The duration of reionization, defined as $z_{10} - z_{90}$, where z_{10} and z_{90} are the redshifts where the Universe is 10 per cent and 90 per cent ionized, respectively.
- (ii) z_{mid} : The redshift where the Universe is 50 per cent ionized.
- (iii) A_z : An asymmetry parameter defined $(z_{10} - z_{\text{mid}})/(z_{\text{mid}} - z_{90})$.
- (iv) M_h : The minimum halo mass hosting ionizing sources (specified in Fig. A1 in M_*).
- (v) λ_{mfp} : The mean free path of ionizing photons (specified in Fig. A1 in $\text{Mpc } h^{-1}$).

¹⁵That is not the Alvarez (2016) simulation, which we use as a fiducial template throughout this paper.

As argued in Section 2, we see a strong dependence for both power spectrum and trispectrum on Δz and z_{mid} , with longer or earlier reionization increasing both signals. We show also, as the blue dashed line, the prediction from Alvarez (2016) that we use as a template throughout this paper. This predicts a higher power spectrum and trispectrum than most of the AMBER variations which can be explained by the relatively high choice of $z_{\text{mid}} = 10$. The other parameters varied in AMBER have less impact on the kSZ power spectrum and trispectrum, but see figs 10 and 13 in Chen et al. (2023) for a more detailed view.

We see also that the response to changing these parameters is different for the power spectrum and trispectrum, demonstrating the complementary information in the two statistics (see Alvarez et al. 2021 for forecasts of degeneracy-breaking power of these two statistics).

When foregrounds are taken into account, which also affect the two signals differently, the gain from including the trispectrum information can be even greater. While the same is true in theory for the other varied parameters, A_z , M_h , and λ_{mfp} , the responses for the trispectrum are relatively small, and thus will be difficult to detect with CMB data.

APPENDIX B: CMB TEMPERATURE l -DEPENDENCE OF FOREGROUND BIASES

In Fig. B1, we demonstrate the dependence of the fractional foreground bias, on the maximum CMB temperature multipole, l used, for the ILC case. For both WEBSKY and S10, there is a significant




Figure B1. Expected fractional foreground biases from the WEBSKY (blue lines) and S10 (orange lines) simulations. Here, we show the impact on the foreground bias of changing the range of scales l entering the quadratic estimator.

increase in the foreground biases when increasing $l_{\max}$, hence we use $l_{\max} = 4000$ for our baseline analysis.

APPENDIX C: NORMALIZATION CONVENTIONS

While we have motivated the quadratic estimator for K by analogy with CMB lensing, the choice of normalization of the K estimator is less obvious than in the case of CMB lensing, where one is reconstructing the lensing potential ϕ . In the main body of this work, we have adopted a convention that is closely related to the Sailer et al. (2020) source estimator, mainly for convenience when adapting the CMB lensing codes to measure and debias the K statistic (and include lensing-hardening etc.). The Sailer et al. (2020) source estimator is defined:

$$s_L^{S20} = N^s(L) \int \frac{d^2l}{(2\pi)^2} \frac{u_l u_{|L-l|} T_l T_{L-l}}{2u_L C_l^{\text{tot}} C_{|L-l|}^{\text{tot}}}, \quad (\text{C1})$$

where u_L is the source profile and

$$[N^s(L)]^{-1} = \int \frac{d^2l}{(2\pi)^2} \frac{1}{2C_l^{\text{tot}} C_{|L-l|}^{\text{tot}}} \left[\frac{u_l u_{|L-l|}}{u_L} \right]^2. \quad (\text{C2})$$

After substituting $u_L = \sqrt{C_l^{\text{pkSZ}}}$, our K estimator is similar, except we divide out dependence on u_L , since our measurement should not depend on C_l^{pkSZ} at large scales:

$$K_L = \left[\int \frac{d^2l}{(2\pi)^2} u_l u_{|L-l|} W_l W_{|L-l|} \right]^{-1}, \quad (\text{C3})$$

$$\times \int \frac{d^2l}{(2\pi)^2} W_l W_{|L-l|} T_l T_{L-l}. \quad (\text{C4})$$

The normalization factor here is somewhat different to that chosen by Smith & Ferraro (2017), who include a C_l^{tot} factor in their normalization, aiming to construct a K statistic that is a

measure of the kSZ non-Gaussianity relative to the total power spectrum:

$$\bar{K} = \int \frac{d^2l}{(2\pi)^2} W_s^2(l) C_l^{\text{tot}}. \quad (\text{C5})$$

C_l^{tot} is however both survey-dependent (if it includes noise) and the signal component is quite uncertain in the $l > 3000$ regime we are considering. This would motivate a normalization where one simply divides out the effect of the filtering, which we call K_{filter} , defined

$$K_{\text{filter}} = \int \frac{d^2l}{(2\pi)^2} (W_l^{\text{eff}} W_{|L-l|}^{\text{eff}}), \quad (\text{C6})$$

$$\approx \int \frac{d^2l}{(2\pi)^2} (W_l^{\text{eff}})^2, \quad (\text{C7})$$

$$\approx \int \frac{ldl}{2\pi} (W_l^{\text{eff}})^2. \quad (\text{C8})$$

In the second line, we have assumed the squeezed limit $L \ll l$.

Raghunathan et al. (2024) adopt a different normalization again:

$$K_L^{R24} = \bar{F}^{-1} \int \frac{d^2l}{(2\pi)^2} W_l W_{|L-l|} T_l T_{L-l}, \quad (\text{C9})$$

where

$$\bar{F} = \int W_l^2 B_{l,\text{eff}}^2 F_l d\ln l. \quad (\text{C10})$$

Note for simplicity we have neglected their f_{sky} corrections, as we have in the discussion above. $B_{l,\text{eff}}^2$ denotes their effective beam, and can be considered as part of an effective filter W_l^{eff} .

In case useful for comparison purposes, we show in Fig. C1 versions of our Fig. 2 using the K_{filter} and $\bar{F}$ normalizations.

¹DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK

²Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

³Center for Data-Driven Discovery, Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, 277-8583, Japan

⁴Astrophysics Group, Cavendish Laboratory, J. J. Thomson Avenue, Cambridge CB3 0HE, UK

⁵Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh, PA 15260 USA

⁶School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, Wales CF24 3AA, UK

⁷Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA

⁸Département de Physique Théorique et CAP, Université de Genève, 24 Quai Ansermet, CH-1211 Genève 4, Switzerland

⁹Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, USA

¹⁰Dunlap Institute for Astronomy and Astrophysics, University of Toronto, Toronto, ON M5S 3H4, Canada

¹¹Department of Physics, Columbia University, New York, NY 10027, USA

¹²Wits Centre for Astrophysics, School of Physics, University of the Witwatersrand, Private Bag 3, Johannesburg 2050, South Africa

¹³Astrophysics Research Centre, School of Mathematics, Statistics, and Computer Science, University of KwaZulu-Natal, Westville Campus, Durban 4041, South Africa

¹⁴Dunlap Institute of Astronomy & Astrophysics, 50 St. George St., Toronto, ON M5S 3H4, Canada

¹⁵David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, 50 St. George St., Toronto, ON M5S 3H4, Canada

¹⁶School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA

¹⁷Department of Physics and Astronomy, University of Pennsylvania, 209 South 33rd Street, Philadelphia, PA 19104, USA

¹⁸A. Astrophysics Research Centre, University of KwaZulu-Natal, Westville Campus, Durban 4041, South Africa

¹⁹B. School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Durban 4041, South Africa

²⁰Physics and Astronomy Department, Stony Brook University, Stony Brook, NY 11794, USA

²¹Instituto de Física, Pontificia Universidad Católica de Valparaíso, Casilla 4059, Valparaíso, Chile

²²Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

²³McWilliams Center for Cosmology and Astrophysics, Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA

²⁴Department of Physics, Cornell University, Ithaca, NY 14853, USA

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.




Figure C1. Two versions of this work's Fig. 2 with alternative normalizations, in case useful for comparison purposes. In the left panel, we normalize K with K_{filter} (defined in equation C6). In the right panel, we use the Raghunathan et al. (2024) normalization, and also their L -scaling of the y-axis.