

New temperature inversion estimator to detect CMB patchy screening by large-scale structure

Theo Schutt^{1,2,3,*}, Abhishek S. Maniyar^{1,2,3}, Emmanuel Schaan^{1,2}, William R. Coulton^{4,5} and Nishant Mishra^{6,7}

¹*Kavli Institute for Particle Astrophysics and Cosmology,
382 Via Pueblo Mall, Stanford, California 94305, USA*

²*SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, California 94025, USA*

³*Department of Physics, Stanford University, Stanford, California 94305, USA*

⁴*Center for Computational Astrophysics, Flatiron Institute,
162 5th Avenue, New York, New York 10010, USA*

⁵*Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom*

⁶*Department of Astronomy, University of Michigan, Ann Arbor, Michigan 48109, USA*

⁷*Department of Physics, University of California, Berkeley, California 94720, USA*



(Received 24 January 2024; accepted 24 April 2024; published 23 May 2024)

Thomson scattering of cosmic microwave background (CMB) photons imprints various properties of the baryons around galaxies on the CMB. One such imprint, called patchy screening, is a direct probe of the gas density profile around galaxies. It usefully complements the information from the kinematic and thermal Sunyaev-Zel'dovich effects and does not require individual redshifts. In this paper, we derive new estimators of patchy screening called the “temperature inversion” (TI) and “signed” estimators, analogous to the gradient inversion estimator of CMB lensing. Pedagogically, we clarify the relation between these estimators and the standard patchy screening quadratic estimator (QE). The new estimators trade optimality for robustness to biases caused by the dominant CMB lensing and foreground contaminants, allowing the use of smaller angular scales. We perform a simulated analysis to realistically forecast the expected precision of patchy screening measurements from four CMB experiments, ACT, SPT, Simons Observatory (SO) and CMB-S4, cross-correlated with three galaxy samples from BOSS, unWISE and the simulated Rubin LSST Data Challenge 2 catalog. Our results give further confidence in the first detection of this effect from the ACT \times unWISE data in the companion paper and show patchy screening will be a powerful observable for future surveys like SO, CMB-S4 and LSST. Implementations of the patchy screening QE and the TI and signed estimators are publicly available in our `LensQuEst`¹ and `ThumbStack`² software packages, respectively.

DOI: [10.1103/PhysRevD.109.103539](https://doi.org/10.1103/PhysRevD.109.103539)

I. INTRODUCTION

While the cosmic microwave background (CMB) is an extremely powerful window onto the initial conditions of our Universe, it also contains a wealth of information on its late-time large-scale structure (LSS). Indeed, large-scale structure properties including gas density, thermal and bulk velocities, gravitational potential and more leave shadows on the CMB called secondary anisotropies; see [1] for a review. One important class of secondary anisotropies is due to Thomson scattering of the CMB photons by free electrons in the ionized plasma (gas) around galaxies and clusters. This class includes the thermal and kinematic Sunyaev-Zel'dovich (tSZ and kSZ) effects [2–4], which are measured with high precision and have already yielded

crucial insights into cosmology [5–9] and the thermodynamics of gaseous halos around galaxies and clusters [e.g., [10–13]].

Another effect in this class is the patchy screening effect [14–18]. This effect is perhaps the simplest, most intuitive one: a free electron along a given line of sight (LOS) scatters CMB photons out of that LOS, and replaces them with CMB photons from the average blackbody distribution, imprinting a shadow. However, patchy screening is smaller in amplitude than the more complex tSZ and kSZ effects since they scale as

$$\begin{cases} \delta T^{\text{tSZ}} \propto \tau (v_{\text{thermal}}/c)^2, & (v_{\text{thermal}}/c)^2 \sim 10^{-2}, \\ \delta T^{\text{kSZ}} \propto \tau (v_{\text{bulk}}/c), & v_{\text{bulk}}/c \sim 10^{-3}, \\ \delta T^{\text{screening}} \propto \tau (\delta T_0/T_0), & \delta T_0/T_0 \sim 10^{-4}, \end{cases} \quad (1)$$

where $\tau \sim 10^{-3}$ for galaxy group-sized halos.

*schutt20@stanford.edu

¹<https://github.com/EmmanuelSchaan/LensQuEst>.

²<https://github.com/EmmanuelSchaan/ThumbStack>.

These scalings also highlight the complementarity of these different effects. Like kSZ, patchy screening is independent of the gas temperature. It thus scales linearly with halo mass, unlike tSZ, and is thus ideally suited for constraining baryonic effects in galaxy lensing by measuring baryon density profiles around the group-sized halos that dominate the matter power spectrum [12].

Like tSZ, patchy screening does not require any individual redshift information to be detected. While the kSZ effect can also be detected without redshift information [19–21], this is done at the cost of squaring the kSZ, which destroys its linear scaling with gas density and halo mass. Combining screening, kSZ and tSZ allows us to form a complete picture of the gas thermodynamics around halos (gas density, temperature). Screening measurements also enables breaking of the kSZ degeneracy between galaxy gas content and galaxy velocity. Thus, by combining them, one may extract the amplitude of large-scale peculiar velocities in the Universe to test the neutrino masses [22], the growth rate of structure [23], general relativity [24], and measure primordial non-Gaussianity [25]. Finally, while screening and kSZ also probe the epoch of reionization [16,26–29], we will not consider this in this paper, focusing instead on the late-time signal from galaxies and clusters.

Observationally, the patchy screening effect is unique in that it changes sign depending on the sign of the local large-scale CMB temperature. This is unlike extragalactic foregrounds like the cosmic infrared background (CIB) or the tSZ effect, whose sign are fixed at any given frequency. This crucial feature is analogous to kSZ, whose signal changes sign based on the sign of the local LOS peculiar velocity. Like for kSZ, it allows us to cleanly distinguish screening from these foregrounds. This property also differentiates screening from CMB lensing, whose sign changes based on the local large-scale temperature gradient, rather than temperature value.

Leveraging this crucial feature of patchy screening, we derive new estimators dubbed “temperature inversion” (TI) and “signed” estimators, analogous to the gradient inversion (GI) estimator of CMB lensing [30,31]. By making use of the sign change of the screening effect, these estimators are naturally robust to contamination from CIB, tSZ and kSZ, as well as CMB lensing.

We carefully present the relation between these new estimators and the standard quadratic estimator (QE) for patchy screening [16,32,33]. We show how the new estimators do not outperform the standard QE in terms of raw statistical power when using same multipole range. However, in its simplest form, the QE is biased by foregrounds and CMB lensing. While CMB lensing contamination to the QE can be removed via “bias hardening”, the situation may be more complex for extragalactic foregrounds.

While we focus on detecting screening from the CMB temperature only, these estimators generalize trivially to

polarization by simply replacing temperature with Stokes Q and U .

The paper is organized as follows. In Sec. II, we review the physical origin of patchy screening. Section III covers the standard QE for patchy screening and its application to reconstructing both large- and small-scale optical depth τ . Section IV introduces the new TI and signed estimators and clarifies their relationship to the small-scale QE. In Sec. V, we study the biases to the QE due to lensing and extragalactic foregrounds and explain the TI estimator’s robustness to these biases. Section VI provides an overview of our realistic forecast simulations and analysis and presents the resulting patchy screening detection signal-to-noise ratios (SNRs) for TI and QE for a set of current and future CMB and LSS surveys. We conclude with Sec. VII.

II. PATCHY SCREENING OF CMB TEMPERATURE AND POLARIZATION

Consider a line of sight where the CMB temperature T^0 deviates from the sky averaged CMB temperature \bar{T}^0 , by an amount $\delta T^0 = T^0 - \bar{T}^0$. If an electron cloud with optical depth τ to Thomson scattering is present along that line of sight, two things occur. First, a fraction $1 - e^{-\tau}$ of the CMB photons with temperature T^0 is deflected out of the line of sight, leaving only a fraction $e^{-\tau}$ of the incident photons at temperature T^0 . Second, an equal fraction $1 - e^{-\tau}$ of CMB photons from all other directions is scattered into the line of sight. Because this effect averages over all directions, as seen by the electron cloud, these photons have temperature \bar{T}^0 . Thus, the observed temperature T along the line of sight, after Thomson screening, is

$$T = T^0 e^{-\tau} + \bar{T}^0 (1 - e^{-\tau}). \quad (2)$$

Considering temperature deviations with respect to the sky average CMB temperature \bar{T}^0 , we thus obtain,

$$\delta T(\mathbf{x}) = \delta T^0(\mathbf{x}) e^{-\tau(\mathbf{x})}, \quad (3)$$

where \mathbf{x} is the angular position on the sky. The screening effect operates in the same way on temperature T , and Q and U polarizations. For ease of notation, we will thus focus on temperature here, but note that all our results apply to polarization as well.

In polarization, an additional “scattering” effect adds to this screening effect, where the local CMB temperature quadrupole observed by the free electrons is converted to linear polarization. This “scattering” occurs most noticeably at the last scattering surface and around the epoch of reionization. It also occurs at later times due to the ionized gas around galaxies and clusters, an effect which has been called polarized Sunyaev-Ze’doovich (pSZ) effect [34–36].

Going further, we decompose the optical depth field $\tau(\mathbf{x})$ into a spatial average $\bar{\tau}$, well measured by *Planck* through

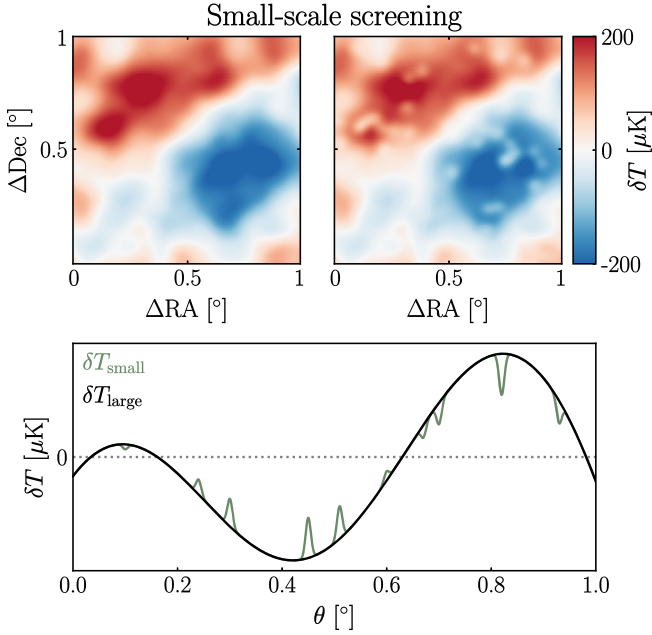


FIG. 1. Top: 1 deg^2 simulated map of a CMB realization unscreened (left) and screened by $5'$ FWHM 2D Gaussian τ profiles placed at CMASS galaxy positions (right). The τ amplitude is greatly exaggerated for visual effect. Bottom: A 1D schematic illustrating the small angular scale temperature deviations δT_{small} induced by the patchy screening effect (green). Unlike any other foreground, these deviations are (anti-)correlated with the large-scale primary CMB anisotropies δT_{large} (black). The tSZ- and CIB-induced anisotropies always have the same sign at a given frequency, lensing is correlated with $\nabla \delta T_{\text{large}}$, not δT_{large} itself, and kSZ-induced anisotropies change sign with the gas's LOS velocity. Thus, stacking the observed CMB temperature patches around galaxies weighted by the sign of the large-scale temperature provides a means to detect screening. Moreover, the patchy screening effect has greater magnitude where δT_{large} is large, motivating weighting by δT_{large} (and not just its sign) when stacking.

the large-scale scattering effect ($\bar{\tau} = 0.054 \pm 0.007$ [37]), and a fluctuating part $\delta\tau$, which traces the inhomogenous projected distribution of free electrons on the sky. Reabsorbing the mean $\bar{\tau}$ into the definition of δT^0 and assuming small fluctuations $\delta\tau(\mathbf{x}) \ll 1$, the screening effect reduces to an additive term:

$$\delta T^{\text{screening}}(\mathbf{x}) = -\delta T^0(\mathbf{x})\delta\tau(\mathbf{x}). \quad (4)$$

This effect is illustrated in Fig. 1. Our goal in this paper is to propose new “temperature inversion” and “signed” estimators in order to reconstruct the optical depth fluctuation field from the observed temperature map. Please note that in the discussion that follows, we use notations τ and $\delta\tau$ interchangeably, and they both represent the fluctuations on top of the mean $\bar{\tau}$ field.

III. STANDARD QUADRATIC ESTIMATOR

A. Review: Derivation of the standard QE

The standard QE for τ [16,32,33] is generally derived starting from

$$\delta T(\mathbf{x}) = \delta T^0(\mathbf{x})[1 - \tau(\mathbf{x}) + \mathcal{O}(\tau^2(\mathbf{x}))]. \quad (5)$$

In Fourier space with a flat-sky approximation, this becomes

$$\delta T_{\ell} \simeq \delta T_{\ell}^0 - \int_{\ell'} \delta T_{\ell'}^0 \tau_{\ell-\ell'}, \quad (6)$$

where $\int_{\ell'} \equiv \int \frac{d^2\ell'}{(2\pi)^2}$ and ℓ is the 2D Fourier wave vector, which we will also refer to interchangeably as angular multipole in what follows, in analogy with curved sky analyses. Thus, screening produces nonzero off-diagonal covariances for the Fourier modes,

$$\langle \delta T_{\ell} \delta T_{L-\ell} \rangle_{\text{fixed } \tau} = f^{\tau_L}(\ell, L-\ell) \tau_L, \quad (7)$$

where

$$f^{\tau_L}(\ell, L-\ell) = -(C_{\ell}^0 + C_{|L-\ell|}^0). \quad (8)$$

An elementary QE for τ_L is thus simply $-\delta T_{\ell} \delta T_{L-\ell} / (C_{\ell}^0 + C_{|L-\ell|}^0)$, with variance $C_{\ell}^{\text{total}} C_{|L-\ell|}^{\text{total}} / (C_{\ell}^0 + C_{|L-\ell|}^0)$. The standard QE for τ_L is therefore obtained from the inverse-variance linear combination of the elementary QEs, giving

$$\begin{aligned} \hat{\tau}_L^{\text{QE}} &\equiv N_L \int_{\ell} F^{\tau_L}(\ell, L-\ell) \delta T_{\ell} \delta T_{L-\ell} \\ &= - \frac{\int_{\ell} \frac{(C_{\ell}^0 + C_{|L-\ell|}^0)}{C_{\ell}^{\text{total}} C_{|L-\ell|}^{\text{total}}} \delta T_{\ell} \delta T_{L-\ell}}{\int_{\ell} \frac{(C_{\ell}^0 + C_{|L-\ell|}^0)^2}{C_{\ell}^{\text{total}} C_{|L-\ell|}^{\text{total}}}}, \end{aligned} \quad (9)$$

where

$$F^{\tau_L}(\ell, L-\ell) = \frac{f^{\tau_L}(\ell, L-\ell)}{C_{\ell}^{\text{total}} C_{|L-\ell|}^{\text{total}}}, \quad (10)$$

and

$$(N_L)^{-1} = \int_{\ell} F^{\tau_L}(\ell, L-\ell) f^{\tau_L}(\ell, L-\ell), \quad (11)$$

where the weights $F^{\tau_L}(\ell, L-\ell)$ are determined such that the variance of Eq. (9), $\mathcal{N}_{\mathcal{L}}^{\tau}$, given below, is minimized.

$$\mathcal{N}_{\mathcal{L}}^{\tau} = 2(N_L)^2 \int_{\ell} (F^{\tau_L}(\ell, L-\ell))^2 C_{\ell}^{\text{total}} C_{|L-\ell|}^{\text{total}}. \quad (12)$$

B. Intuition: Large-scale vs small-scale τ reconstruction

The standard QE derived above [Eq. (9)] can be understood more simply in two limiting cases, where large-scale and small-scale Fourier modes are being reconstructed.

In the former limit, where large-scale τ_L Fourier modes are reconstructed from small-scale temperature modes δT_ℓ , with $L \ll \ell$, the standard QE is effectively looking at the local power spectrum and searching for a change of the form

$$C_\ell \leftarrow C_\ell e^{-2\delta\tau}, \quad \text{i.e.,} \quad \frac{d \ln C_\ell}{d\tau} = -2. \quad (13)$$

This is illustrated in Fig. 2.

In the latter limit, where we reconstruct the optical depth field τ_L on small scales, i.e., $L \gg \ell$, we can approximate the power spectra as smoothly varying functions; $C_{|L-\ell|}^0 \simeq C_L^0$ and $C_{|L-\ell|}^{\text{total}} \simeq C_L^{\text{total}}$ to lowest order. Furthermore, the CMB power spectrum is very red, such that $C_\ell^0 \gg C_L^0$. In this limit, the standard QE takes the suggestive form,

$$\hat{\tau}_L^{\text{QE}} \simeq - \frac{\int_\ell \left(\frac{C_\ell^0}{C_\ell^{\text{total}}} \delta T_\ell \right) \delta T_{L-\ell}}{\int_\ell \frac{(C_\ell^0)^2}{C_\ell^{\text{total}}}}, \quad (14)$$

or in configuration space:

$$1 - \hat{\tau}^{\text{QE}}(\mathbf{x}) \simeq \delta T(\mathbf{x}) \frac{\delta T^{\text{large-WF}}(\mathbf{x})}{\langle (\delta T^{\text{large-WF}})^2 \rangle}. \quad (15)$$

In this limit, the standard QE is thus simply multiplying the local small-scale temperature $\delta T(\mathbf{x})$ by the Wiener-filtered large-scale temperature $\delta T^{\text{large-WF}}$, and normalizing it by the expected variance of the Wiener filter. The significance of this expression will become clear in Sec. IV, where we connect it to the new TI and signed estimators.

IV. “TEMPERATURE-INVERSION” ESTIMATOR & APPROXIMATE SMALL-SCALE QE

A. New “temperature-inversion” estimator

The optical depth fluctuation field possesses power on very small scales since it traces galaxy halos, whereas the primary CMB temperature has almost none due to Silk damping. Denoting small and large scales with subscripts S and L respectively, the screening effect on small scales is thus almost exclusively sourced by the following configurations:

$$\delta T_S^{\text{screening}}(\mathbf{x}) = -\delta T_L^0(\mathbf{x}) \delta \tau_S(\mathbf{x}). \quad (16)$$

This separation of scales is crucial; thanks to it, both the unscreened and screened temperature fluctuations can be measured from the same temperature map [27].

Indeed, we can low- and high-pass filter the observed temperature map to estimate the large- and small-scale temperature fluctuations $\delta \hat{T}_L(\mathbf{x})$ and $\delta \hat{T}_S(\mathbf{x})$. A natural, if somewhat singular, estimator for τ is then obtained by simply taking the ratio:

$$\hat{\tau}^{\text{TI}}(\mathbf{x}) \equiv - \frac{\delta \hat{T}_S(\mathbf{x})}{\delta \hat{T}_L(\mathbf{x})}. \quad (17)$$

We call this the TI estimator in analogy with the gradient inversion (GI) estimator of CMB lensing [30,31]. A natural




FIG. 2. Schematic illustration of the effect of large-scale optical depth fluctuations on the small-scale power spectrum: the locally observed power spectrum is simply modulated in amplitude. Top: A large-scale optical depth fluctuation $\delta\tau$ spanning 10° in RA. Middle: A $10^\circ \times 5^\circ$ simulated temperature map of a CMB realization screened by the large-scale $\delta\tau$ fluctuation shown in the top panel. The dotted (dashed) square corresponds to the area with suppressed (enhanced) anisotropies with respect to the full-sky average optical depth $\bar{\tau}$. Bottom: The power spectra corresponding to the two regions where $\delta\tau > 0$ (dotted) and $\delta\tau < 0$ (dashed). The power spectrum for the CMB screened only by $\bar{\tau}$ (i.e. $\delta\tau = 0$, solid) is shown for comparison. The screening of the CMB by a large-scale mode of the optical depth modulates the locally measured power spectrum as $C_\ell \leftarrow C_\ell e^{-2\delta\tau}$. The standard QE for $\delta\tau_L$ is looking precisely for this effect when reconstructing from temperature modes δT_ℓ with $L \ll \ell$.

choice of low-pass filter in this case would be a Wiener filter, the linear estimator for $\delta T_L^0(\mathbf{x})$ that has the lowest mean-squared error. A natural choice for the high-pass filter would be an inverse-variance weighted average of the small scales, the unbiased linear estimator of $\delta \hat{T}_S(\mathbf{x})$ with the lowest variance. In general, we can define linear filters W_L and W_S such that,

$$\begin{cases} \delta \hat{T}_L(\mathbf{x}) = \int d\mathbf{x}' W_L(\mathbf{x} - \mathbf{x}') \delta T(\mathbf{x}), \\ \delta \hat{T}_S(\mathbf{x}) = \int d\mathbf{x}' W_S(\mathbf{x} - \mathbf{x}') \delta T(\mathbf{x}). \end{cases} \quad (18)$$

In practice, we apply low- and high-pass filters (see Sec. VI A) to the temperature map to isolate δT_L and δT_S , respectively. If the shape of the τ profile for the objects of interest is known (e.g. from kSZ, which has higher SNR), we can use a matched filter with this profile to estimate the amplitude of the signal from the data. However, our filters, with variable size, allow us to actually measure the scale dependence of the electron density profile. When jointly analyzed, including their covariance matrix, they are as statistically optimal as the matched filter.

In Appendix A, we further derive the filter W_L which minimizes the mean-squared error in the presence of foreground contamination with known profiles (a combination of matched filter and Wiener filter), but do not implement it in what follows.

B. Comparison to the approximate small-scale QE

In the limit where we are reconstructing small-scale optical depth fluctuations, we have shown that the QE takes the approximate form,

$$1 - \hat{\tau}^{\text{QE}}(\mathbf{x}) \simeq \delta T(\mathbf{x}) \frac{\delta T^{\text{large-WF}}(\mathbf{x})}{\langle (\delta T^{\text{large-WF}})^2 \rangle}, \quad (19)$$

where $\delta T^{\text{large-WF}}$ is the Wiener filter estimator for the unscreened CMB, given the map noise. This can be contrasted with our TI estimator of the form,

$$1 - \hat{\tau}^{\text{TI}}(\mathbf{x}) \simeq \frac{\delta T(\mathbf{x})}{\delta T^{\text{large-WF}}(\mathbf{x})}. \quad (20)$$

In other words, in the small-scale limit, the two estimators are related via

$$1 - \hat{\tau}^{\text{QE}}(\mathbf{x}) = [1 - \hat{\tau}^{\text{TI}}(\mathbf{x})] \frac{(\delta T^{\text{large-WF}}(\mathbf{x}))^2}{\langle (\delta T^{\text{large-WF}})^2 \rangle}. \quad (21)$$

Thus, even if the large-scale temperature fluctuation $\delta T^{\text{large-WF}}$ is measured with infinite precision, the QE suffers from an irreducible variance due to the ratio between $(\delta T^{\text{large-WF}}(\mathbf{x}))^2$ and its mean. This is completely analogous to the GI estimator of CMB lensing. In Appendix B, we expand on this comparison, showing that

the SNR is higher for TI than the QE when the large-scale temperature leg is measured at high SNR (already true for *Planck* data) and the small-scale temperature leg is dominated by the screening effect. This latter condition is likely not satisfied; even in the absence of detector noise, and assuming that all nonblackbody foregrounds have been subtracted, kSZ and CMB lensing will likely dominate the noise unless they can be largely subtracted.

If TI does not outperform the QE in terms of SNR when considering the same set of temperature Fourier modes, why consider it further? As we show in Secs. VB and VD, TI is more robust to extragalactic foreground biases and naturally robust to CMB lensing bias. This means that TI can use more temperature Fourier modes (higher ℓ_{max}^T) and thus potentially recover a higher total SNR while still being as or more robust to biases.

C. Stacked TI estimator for cross-correlation and comparison with small-scale QE

For a single galaxy, the naïve TI estimator [Eq. (17)] is clearly singular for lines of sight where the unscreened temperature fluctuation is zero. This issue is automatically addressed when stacking around a set of galaxies, as we now show. The large-scale temperature $\delta \hat{T}_L(\mathbf{x})$ can be measured to high precision from the temperature map, such that we ignore the noise on $\delta \hat{T}_L(\mathbf{x})$ for now. On the other hand, the small-scale temperature $\delta \hat{T}_S(\mathbf{x})$ is typically affected by a non-negligible noise n_S (with variance $\sigma_{n_S}^2$) such that,

$$\hat{\tau}^{\text{TI}}(\mathbf{x}) = \tau(\mathbf{x}) + \frac{n_S(\mathbf{x})}{\delta \hat{T}_L(\mathbf{x})}. \quad (22)$$

In this case, the estimator is unbiased, and its noise is smallest along lines of sight with the highest large-scale temperature, as expected intuitively.

To extract the mean optical depth profile of a sample of tracers (galaxies, galaxy groups or clusters), we look for the unbiased linear combination of the $\hat{\tau}^{\text{TI}}$ around each cluster that has the smallest noise variance. The answer is the inverse-variance weighted average:

$$\hat{\tau}^{\text{TI stack}} \equiv \frac{\sum_i \hat{\tau}_i^{\text{TI}} \times \delta \hat{T}_{Li}^2 / \sigma_{n_S}^2}{\sum_i \delta \hat{T}_{Li}^2 / \sigma_{n_S}^2}, \quad (23)$$

which can be rewritten as

$$\hat{\tau}^{\text{TI stack}} = - \frac{\sum_i \delta \hat{T}_{Si} \delta \hat{T}_{Li} / \sigma_{n_S}^2}{\sum_j \delta \hat{T}_{Lj}^2 / \sigma_{n_S}^2}. \quad (24)$$

Conveniently, the inverse-variance weighting regularizes the singularity of the naïve estimator, such that lines of

sight where the unscreened CMB temperature $\delta T^{\hat{0}}$ is estimated to vanish are given zero weight in the sum.

Although the numerator of Eq. (24) is now quadratic in the temperature map, the estimator as a whole is not since the denominator also involves a quadratic combination of the data. As derived in Appendix C, the stacked QE takes a very similar form in the limit of small-scale screening:

$$\hat{\tau}^{\text{QE stack}} = -\frac{\sum_i \delta \hat{T}_{Si} \delta \hat{T}_{Li} / \sigma_{n_s}^2}{\sum_j \sigma_{T_L}^2 / \sigma_{n_s}^2}. \quad (25)$$

The numerator is identical to that of the stacked TI, and the denominator only differs by replacing $\delta \hat{T}_{Lj}^2$ with $\sigma_{T_L}^2$. Thus, the ratio of the stacked QE to the stacked TI is simply

$$\frac{\hat{\tau}^{\text{QE stack}}}{\hat{\tau}^{\text{TI stack}}} = \frac{\sum_j \delta \hat{T}_{Lj}^2 / \sigma_{n_s}^2}{\sum_j \sigma_{T_L}^2 / \sigma_{n_s}^2}. \quad (26)$$

This ratio has mean 1, but its variance is nonzero, equal to $2/N_{\text{galaxies}}$. Thus, in the limit of zero temperature noise, the stacked QE retains an irreducible variance $2\sigma_{\tau}^2/N_{\text{galaxies}}$, whereas the stacked TI yields a noiseless reconstruction of the true τ . While this variance vanishes as $\propto 1/N_{\text{galaxies}}$ for large galaxy samples, the other sources of variance (e.g., from detector noise) also share the same scaling, such that this additional noise remains relevant for all sample sizes.

D. “Signed estimator”: Suboptimal to TI, but potentially even more robust to foregrounds

The key intuition behind the derivation of the stacked TI estimator [Eq. (24)] is shown in Fig. 1; unlike all other imprints of galaxies on the CMB, the screening signal changes sign and amplitude depending on those of the local large-scale temperature δT_L . Discarding the amplitude information and weighting only by the sign of δT_L leads to a different screening estimator. We name it the “signed” estimator:

$$\hat{\tau}^{\text{sgn}} \equiv \frac{\sum_i \text{sgn}(\delta T_{L,i}) \delta T_{S,i}}{\sum_i |\delta T_{L,i}|}, \quad (27)$$

where $\delta T_{L,i}$ (or $\delta T_{S,i}$) is the large-scale (or small-scale) temperature fluctuation at the position of galaxy i .

While the TI and signed estimators are both unbiased estimators of τ , their variances differ such that,

$$\frac{\text{var} \hat{\tau}^{\text{sgn}}}{\text{var} \hat{\tau}^{\text{TI}}} = \frac{\langle (\delta T_{L,i})^2 \rangle}{\langle |\delta T_{L,i}| \rangle^2} = \frac{\pi}{2}, \quad (28)$$

for the Gaussian distributed large-scale temperatures δT_L . Thus, the signed estimator has lower SNR than TI by a factor of $\sqrt{\pi/2} = 1.25$.

On the other hand, this signed estimator enables two interesting avenues for controlling potential foregrounds, which may be worth the cost in SNR:

(i) **Sign balancing.** The number of galaxies coinciding with either positive or negative CMB patches will follow a binomial distribution, and thus will not perfectly cancel for a finite galaxy sample. We can discard a very small number of galaxies on CMB patches with the excess sign to ensure that $\sum_i \text{sgn}(\delta T_{L,i}) = 0$. Performing this “sign balancing” step exactly nulls any part of the foreground emission (e.g., CIB, tSZ) that is constant across galaxies. This does not lead to any bias in the estimated screening since this selection depends on the large-scale temperature only, statistically independent from the screening signal. Furthermore, this typically only requires discarding a fraction $\sim 1/N_{\text{gal}}$ of the galaxy sample, completely negligible for the large galaxy catalogs of interest here. This is described in more detail in Appendix D.

(ii) **Sign thresholding.** After masking bright point sources and clusters from the temperature map, any residual foreground emission from our galaxies will be small ($\sim \mu\text{K}$) compared to the typical large-scale temperature fluctuations ($\sim 100 \mu\text{K}$ for the CMB). Thus, we may select for the stack only the objects for which $|\delta T_{L,i}|$ is larger than a few times the typical foreground signal [e.g., ($\sim 10 \mu\text{K}$)]. This ensures that residual foregrounds are not large enough to affect the sign weighting used in the signed estimator. This is in contrast with TI, where residual foregrounds in $\delta T_{L,i}$, however small, remain present. This is described in more detail in Appendix E.

Including tests on simulations for biases from extragalactic sources, this estimator is presented in detail in [38], where we test for foreground biases and validate these arguments by applying the signed estimator to realistic, non-Gaussian sky simulations.

V. BIASES TO QE AND TI FROM CMB LENSING AND EXTRAGALACTIC FOREGROUNDS

A. CMB lensing bias to QE: Lens hardening needed

From Eq. (9), we can see that the QE for τ is of the form,

$$\hat{\tau}_L^{\text{QE}} \equiv -N_L \int_{\ell} \frac{(C_{\ell}^0 + C_{|L-\ell|}^0)}{C_{\ell}^{\text{total}} C_{|L-\ell|}^{\text{total}}} \delta T_{\ell} \delta T_{L-\ell}. \quad (29)$$

This has a similar form as that of the QE for the lensing potential; $\hat{\phi}_L^{\text{QE}} \propto \int_{\ell} F(\ell, L-\ell) \delta T_{\ell} \delta T_{L-\ell}$, where $F(\ell, L-\ell)$ are optimized weights for the pair of multipoles δT_{ℓ} and $\delta T_{L-\ell}$, such that the variance of the estimator is minimized. Thus, Eq. (9) can be interpreted as an estimator for the CMB lensing potential reconstruction

TABLE I. The amplitude of the cross-correlation of galaxies with the screening QE suffers large fractional biases from CMB lensing and tSZ, using $30 < \ell < 3000$. Appendix F 2 describes the “low” and “high” tSZ modeling in detail.

Galaxy sample	Fore-ground	SPT	ACT	SO	CMB-S4
CMASS	κ	410.9%	346.9%	394%	433.5%
	tSZ low	-2.7%	-2.5%	-2.7%	-2.8%
	tSZ high	-39.7%	-28.4%	-36.5%	-44.2%

with nonoptimized weights, and as such, has a nonzero response to lensing. We can rewrite Eq. (29) as

$$\langle \hat{\tau}_L^{\text{QE}} \rangle = \tau_L^{\text{QE}} + N_L \underbrace{\left(\int_{\mathcal{L}} F^r(\boldsymbol{\ell}, \mathbf{L} - \boldsymbol{\ell}) f_{TT}^{\phi}(\boldsymbol{\ell}, \mathbf{L} - \boldsymbol{\ell}) \right)}_{\text{Response of QE to lensing}} \phi(\mathbf{L}), \quad (30)$$

where we have used $\langle \delta T_{\ell} \delta T_{L-\ell} \rangle = f_{TT}^{\phi}(\boldsymbol{\ell}, \mathbf{L} - \boldsymbol{\ell}) \phi(\mathbf{L})$ with $f_{TT}^{\phi}(\boldsymbol{\ell}, \mathbf{L} - \boldsymbol{\ell}) = C_{\ell}^0(\mathbf{L} \cdot \boldsymbol{\ell}) + C_{|\mathbf{L}-\boldsymbol{\ell}|}^0(\mathbf{L} \cdot (\mathbf{L} - \boldsymbol{\ell}))$, and $F^r(\boldsymbol{\ell}, \mathbf{L} - \boldsymbol{\ell})$ is given by Eq. (10). In practice, the lensing bias in Eq. (30) is nonzero and can be dominant even in cross-correlations with other tracers, as shown in Table I and Fig. 3 (see also Fig. 1 of [39]).

This large lensing bias to the QE is typically avoided via “lens hardening” [40]. The idea behind the lens-hardened QE is simple. Just as we can construct a $\hat{\tau}$ QE, we can also construct a lensing convergence $\hat{\kappa}$ QE, which can then be used to subtract off the lensing bias in Eq. (30). Therefore,




FIG. 3. The expected cross-correlation signal between BOSS CMASS galaxies and the screening QE from Advanced ACT-like temperature (black) is compared to the CMB lensing bias (red) and the tSZ foreground bias (cyan area) for a range of tSZ models. The importance of these biases highlight the need for lens hardening and foreground hardening for the τ QE. The corresponding biases to the overall amplitude of the signal are shown in Table I. The simulations used to create these correlated lensing and tSZ maps are described in detail in Appendix F.

the lens-hardened τ QE minimizes contamination from lensing. However, this comes at an increased reconstruction noise cost when compared with the standard QE variance $\mathcal{N}_{\ell}^{\tau}$ given by Eq. (12). For a *Planck*-like experiment, this noise cost is ~ 0 –40% depending on scales [39], whereas the reconstruction noise from temperature-only (no polarization) QE increases by a factor of ~ 2 –3 across all ℓ values for a CMB-S4-like experiment [41].

Intuitively, this lensing bias to QE can be understood from Fig. 2 when reconstructing large-scale optical depth fluctuations τ_L from small-scale temperature modes T_{ℓ} with $L \ll \ell$. Both lensing and screening cause variations in the local power spectrum from patch to patch. While these variations are different, they have nonzero overlap.

The opposite regime, where we reconstruct small-scale optical depth fluctuations τ_L from temperature modes T_{ℓ} and $T_{L-\ell}$ with $L \gg \ell$, is different. In this limit, the lensing term in Eq. (30) vanishes. Thus, the small-scale part of the QE is unbiased with respect to lensing. This part of the QE is very similar to the TI estimator, only with a slightly different normalization [see Eqs. (19)–(21)], such that both are unbiased by lensing. We explore the robustness of TI to CMB lensing in more detail in the next subsection.

B. TI is robust to CMB lensing

The lensed temperature map can be written, to first order, as

$$\delta \tilde{T}(\mathbf{x}) \approx \delta T(\mathbf{x}) + \nabla \phi(\mathbf{x}) \nabla \delta T(\mathbf{x}), \quad (31)$$

where \tilde{T} is the lensed temperature field and ϕ is the lensing potential. Thus, lensing by a galaxy of interest produces a local small-scale dipole in the CMB map, whose direction and amplitude matches those of the local large-scale unlensed CMB gradient. The spatial symmetry of this dipole effect means that lensing is nulled by our azimuthally symmetric small-scale temperature filter, typically a disk or ring aperture around the galaxy, which only extracts the local small-scale monopole. Thus, to first order, lensing from the galaxies in our sample does not bias the TI or signed estimators and does not even contribute noise.

The situation is different for the lensing caused by objects not aligned with the aperture filter. For a single galaxy, nearby objects (including other galaxies in the sample) can contribute a lensing dipole to the aperture photometry. Since these small-scale dipoles are no longer aligned with our small-scale apertures (rings or disks), lensing does enter the TI and signed estimators. This lensing is not a bias, however, as it is correlated with $\nabla \delta T_L$, a random variable uncorrelated with δT_L or δT_S . Since there is an equal number of galaxies on either side of the gradient on average, lensing from objects not aligned with the filter only contributes noise to the TI and signed estimators. In fact, on small scales $\ell \gtrsim 5000$, the CMB power spectrum is dominated by lensing, meaning that

lensing may be the dominant source of noise (other than kSZ, other foregrounds and detector noise potentially present in the map).

In summary, TI and signed estimators null the lowest order bias due to lensing $\sim \mathcal{O}(\phi g)$. At higher order in lensing [$\mathcal{O}(\phi^2)$], a bias to the TI and signed estimators may be present if:

- (i) Lensing is non-negligible in the large-scale temperature map as well;
- (ii) This lensing is not due to the galaxy we are centered on, such that the corresponding dipole is not nulled by our small-scale aperture filter;
- (iii) The lensing is still caused by objects correlated with the galaxy we are centered on, e.g., two-halo term or other correlated galaxies in the sample.

We leave the exploration of these higher-order bias terms in simulations to future work.

C. Extragalactic foregrounds bias QE

Similar to CMB lensing, foregrounds like tSZ, kSZ and CIB cause mode coupling of temperature anisotropies. As a result, they bias the QE for screening. Again, similar to CMB lensing, this bias also affects cross-correlation of the reconstructed τ field with other tracers. Figure 3 shows the expected cross-correlation of the τ field reconstructed from a QE using an Advanced ACT-like experiment with CMASS galaxies. Here, ‘‘tSZ high’’ and ‘‘tSZ low’’ are two tSZ maps we simulate which give an upper and lower bound respectively on this bias (see Appendix F for details). We can see that the bias from CMB lensing is higher than the signal itself. The bias from tSZ emission is not negligible either. While we only discuss the foreground bias due to CMB lensing and tSZ here, foregrounds like CIB could also have non-negligible bias to τ QE. Such foreground biases have been extensively studied and shown to be important in the context of CMB lensing QE [42,43]. Similar conclusions hold for the τ QE as well due to the similarity between the two estimators. It should also be noted that these biases are likely bigger for lower-noise CMB experiments as they will have increased weights from small scales where the foregrounds are dominant.

In order to quantify the bias from foregrounds like CMB lensing and tSZ to the τ QE in cross-correlation with galaxies, we define an amplitude parameter A for a given power spectrum measurement \hat{C}_L such that,

$$\hat{C}_L = AC_L^{\text{theory}} + n_L, \quad (32)$$

where C_L^{theory} is the theory power spectrum and n_L is noise. The minimum-variance estimator for A , linear in the measured power spectrum, is then

$$\hat{A} = \frac{\sum_L \hat{C}_L (C_L^{\text{theory}}) / \sigma_L^2}{\sum_L (C_L^{\text{theory}})^2 / \sigma_L^2}. \quad (33)$$

The nominal value of \hat{A} is 1; however, due to biases from foregrounds, the value differs from 1. Thus, bias from a foreground X to cross-correlation of the QE reconstructed τ field with galaxies δ_g ($C_L^{\hat{\tau} \times \delta_g}$) can be calculated in the continuum limit as

$$\text{bias}(C_L^{\hat{\tau} \times \delta_g}) = \frac{\int \frac{d^2 \mathbf{L}}{(2\pi)^2} \text{bias}(C_L^{X \times \delta_g}) \frac{(C_L^{\hat{\tau} \times \delta_g})^2}{\sigma_L^2}}{\int \frac{d^2 \mathbf{L}}{(2\pi)^2} \frac{(C_L^{\hat{\tau} \times \delta_g})^2}{\sigma_L^2}}, \quad (34)$$

where $\text{bias}(C_L^{X \times \delta_g})$ is the bias introduced by the cross-correlation between the galaxies and the foreground X passed through a τ QE. Table I shows that these foreground biases can be larger than the screening signal and highly significant. Some form of foreground mitigation is therefore likely required for the QE. This may include multi-frequency component separation, masking, template deprojection, using shear only estimators, and bias hardening with respect to point sources or specific profiles. Such approaches and their challenges have been explored in the context of CMB lensing and τ QE [39,40,44–46] and similar results hold here too. We do not explore these further in this paper.

D. TI is robust to extragalactic foregrounds

The TI and signed estimators are automatically robust to extragalactic foreground biases, and this can be further enhanced by sign balancing and sign thresholding for the signed estimator. We summarize the reasons here and present more detailed derivations in Appendixes D and E.

Unlike the patchy screening effect, foregrounds like the CIB, tSZ and kSZ do not change sign depending on the sign of the local large-scale temperature δT_L . In the absence of foreground contamination to the large-scale leg, the weights in the TI and signed estimators, which change sign and have zero mean value, automatically average out these foregrounds. For a finite galaxy sample size, this foreground reduction may leave a residual bias $\sim f_s / \sqrt{N_{\text{patch}}}$, where f_s is the typical small scale foreground emission, and N_{patch} is the number of patches with independent large-scale CMB temperature δT_L . Since the CMB is coherent on degree scales, N_{patch} is roughly equal to the area of the CMB map in square degrees.

For the signed estimator, this residual can be further reduced by sign balancing the weights, as described in Sec. IV D. If the foreground signal is identical for all galaxies in the sample, sign balancing exactly nulls the foreground bias. While foreground emissions actually vary across galaxies, sign balancing is still effective in realistic scenarios, as shown in Appendix D.

In practice, foreground contamination may be present both in the large and small-scale temperature legs. In this case, the TI and signed estimators may suffer a bispectrum

bias $\sim f_L f_S g_S$, where f is the foreground and g the galaxy map. However, for the signed estimator, this bias can be nulled exactly with the sign thresholding procedure presented in Sec. IV D. This procedure ensures that absolutely no foreground signal leaks into $\text{sgn}(\delta T_L)$, thus removing any bispectrum foreground bias.

VI. SNR FORECASTS

We produce SNR forecasts for the TI and τ QE estimators using realistic simulations for four CMB experiments, South Pole Telescope (SPT) [47–49], Atacama Cosmology Telescope (ACT) [50], SO [51] and CMB-S4 [52], cross-correlated with three galaxy survey samples, the Baryonic Oscillation Spectroscopic Survey (BOSS) CMASS sample [53,54], the Wide-field Infrared Survey Explorer (WISE) unWISE “blue” and “green” samples [55,56], and the simulated Legacy Survey of Space and Time (LSST) Data Challenge 2 (DC2) extragalactic sample [57]. Details of the sky fraction, galaxy number density and number of galaxies overlapping the CMB survey footprints are listed in Table II.

A. Simulated analysis for TI and QE

Here we give an overview of the methods used to generate the simulations for the TI and QE detection forecast. See Appendix G for further details. We create realistic simulated τ signal maps by placing identical, isotropic Gaussian τ profiles at CMASS galaxy positions as described in Appendix G 1. The parameters of the τ profiles are matched to the best-fit values from the ACT \times unWISE first measurement [38].

To calculate the covariance matrix, we simulate a set of 128 maps for each CMB experiment, using the power spectra shown in Fig. 4. All of the spectra are beamed and include noise contributions from the CMB, detector noise and—with the exception of CMB-S4—extragalactic foreground contamination (e.g. lensing, CIB, tSZ, kSZ, radio sources). More details may be found in Appendix G 2.

To separate the large- and small-scale temperature fluctuations δT_L and δT_S , we apply low- and high-pass filters (see Eqs. (G1) and (G2) to generate two separate maps. We have implemented the stacked TI and signed estimators in the aperture photometry filtering and stacking software, ThumbStack [11]. We use this code to generate

TABLE II. Number of galaxies assumed to overlap with each CMB experiment used in the SNR forecasts. The shot noise, \bar{n} , for each galaxy sample and the sky fraction, f_{sky} , for each CMB experiment are also listed. ACT, SO, and CMB-S4 are assumed to have equal f_{sky} .

Galaxy sample	\bar{n} arcmin ⁻²	SPT ($f_{\text{sky}} = 0.036$)	ACT/SO/CMB-S4 ($f_{\text{sky}} = 0.315$)
CMASS	0.012	65,600	568,776
unWISE	1.24	6.8M	59M
LSST	34.1	187M	1.62B




FIG. 4. Power spectra used in the SNR forecasts for ACT (solid orange), SPT (dashed green), SO (dash-dotted blue), and CMB-S4 (dotted purple). Details for the spectra are given in Appendix G. The shaded bands show the ℓ ranges used by the TI estimator; $30 < \ell \lesssim 2075$ (grey) and $2425 \lesssim \ell < 5950$ (green) for large- and small-scale temperature measurements, respectively. See Eqs. (G1) and (G2) for the exact ℓ filtering used. The QE uses $30 < \ell < 3000$ (black dotted line) as it is sensitive to foreground contamination at higher ℓ . Since the TI estimator is robust to foregrounds at high ℓ , it can take advantage of future experiments’ lower detector noise and smaller beam widths.

small postage stamps centered on each galaxy position for both maps, apply annular (i.e. “ring”) aperture photometry filters to measure the δT_S radial profile for each stamp, and stack the stamps according to Eq. (24), using each stamp’s average δT_L as weights.

We construct the covariance matrix taking into account the clustering of galaxies and scaling by the proper Hartlap correction factor [58]. Details are provided in Appendix G 4.

We then measure the SNR for each CMB-galaxy sample combination, scaling to each experiment’s full galaxy sample and sky fraction, f_{sky} , shown in Table II. The resulting SNR values for the surveys are shown in Table IV.

For the QE of τ , details of the SNR estimation for the cross-power spectrum of the τ map reconstructed using QE and galaxy map g are presented in Appendix H. The SNR is calculated analytically as given in Eq. (H1) where the noise terms contain the total power spectrum of the reconstructed τ map along with its reconstruction noise and the total galaxy power spectrum, which includes both the clustering and shot noise contributions. The shot noise levels for each galaxy sample are calculated using the number density values provided in Table II.

B. Galaxy clustering effects

As we go towards higher density galaxy samples, the shot noise of the sample decreases by $\sqrt{N_{\text{gal}}}$, but the clustering power spectrum still adds noise to the covariance. Here we assess the impact of clustering effects for the TI and QE forecasts.

TABLE III. Summary of the detection SNRs for TI and QE varying the noise contribution due to galaxy clustering, gg . The first value in each pair corresponds to the case with no clustering ($gg = 0$) power spectrum contribution to the total galaxy power spectrum; the second value includes it ($gg \neq 0$). For QE, we use $60 < \ell < 6000$. For TI, we use $30 < \ell \lesssim 2075$ for δT_L weights and $2425 \lesssim \ell < 5950$ for τ reconstruction. Ratios of the $gg = 0$ and $gg \neq 0$ SNR values are visualized in Fig. 5.

CMB Expt.	Estimator	$gg = 0/gg \neq 0$		
		CMASS	unWISE	LSST
SPT	TI ₅₉₅₀	0.34/0.37	3.4/3.0	18/11
	QE ₆₀₀₀	0.54/0.53	5.5/4.7	29/12
ACT	TI ₅₉₅₀	0.60/0.65	6.1/5.2	31/18
	QE ₆₀₀₀	1.1/1.1	11/9.3	58/22
SO	TI ₅₉₅₀	0.84/0.92	8.5/7.5	45/25
	QE ₆₀₀₀	1.4/1.4	14/12	73/29
CMB-S4	TI ₅₉₅₀	1.5/1.7	16/13	82/49
	QE ₆₀₀₀	2.7/2.7	28/25	145/65

For the QE, we complete the simulated analysis with and without the galaxy clustering (gg) contribution at the power spectrum level, where the total galaxy power spectrum is calculated with and without the clustering power spectrum. The results of this study are summarized in Table III.

For the TI estimator, we compare the SNR values obtained with a covariance matrix estimated from bootstrapping from a single map realization (thus neglecting clustering effects) versus that estimated from the sample variance over 128 map realizations.

As shown in Fig. 5, neglecting clustering effects results in overestimation of the SNR for both the QE and TI. While




FIG. 5. Ratio of the SNRs when accounting for galaxy shot noise only versus galaxy shot noise and clustering. For all three galaxy samples, the difference in SNR ratio between CMB experiments (same colors as Fig. 4) is relatively small. However, the impact of clustering on SNR increases significantly for higher density galaxy samples, especially LSST. When neglecting clustering in the LSST sample, the SNR is overestimated by $\sim 60\%$ for TI (solid lines) and $\sim 150\%$ for QE (dashed lines).

the CMASS sample is sparse enough to not show significant differences, using unWISE results in a moderate $\sim 20\%$ overestimated SNR. Ignoring clustering in the LSST sample is a more dramatic oversight; $\sim 60\%$ for TI and $\sim 150\%$ for QE. The marked differences in sensitivity to clustering effects in TI and QE, especially for an LSST-like experiment, is unclear. We leave an investigation of this difference to future work. Despite these uncertainties, it is clear clustering has a significant effect on noise estimation. Thus, we include these effects in the present forecasts and emphasize they must be properly accounted for in future analyses.

We also study the effect of neglecting the cosmic variance of the signal itself, i.e., $C_{\rho}^{\tau gg}$, and find it is negligible for all the experiments considered.

Figure 6 shows the comparison of the stacked high-pass filtered radial τ profile on the unWISE (orange) and CMASS (red) locations in our simulations where the τ profile has been reconstructed using the TI estimator with an ACT-like experiment. The input profile used when painting each galaxy in both maps is shown for comparison in black. The average τ profile stacked on CMASS galaxy locations matches the profile painted on an individual galaxy, i.e., the CMASS sample is sparse enough for the two-halo term to have negligible contribution. However, after stacking on




FIG. 6. Comparison of the high-pass filtered radial $\langle \tau \rangle$ profile when stacking on the unWISE (red) and CMASS (orange) τ maps with the TI estimator. The input profile used when painting each galaxy in both maps is shown for comparison (black). Note the shape of the profiles dips below zero due to the high-pass filtering applied. The τ profile stacked on CMASS galaxy locations matches the profile painted on an individual galaxy, i.e., the CMASS sample is sparse enough for the two-halo term to have negligible contribution. However, after stacking on the unWISE sample, the profile amplitude shows a slight excess, which reflects the two-halo term contribution. The error bars on the CMASS and unWISE profiles correspond to the expected noise from ACT. A tentative detection of this two-halo term may be possible with ACT \times unWISE.

TABLE IV. Summary of the forecast detection SNRs for TI and QE accounting for galaxy clustering and using realistic ℓ ranges for each estimator. We do not include cosmic variance on the τg signal, which is negligible for all experiments considered. For the standard QE and lens-hardened QE (LHQE), we use $60 < \ell < 3000$ where these estimators have limited response to foreground contamination. For TI, we use the same ℓ range as in Table III. TI can use higher ℓ modes as it is robust to foregrounds; the SNR for TI can thus exceed QE because of this added information. All SNR values are visualized in Fig. 7.

Galaxy sample	Estimator	SPT	ACT	SO	CMB-S4
CMASS	TI ₅₉₅₀	0.37 σ	0.65 σ	0.92 σ	1.7 σ
	QE ₃₀₀₀	0.34 σ	0.87 σ	0.96 σ	1.1 σ
	LHQE ₃₀₀₀	0.30 σ	0.78 σ	0.84 σ	0.94 σ
unWISE	TI ₅₉₅₀	3.0 σ	5.2 σ	7.5 σ	13 σ
	QE ₃₀₀₀	2.9 σ	7.3 σ	8.1 σ	9.2 σ
	LHQE ₃₀₀₀	2.6 σ	6.6 σ	7.2 σ	8.0 σ
LSST	TI ₅₉₅₀	11 σ	18 σ	25 σ	49 σ
	QE ₃₀₀₀	6.6 σ	16 σ	18 σ	21 σ
	LHQE ₃₀₀₀	5.9 σ	15 σ	16 σ	18 σ

the unWISE sample, there is a slight excess in the amplitude of the profile, which reflects the two-halo term contribution. This shows that the clustering term already starts becoming important for the unWISE-like galaxy sample and will be even more important for an LSST-like sample. A tentative detection of this two-halo term might be possible using ACT \times unWISE data.

C. Results

The forecast SNRs for TI, QE, and lens-hardened QE (LHQE), using realistic simulations and accounting for clustering effects as described above, are summarized in Table IV and visualized in Fig. 7. Unlike for the clustering study, we limit ℓ_{\max} to 3000 for the QE and LHQE to mitigate foreground contamination. As TI is robust to foregrounds, we use scales up to $\ell_{\max} = 5950$.³ As expected, we find that the LHQE degrades the SNR to some extent ($\sim 10\text{--}15\%$) when compared to the standard QE as the noise in the reconstruction of τ for a LHQE goes up. As this is the forecast SNR for cross-correlation of τ and galaxies $C_{\ell}^{\tau g}$ rather than auto-power spectrum of $\tau C_{\ell}^{\tau\tau}$, effects of the noise penalty for hardening against lensing are slightly less dramatic. Factors contributing to the noise [denominator of Eq. (H1)] are total galaxy and τ power spectra. As expected, we find that for a given galaxy sample, as we go towards more sensitive CMB experiments, effects of lens-hardening are felt slightly more as the noise penalty is higher for more sensitive experiments [41].

³This rather specific value corresponds to the most conservative ℓ_{\max} of the input ILC power spectra used in the simulations.




FIG. 7. Visualization of the forecast SNR values given in Table IV. The LHQE (dotted lines) results in a SNR reduction compared to the standard QE (dashed lines). For current experiments, the QE and TI (solid lines) give similar SNRs. Cross-correlation with CMASS (orange circles) does not yield a significant measurement for any CMB experiment, whereas the signal is detectable using the unWISE sample (red diamonds). For future experiments, TI can attain higher SNR than the QE, most dramatically when using the LSST galaxy sample (blue pluses), because it can use smaller CMB scales.

As a result, SNR is affected more by lens-hardening for a CMB-S4-like experiment than an ACT-like experiment. In the following discussion, we limit our comparisons to TI and QE.

For current experiments like ACT and SPT in combination with CMASS and unWISE galaxies, TI and QE give similar results. Due to fairly high shot noise, the cross-correlation with the CMASS-like sample does not yield a significant SNR for any CMB experiment, whereas the signal is detectable using an unWISE-like sample. We note this result is consistent with [38], which uses ACT \times unWISE data to make the first detection of this effect.

For future CMB experiments, the TI estimator can attain higher SNR than the QE, most dramatically when using an LSST-like galaxy sample (see Table IV and Fig. 7). This is due to the use of smaller scales of the CMB. Future CMB experiments, particularly CMB-S4, have exquisite projected sensitivity and beam size. With lower noise and a smaller beam, the TI estimator can extract information from the small scales that QE cannot reliably use due to foreground biases. Thus, the SNR values across all cross-correlations with LSST are improved for TI compared to the QE.

VII. CONCLUSIONS

In this paper, we have derived and studied new estimators for the patchy screening of the CMB, called TI and signed estimators. These estimators are modifications of the standard screening QE, analogous to the gradient inversion estimators for CMB lensing.

The TI and signed estimators make use of a unique feature of patchy screening: it is the only effect which changes sign based on the local large-scale temperature. This allows TI and signed estimators to be automatically unbiased to extragalactic foregrounds like the CIB, tSZ and kSZ, unlike the standard QE. This foreground robustness can be further enhanced for the signed estimator via “sign balancing” and “sign thresholding”, which further immunize it against foregrounds at a negligible noise cost. For the QE, the corresponding biases can be large and likely require some form of hardening at some cost in SNR.

The TI and signed estimators are also robust to CMB lensing bias, unlike the standard QE. Lens hardening remedies this issue for QE, again, with higher noise penalty.

We also clarify the relationship between the TI estimator, signed estimator and QE. We show that the QE reconstructs both small- and large-scale screening information and present simple intuition about these two regimes. We show how TI is a subtle modification to the small-scale-only QE, which in principle avoids the noise floor above. Finally, the signed estimator is a modification to TI which allows even more robustness to foregrounds (via sign balancing and sign thresholding), at the cost of a factor of $\sqrt{\pi/2} \approx 1.25$ reduction in SNR. Even in the absence of detector noise, foregrounds, and lensing, the τ QE has an irreducible source of variance and hits a noise floor, whereas the TI and signed estimators could perform arbitrarily better under such noiseless conditions. However, this noise floor is not reached in practice. For a fixed set of Fourier scales included in the analysis, the statistical SNR in the TI or signed estimators are similar to the QE but do not outperform it. However, by being robust to lensing and foregrounds, TI and signed estimators can be applied to more Fourier modes (pushing to smaller scales), thus potentially increasing the overall SNR.

We perform realistic forecasts for current and upcoming CMB and LSS experiments (see Table IV and Fig. 7). Our results highlight the large noise enhancement due to the clustering of galaxies once the galaxy number density reaches LSST-like values. Further, our SNR forecast for ACT×unWISE is consistent with the measurement found in the companion paper [38], which is the first detection of this signal.

We have focused on measuring the patchy screening from the CMB temperature only. The TI and signed estimators should work identically when substituting temperature T with Stokes Q or U polarization maps. We leave this exploration to future work.

In the future, patchy screening measurements around galaxies will be a powerful probe of the gas density profile around them, highly complementary with kSZ and tSZ effects. Joint analyses of these signals around the same galaxies will yield valuable insights into the thermodynamical properties of baryons within these halos and shed more light on the lesser understood properties of feedback

processes. This will have a direct impact on weak galaxy lensing studies where baryonic effects are a major systematic. They will also help in breaking the velocity-optical depth degeneracy in kSZ, thus expanding its power as a probe of the large-scale velocities to constrain the growth rate of structure, modifications to general relativity and local primordial non-Gaussianity. In particular, as the uncertainty on the kSZ ($\propto \tau v_{\text{LOS}}$) measurements will be negligible compared to that of patchy screening ($\propto \tau$), the SNR values presented here translate as the expected measurement SNR on the growth rate of structure ($\propto v_{\text{LOS}}$).

The `LensQuEst` and `ThumbStack` software packages are publicly available under an open source license.

ACKNOWLEDGMENTS

We thank Nicholas Battaglia, Simone Ferraro, Darby Kramer, Srinivasan Raghunathan, Rashid Sunyaev, Alexander van Engelen and Ziang Yan for useful discussions. We thank Joanna Dunkley, Lyman Page and Edward Wollack for helpful comments on the manuscript. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2146755. T. S. thanks the LSST Corporation (LSSTC) Data Science Fellowship Program, which is funded by LSSTC, NSF Cybertraining Grant No. 1829740, the Brinson Foundation, and the Moore Foundation; their participation in the program has benefited this work. This work received support from the U.S. Department of Energy under Contract No. DE-AC02-76SF00515 to SLAC National Accelerator Laboratory. Some of the computing for this project was performed on the Sherlock cluster. We thank Stanford University and the Stanford Research Computing Center for providing computational resources and support.

APPENDIX A: OPTIMAL δT_L ESTIMATOR IN THE PRESENCE OF FOREGROUND BIAS

In this section, we derive estimators of the large-scale CMB temperature around a given galaxy. To declutter the derivation, we simplify our notation as

$$d_\ell = T_\ell + n_\ell + Au_\ell, \quad (\text{A1})$$

where d_ℓ is the temperature data as a function of Fourier multipole ℓ , T_ℓ is the true CMB (with known power spectrum C_ℓ) to be estimated, n_ℓ is the noise (with known power spectrum N_ℓ), and Au_ℓ is a foreground emission, with a known profile u_ℓ (e.g., the beam for a point source) and unknown amplitude A . We would like to build an estimator \hat{T}_ℓ of T_ℓ , which minimizes the mean squared error ($\langle (\hat{T}_\ell - T_\ell)^2 \rangle$). We restrict ourselves to estimators linear in the data, i.e. we wish to determine the weights

W_L for each ℓ such that,

$$\hat{T}_\ell = \sum_L W_L d_{\ell-L}. \quad (\text{A2})$$

In the absence of foregrounds, i.e., if A is known to be zero, minimizing the mean-squared error leads to the well-known Wiener filter:

$$\hat{T}_\ell = \frac{C_\ell}{C_\ell + N_\ell} d_\ell. \quad (\text{A3})$$

The Wiener filter can also be derived as the maximum *a posteriori* estimator for T_ℓ , in the absence of foregrounds ($A = 0$), as well as in the presence of foregrounds for which the prior on the amplitude A is flat. Such foregrounds would have a random sign, which is typically not the case in our problem.

We now consider the case of nonzero foreground and impose the additional constraint that the estimator \hat{T}_ℓ should have zero response to the foreground amplitude:

$$\frac{d\hat{T}_\ell}{dA} = 0. \quad (\text{A4})$$

In what follows, we show that the linear estimator which minimizes the mean squared error, while having zero response to the foreground emission, is simply obtained by Wiener-filtering the data, after having subtracted the matched-filtered foreground template.

We introduce a Lagrange multiplier λ , to minimize the mean-squared error under the constraint,

$$\mathcal{L} = \langle (\hat{T}_\ell - T_\ell)^2 \rangle - \lambda \frac{d\hat{T}_\ell}{dA}. \quad (\text{A5})$$

In terms of the weights W_L , this loss function becomes

$$\begin{aligned} \mathcal{L} = & (W_0 - 1)^2 C_\ell + W_0^2 N_\ell + \sum_{L \neq 0} W_L^2 (C_{\ell-L} + N_{\ell-L}) \\ & + A^2 \sum_{L, L'} W_L W_{L'} u_{\ell-L} u_{\ell-L'} - \lambda \sum_L W_L u_{\ell-L}. \end{aligned} \quad (\text{A6})$$

Setting $\partial \mathcal{L} / \partial W_L$ to zero gives,

$$W_L = \frac{C_\ell}{C_\ell + N_\ell} \delta_{L,0}^K - (\alpha + \lambda) \frac{u_{\ell-L}}{C_{\ell-L} + N_{\ell-L}}, \quad (\text{A7})$$

where we have defined

$$\alpha \equiv A \sum_{L'} W_{L'} u_{\ell-L'}. \quad (\text{A8})$$

Setting $\partial \mathcal{L} / \partial \lambda$ to zero gives $\sum_L W_L u_{\ell-L} = 0$, which means that $\alpha = 0$ and lets us solve for λ in Eq. (A6):

$$\lambda = \frac{C_\ell}{C_\ell + N_\ell} \frac{u_\ell}{\left(\sum_L \frac{u_L^2}{C_L + N_L} \right)}. \quad (\text{A9})$$

Plugging this into Eq. (A6) yields the expressions for the weights:

$$W_L = \frac{C_\ell}{C_\ell + N_\ell} \left[\delta_{L,0}^K - \frac{\frac{u_{\ell-L}}{C_{\ell-L} + N_{\ell-L}}}{\left(\sum_{L'} \frac{u_{L'}^2}{C_{L'} + N_{L'}} \right)} u_\ell \right]. \quad (\text{A10})$$

In conclusion, the estimator for T_ℓ , linear in the data d_ℓ , which minimizes the mean squared error while having zero response to the foreground emission is

$$\hat{T}_\ell = \frac{C_\ell}{C_\ell + N_\ell} [d_\ell - \hat{A} u_\ell], \quad (\text{A11})$$

where \hat{A} is the matched filter estimator for the foreground amplitude A :

$$\hat{A} \equiv \frac{\left(\sum_L \frac{d_L u_L}{C_L + N_L} \right)}{\left(\sum_{L'} \frac{u_{L'}^2}{C_{L'} + N_{L'}} \right)}. \quad (\text{A12})$$

As stated, this reduces to subtracting the foreground template $\hat{A} u_\ell$ from the data, where \hat{A} is the matched filter for the profile u_ℓ , and then Wiener-filtering the result.

APPENDIX B: VARIANCE COMPARISON: SMALL-SCALE QE VS TI

We wish to get intuition about the noise properties of the QE and TI, in order to understand in which case one may be a better estimator than the other.

We decompose the observed small- and large-scale temperatures around a given galaxy into signal and noise:

$$\begin{cases} \delta \hat{T}_S = -\delta T_L \tau + n_S \\ \delta \hat{T}_L = \delta T_L + n_L \end{cases} \quad (\text{B1})$$

Here, the noise n_S includes everything in the small-scale map other than our screening effect, i.e. detector noise, foregrounds, lensing and the small, Silk-damped primary CMB. The large-scale noise n_L should be thought of as very small, since the primary CMB is very well-measured on degree scales.

To get a sense of the relevant noise regimes, we introduce the variances $\sigma_{T_L}^2$, σ_τ^2 , $\sigma_{n_L}^2$ and $\sigma_{n_S}^2$ of $\delta \hat{T}_L$, the true τ , n_L and n_S . In what follows, we assume $\sigma_{n_L}^2 \ll \sigma_{T_L}^2$, realistic for current and upcoming experiments, and that δT_L , τ , n_L and n_S are statistically independent.

We can now express the TI and QE estimators in real space as

$$\begin{cases} \hat{\tau}^{\text{TI}} = -\delta\hat{T}_S/\delta\hat{T}_L \\ \hat{\tau}^{\text{QE}} = -\delta\hat{T}_S\delta\hat{T}_L/\sigma_{T_L}^2 \end{cases} \quad (\text{B2})$$

We then plug in the expressions for the observed quantities Eq. (B1), in terms of the true ones, into the QE and TI estimators. As expected, we obtain a linear response in the true τ and a noise term:

$$\begin{cases} \hat{\tau}^{\text{TI}} = \alpha^{\text{TI}}\tau + n^{\text{TI}} \\ \hat{\tau}^{\text{QE}} = \alpha^{\text{QE}}\tau + n^{\text{QE}} \end{cases} \quad (\text{B3})$$

1. Irreducible noise in the QE response

For the TI estimator,

$$\alpha^{\text{TI}} = \frac{\delta T_L}{\delta T_L + n_L} = 1 - \frac{n_L}{\delta T_L} + \left(\frac{n_L}{\delta T_L}\right)^2 + \dots, \quad (\text{B4})$$

such that the mean response $\langle\alpha^{\text{TI}}\rangle$ differs from unity by a bias $\mathcal{O}(\sigma_{n_L}^2/\sigma_{T_L}^2)$. The variance of the response α^{TI} is similarly small, with

$$\text{var}(\alpha^{\text{TI}}) = \mathcal{O}(\sigma_{n_L}^2/\sigma_{T_L}^2). \quad (\text{B5})$$

Thus, in practice, the bias and the variance of the TI response are negligibly small, i.e. $\alpha^{\text{TI}} \simeq 1$.

In contrast, the linear response of the QE estimator is

$$\alpha^{\text{QE}} = \frac{\delta T_L(\delta T_L + n_L)}{\sigma_{T_L}^2}. \quad (\text{B6})$$

This response is exactly unbiased, i.e. $\langle\alpha^{\text{QE}}\rangle = 1$. However, its variance does not go to zero in the limit of low noise:

$$\text{var}(\alpha^{\text{TI}}) = 2 + \sigma_{n_L}^2/\sigma_{T_L}^2. \quad (\text{B7})$$

Thus, if the noises n^{TI} and n^{QE} are negligible, the TI estimator converges to a perfect reconstruction of the true τ , whereas the QE estimator converges to $\alpha^{\text{QE}}\tau$, where α^{QE} is a random number with variance 2.

2. Comparable estimator noises

The noise n^{TI} of the TI and QE estimators are

$$\begin{cases} n^{\text{TI}} = \frac{n_S}{\delta T_L + n_L} \\ n^{\text{QE}} = \frac{n_S(\delta T_L + n_L)}{\sigma_{T_L}^2} \end{cases} \quad (\text{B8})$$

Since $\langle 1/\delta T_L^2 \rangle \gtrsim 1/\sigma_{T_L}^2$ by convexity of $x \rightarrow 1/x$, we obtain,

$$\text{var}(n^{\text{TI}}) \gtrsim \frac{\sigma_{n_S}^2}{\sigma_{T_L}^2} \left(1 + \frac{\sigma_{n_L}^2}{\sigma_{T_L}^2}\right) = \text{var}(n^{\text{QE}}). \quad (\text{B9})$$

Thus, while the TI noise is larger than the QE noise, they are comparable.

3. Total variance: TI outperforms QE in the low-noise regime

To decide in which regime one estimator might be preferred over the other, we take into account all the contributions to the estimator variance, including the true signal τ , the noise $n^{\text{TI/QE}}$ and the response $\alpha^{\text{TI/QE}}$. This leads to

$$\text{var}(\hat{\tau}^{\text{TI/QE}}) = \text{var}(\alpha^{\text{TI/QE}})\sigma_\tau^2 + \text{var}(n^{\text{TI/QE}}), \quad (\text{B10})$$

such that we find:

$$\begin{cases} \text{var}(\hat{\tau}^{\text{TI}}) = \frac{\sigma_{n_L}^2}{\sigma_{T_L}^2}\sigma_\tau^2 + \mathcal{O}\left(\frac{\sigma_{n_S}^2}{\sigma_{T_L}^2}\right) \\ \text{var}(\hat{\tau}^{\text{QE}}) = \left(2 + \frac{\sigma_{n_L}^2}{\sigma_{T_L}^2}\right)\sigma_\tau^2 + \frac{\sigma_{n_S}^2}{\sigma_{T_L}^2}. \end{cases} \quad (\text{B11})$$

It thus appears that the TI is preferred over the QE when the ‘‘2’’ term is relevant, i.e. when

$$\sigma_{n_L}^2 \lesssim \sigma_{T_L}^2 \quad \text{and} \quad \sigma_{n_S}^2 \lesssim \sigma_{T_L}^2\sigma_\tau^2. \quad (\text{B12})$$

The first condition is easily satisfied, since the large scale temperature is measured at high SNR in current experiments. The second condition effectively requires that the n_S be smaller than the screening effect $\delta T_L\tau$ in the measured small-scale temperature. However, the noise n_S includes not only detector noise, which could in principle be made arbitrarily small, but also foregrounds and lensing. Thus, going beyond QE with TI requires significant foreground cleaning. Lensing is much larger than screening in the small-scale temperature power spectrum, which violates $\sigma_{n_S}^2 \lesssim \sigma_{T_L}^2\sigma_\tau^2$. Thus, the TI SNR should only outperform the QE’s if a large amount of delensing can be achieved. We leave quantifying this required amount of delensing to future work.

APPENDIX C: TWO VERSIONS OF THE APPROXIMATE SMALL-SCALE QE; SAME COMPARISON WITH TI

Here we give the expressions for the stacked TI and QE estimators. These show that the key difference between TI and QE remains true for the stacked estimators; the QE’s response to the true τ has an irreducible variance, even in the absence of any noise.

As we show in the main text, the individual TI estimator for galaxy i is $\tau_i^{\text{TI}} \equiv \delta\hat{T}_{Si}/\delta\hat{T}_{Li}$, with variance $\sigma_{n_S}^2/\delta\hat{T}_{Li}^2$,

such that the inverse-variance weighted average over all galaxies is given by Eq. (24),

$$\hat{\tau}^{\text{TI stack}} = \frac{\sum_i \delta\hat{T}_{Si} \delta\hat{T}_{Li} / \sigma_{n_s}^2}{\sum_j \delta\hat{T}_{Lj}^2 / \sigma_{n_s}^2}. \quad (\text{C1})$$

Ignoring the small large-scale noise n_L , but taking into account the small-scale noise n_s , this can be rewritten as

$$\hat{\tau}^{\text{TI stack}} = \underbrace{1}_{\text{TI response}} \times \tau + \underbrace{\frac{\sum_i n_{Si} \delta\hat{T}_{Li} / \sigma_{n_s}^2}{\sum_j \delta\hat{T}_{Lj}^2 / \sigma_{n_s}^2}}_{\text{TI noise}}, \quad (\text{C2})$$

i.e. the response to the true τ is unity (again, for small n_L).

On the other hand, the small-scale QE estimator for a given galaxy i takes the form $\tau_i^{\text{QE}} \equiv \delta\hat{T}_{Si} \delta\hat{T}_{Li} / \sigma_{T_L}^2$. When quantifying the variance of τ_i^{QE} , two options are possible. A “local” variance estimate takes into account the local value of $\delta\hat{T}_{Li}$, whereas a “global” variance estimate marginalizes over $\delta\hat{T}_{Li}$, given its variance $\sigma_{T_L}^2$.

The global variance estimate $\sigma_{n_s}^2 / \sigma_{T_L}^2$ is in the spirit of the standard QE, as it preserves the quadratic nature of the estimator, by not including the data itself in the weighting. In this case, the stacked estimator is then

$$\begin{aligned} \hat{\tau}^{\text{QE “global” stack}} &= \frac{\sum_i \delta\hat{T}_{Si} \delta\hat{T}_{Li} / \sigma_{n_s}^2}{\sum_j \sigma_{T_L}^2 / \sigma_{n_s}^2} \\ &= \underbrace{\frac{\sum_i \delta\hat{T}_{Li}^2 / \sigma_{n_s}^2}{\sum_j \sigma_{T_L}^2 / \sigma_{n_s}^2}}_{\text{response, var} \neq 0} \times \tau + \underbrace{\frac{\sum_i n_{Si} \delta\hat{T}_{Li} / \sigma_{n_s}^2}{\sum_j \sigma_{T_L}^2 / \sigma_{n_s}^2}}_{\text{noise}}. \end{aligned} \quad (\text{C3})$$

If one were to instead weigh by the local variance $\sigma_{n_s}^2 / \sigma_{T_L}^2 \delta\hat{T}_{Li}$, one obtains the “local” stacked QE, no longer quadratic in the data:

$$\begin{aligned} \hat{\tau}^{\text{QE “local” stack}} &= \frac{\sum_i \frac{\delta\hat{T}_{Si}}{\delta\hat{T}_{Li}} \frac{\sigma_{T_L}^2}{\sigma_{n_s}^2}}{\sum_j \frac{1}{\delta\hat{T}_{Lj}^2} \frac{\sigma_{T_L}^4}{\sigma_{n_s}^2}} \\ &= \underbrace{\frac{\sum_i \sigma_{T_L}^2 / \sigma_{n_s}^2}{\sum_j \frac{\sigma_{T_L}^2 \sigma_{T_L}^2}{\delta\hat{T}_{Lj}^2 \sigma_{n_s}^2}}}_{\text{response, var} \neq 0} \times \tau + \underbrace{\frac{\sum_i \frac{n_{Si}}{\delta\hat{T}_{Li}} \frac{\sigma_{T_L}^2}{\sigma_{n_s}^2}}{\frac{\sigma_{T_L}^2 \sigma_{T_L}^2}{\delta\hat{T}_{Lj}^2 \sigma_{n_s}^2}}}_{\text{noise}}. \end{aligned} \quad (\text{C4})$$

Thus, both for the “local” and “global” weightings, the stacked QE estimator differs from the stacked TI estimator. Furthermore, the response to the true τ of the stacked QE estimators still has an irreducible variance, even as the small-scale noise $\sigma_{n_s}^2$ goes to zero.

As discussed in the main text, the stacked global QE and stacked TI appear very similar, differing only in their denominator. The fractional difference between the denominators scales as $1/\sqrt{N_{\text{galaxies}}}$, such that it becomes smaller and smaller with larger galaxy samples. However, this difference is not irrelevant, because the variance of the estimators (stacked QE and TI) also decreases as $1/\sqrt{N_{\text{galaxies}}}$. Thus, this difference is actually significant.

APPENDIX D: τ^{sgn} ESTIMATOR: SIGN BALANCING HELPS WITH IMPERFECT FOREGROUND CANCELLATION

“Sign balancing” refers to requiring that the number of positively and negatively weighted galaxies in the stack is the same. Sign balancing is effective as long as the contribution to the foreground emission comes from $> 0.1\%$ of the total number of galaxies. This appendix presents the evidence for this statement. The signed estimator $\hat{\tau}^{\text{sgn}}$ for the foreground emission becomes

$$\hat{\tau}^{\text{sgn}} \equiv \frac{\sum_i \text{sgn}(T_{L,i}) f_i}{\sum_i |T_{L,i}|}, \quad (\text{D1})$$

where $T_{L,i}$ is the large-scale temperature at the position of galaxy i , and f_i is the small-scale temperature due to foreground emission from galaxy i . Simplifying the notation we write,

$$\begin{aligned} \hat{\tau}^{\text{sgn}} &= \frac{\frac{1}{n_{\text{gal}}} \sum_i \text{sgn}(T_{L,i}) f_i}{\frac{1}{n_{\text{gal}}} \sum_i |T_{L,i}|} \\ &\equiv \frac{T^{\text{N}}}{T^{\text{D}}}, \end{aligned} \quad (\text{D2})$$

where n_{gal} is the number of galaxies in the stack, and T^{N} and T^{D} are defined as the numerator and denominator of the r.h.s., respectively. We have the following relations:

$$\begin{cases} \text{sgn}(T_{L,i}) \perp |T_{L,j}| \\ \text{sgn}(T_{L,i}) \perp f_j \\ |T_{L,i}| \perp f_j \\ |T_{L,i}| \not\perp |T_{L,j}|, \end{cases} \quad \forall i, j, \quad (\text{D3})$$

where \perp represents statistical independence of the lhs and rhs. The last relation comes from the fact that galaxies may be located in a region where the CMB temperature is highly correlated (i.e. in $\sim 1 \text{ deg}^2$ patches). Since we assume $\text{sgn}(T_{L,i}) \perp f_j$, there is no mean bias from foregrounds, i.e. $\langle \hat{\tau} \rangle = 0$. However, for a given realization we may find $\langle \hat{\tau} \rangle \neq 0$ due to imperfect cancellation among the n_{gal} sources. Here we derive the typical value of $\hat{\tau}_{\text{bias}}$, i.e. the RMS error of $\hat{\tau}_{\text{bias}}$, to estimate the bias due to imperfect

cancellation. First, we simplify our expression for $\text{var}(\hat{\tau}_{\text{bias}})$:

$$\begin{aligned}\text{var}(\hat{\tau}_{\text{bias}}) &= \text{var}\left(\frac{1}{T^{\text{D}}} T^{\text{N}}\right) \\ &= \text{var}\left(\frac{1}{T^{\text{D}}}\right) \text{var}(T^{\text{N}}) + \left\langle \frac{1}{T^{\text{D}}} \right\rangle^2 \text{var}(T^{\text{N}}) \\ &\quad + \underbrace{\langle T^{\text{N}} \rangle^2}_{=0} \text{var}\left(\frac{1}{T^{\text{D}}}\right) \\ &= \left(\text{var}\left(\frac{1}{T^{\text{D}}}\right) + \left\langle \frac{1}{T^{\text{D}}} \right\rangle^2 \right) \text{var}(T^{\text{N}}) \\ &= \left\langle \frac{1}{(T^{\text{D}})^2} \right\rangle \text{var}(T^{\text{N}}).\end{aligned}\quad (\text{D4})$$

Since $\text{var}(\hat{\tau}_{\text{bias}}) \propto \text{var}(T^{\text{N}})$, we can simply use the variance of the numerator $T^{\text{N}} \equiv \frac{1}{n_{\text{gal}}} \sum \text{sgn}(T_{L,i}) f_i$ for the remainder of the analysis. We first determine the typical value of T^{N} in the absence of fluctuations in foreground contamination, i.e., $f_i = \langle f_i \rangle = f \forall i$. Within a correlated $\sim 1 \text{ deg}^2$ patch of the CMB, $\text{sgn}(T_{L,i})$ are identical. For n_{patch} patches, we expect $\sqrt{n_{\text{patch}}}$ excess patches of one sign (positive or negative). Hence, the total excess number of galaxies with positive or negative weight is

$$n_{\text{gal,excess}} = \sqrt{n_{\text{patch}}} \left(\frac{n_{\text{gal}}}{n_{\text{patch}}} \right) = \frac{n_{\text{gal}}}{\sqrt{n_{\text{patch}}}}. \quad (\text{D5})$$

Therefore, the typical value of T^{N} is

$$\begin{aligned}T_{\text{typical}}^{\text{N}} &= \frac{1}{n_{\text{gal}}} (n_{\text{gal,excess}} f) \\ &= \frac{1}{n_{\text{gal}}} \left(\frac{n_{\text{gal}} f}{\sqrt{n_{\text{patch}}}} \right) \\ &= \frac{f}{\sqrt{n_{\text{patch}}}}.\end{aligned}\quad (\text{D6})$$

By ‘‘sign balancing,’’ i.e., ensuring $n_{\text{gal,excess}} = 0$ before stacking the sources, the bias due to imperfect cancellation of sources is completely negated in the regime where all sources have equal foreground emission. We can relax this assumption and allow f_i to be distinct for all i . We define $\sigma_f^2 \equiv \text{var}(f_i)$ and assume the source weights are already sign balanced such that $\sum_i \text{sgn}(T_{L,i}) = 0$. In this case,

$$\begin{aligned}\text{var}(T^{\text{N}}) &= \frac{1}{n_{\text{gal}}^2} \sum_i \text{sgn}(T_{L,i})^2 \sigma_f^2 \\ &= \frac{\sigma_f^2}{n_{\text{gal}}}.\end{aligned}\quad (\text{D7})$$

Thus,

$$T_{\text{typical}}^{\text{N}} = \frac{\sigma_f}{\sqrt{n_{\text{gal}}}}. \quad (\text{D8})$$

We see sign balancing does not fully negate the imperfect cancellation caused by foreground fluctuations. It is therefore only effective if the bias due to mismatched positive/negative weights is much greater than the bias due to scatter in foreground levels, i.e.,

$$\begin{aligned}\frac{f}{\sqrt{n_{\text{patch}}}} &\gg \frac{\sigma_f}{\sqrt{n_{\text{gal}}}} \\ \Rightarrow \frac{\sigma_f}{f} &\ll \sqrt{\frac{n_{\text{gal}}}{n_{\text{patch}}}} \\ &\ll \sqrt{\rho_{\text{gal}} A_{\text{patch}}},\end{aligned}\quad (\text{D9})$$

where ρ_{gal} is the galaxy number density and A_{patch} is the area of the correlated temperature patch. $A_{\text{patch}} \sim 1 \text{ deg}^2$, and $\rho_{\text{gal}} \sim 1000 \text{ deg}^{-2}$ for current galaxy surveys. Equation (D9) is not satisfied if a small fraction of the galaxy sample dominates the foreground signal. We use the following toy model to determine the minimum fraction of foreground emitting galaxies required for sign balancing to be effective. Consider the case where a fraction x of the galaxy sample emits foregrounds, and $1-x$ emits no foregrounds. The population of emitting galaxies must contribute a total of f/x foreground emission such that $\langle f_i \rangle = (f/x) * x + (0) * (1-x) = f$. The variance of the foreground emission is

$$\begin{aligned}\text{var}(f_i) &= \langle f_i^2 \rangle - \langle f_i \rangle^2 \\ &= \left[\left(\frac{f}{x} \right)^2 x + 0^2 (1-x) \right] - f^2 \\ &= f^2 \left(\frac{1}{x} - 1 \right).\end{aligned}\quad (\text{D10})$$

Thus,

$$\frac{\sigma_f}{f} = \sqrt{\frac{1}{x} - 1}. \quad (\text{D11})$$

Substituting this result into Eq. (D9) and squaring both sides, we determine

$$\begin{aligned}\frac{1}{x} - 1 &\ll \rho_{\text{gal}} A_{\text{patch}} \\ &\ll 1000 \\ \Rightarrow x &\gg 10^{-3}.\end{aligned}\quad (\text{D12})$$

Thus, sign balancing is effective as long as the foreground emission comes from $>0.1\%$ of the total galaxies. To

ensure this condition is satisfied, one can exclude clusters or high-mass groups from the galaxy sample, which would dominate the foreground contamination.

APPENDIX E: MEAN FOREGROUND CONTAMINATION COMPARISON FOR STACKED TI AND SIGNED ESTIMATORS

There are three regimes of foreground contamination for the TI and signed estimators: (1) present small-scale foreground f_S but negligible large-scale foreground f_L ; (2) present large-scale foreground f_L but negligible small-scale foreground f_S ; and (3) non-negligible correlated foregrounds f_S and f_L . In the first two cases, f_S (or f_L) is uncorrelated with the temperature fluctuation δT_L (or δT_S) and thus introduces no bias to either the TI or signed estimators. We study the third regime for both estimators below.

The stacked TI estimator including both large- and small-scale foreground contamination is

$$\hat{\tau}^{\text{TI}} = -\frac{\sum(\delta T_L + f_L)(\delta T_S + f_S)}{\sum(\delta T_L + f_L)^2}. \quad (\text{E1})$$

We assume $f_L \ll \sigma_{T_L}$ (for the typical values of $\sigma_{T_L} \sim 100 \mu\text{K}$ and $f_L \sim 1 \mu\text{K}$ this is quite accurate). For simplicity, we also assume all objects have equal f_L and f_S . The bias to the stacked TI estimator as a function of radial distance r in the 1D stacked photometry profile to first order in f_L/σ_{T_L} is

$$\begin{aligned} \hat{\tau}_{\text{bias}}^{\text{TI}}(r) &= -\left\langle \frac{\sum f_S(r) f_L}{\sum \delta T_L^2} \right\rangle + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right) \\ &= -f_S(r) f_L \left\langle \frac{N_{\text{object}}}{\sum \delta T_L^2} \right\rangle + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right) \\ &\simeq -f_S(r) f_L \frac{1}{\langle \delta T_L^2 \rangle} + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right) \\ &\simeq -\frac{f_S(r) f_L}{\sigma_{T_L}^2} + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right). \end{aligned} \quad (\text{E2})$$

Next, we evaluate the foreground bias for the signed estimator. The stacked signed estimator including large- and small-scale foregrounds is

$$\hat{\tau}^{\text{Sgn}} = -\frac{\sum \text{sgn}(\delta T_L + f_L)(\delta T_S + f_S)}{\sum |\delta T_L + f_L|}. \quad (\text{E3})$$

The bias to the stacked signed estimator to first order in f_L/σ_{T_L} is

$$\begin{aligned} \hat{\tau}_{\text{bias}}^{\text{sgn}}(r) &= -\left\langle \frac{\sum \text{sgn}(\delta T_L + f_L) f_S(r)}{\sum |\delta T_L|} \right\rangle + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right) \\ &\simeq -f_S(r) \frac{\langle \sum \text{sgn}(\delta T_L + f_L) \rangle}{\langle \sum |\delta T_L| \rangle} + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right). \end{aligned} \quad (\text{E4})$$

To determine the value of $\langle \sum \text{sgn}(\delta T_L + f_L) \rangle$, consider the case where $\delta T_L \sim \mathcal{N}(0, \sigma_{T_L})$ and f_L is a small, positive, fixed value. In this case, the distribution of $\delta T_L + f_L$ is just a translated Gaussian, i.e., $\delta T_L + f_L \sim \mathcal{N}(f_L, \sigma_{T_L})$. Thus, the average value of the signed estimator is

$$\begin{aligned} \left\langle \sum \text{sgn}(\delta T_L + f_L) \right\rangle &= A \int_0^\infty \exp\left[-\frac{(x-f_L)^2}{2\sigma_{T_L}^2}\right] dx \\ &\quad - A \int_{-\infty}^0 \exp\left[-\frac{(x-f_L)^2}{2\sigma_{T_L}^2}\right] dx \\ &= A \int_{-f_L}^{f_L} \exp\left[-\frac{x^2}{2\sigma_{T_L}^2}\right] dx \\ &= \sqrt{\frac{2}{\pi}} \frac{f_L}{\sigma_{T_L}} + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right), \end{aligned} \quad (\text{E5})$$

where $A \equiv \frac{1}{\sqrt{2\pi}\sigma_{T_L}}$.

Substituting the result from Eq. (E5) into Eq. (E4) and using the relation $\langle |\mathcal{N}(\mu, \sigma)|^2 \rangle = \frac{2}{\pi} \langle \mathcal{N}(\mu, \sigma)^2 \rangle$, we find

$$\begin{aligned} \hat{\tau}_{\text{bias}}^{\text{Sgn}}(r) &\simeq -\sqrt{\frac{2}{\pi}} \frac{f_L}{\sigma_{T_L}} \frac{f_S(r)}{\langle |\delta T_L| \rangle} + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right) \\ &\simeq -\frac{f_S(r) f_L}{\sigma_{T_L}^2} + \mathcal{O}\left(\left(\frac{f_L}{\sigma_{T_L}}\right)^2\right). \end{aligned} \quad (\text{E6})$$

Thus, to first order in f_L/σ_{T_L} , the TI and signed estimators have the same sensitivity to foreground contamination. For the signed estimator, this bias is only caused by regions where $|\delta T_L| < f_L$. As sources are weighted by δT_L in the TI estimator, these regions contribute very little to the stacked TI profile. For the signed estimator, the mean bias can be completely mitigated by enforcing a minimum δT_L cut such that $|\delta T_L| > |f_L|$ for all sources.

The analysis in this appendix assumes no overlapping of sources, which would contribute a nontrivial bias. A detailed analysis of the effect of overlapping sources on the stacked estimators is left to future work.

APPENDIX F: SIMULATIONS OF CORRELATED EXTRAGALACTIC FOREGROUNDS

1. Simulating a lensing map correlated with τ map

The deflection angle $\vec{\alpha}$ due to the gravitational lensing of photons is given by

$$\vec{\alpha} = \int \frac{d^2\vec{\theta}}{\pi} \kappa(\vec{\theta}) \frac{\hat{\theta}}{\theta}, \quad (\text{F1})$$

where

$$\kappa(\vec{\theta}) = \frac{4\pi G \Sigma(\vec{\theta}) d_L d_{LS}}{c^2 a d_S}. \quad (\text{F2})$$

In the above equations, $\vec{\theta}$ are the 2D angles in the plane of the sky, κ is the convergence field, G is Newton's gravitational constant, c is the speed of light, a is the scale factor, d_L , d_{LS} , and d_S are the comoving distance from observer to the lens, between the lens and the source, and to the source plane respectively. $\Sigma(\vec{\theta})$ is the comoving angular surface density. We therefore need $\Sigma(\vec{\theta})$ to get a lensing map $\kappa(\vec{\theta})$ correlated with the τ map. For this, we first assume a Gaussian projected profile for the angular surface density which we apply to the pointlike galaxy maps. For CMASS galaxies with typical mass $M = 2 \times 10^{13} M_\odot$ at $z = 0.55$, the virial radius is $\approx 1.6'$ [11]. If we assume that the 3σ span of the Gaussian surface density profile should cover most of the profile and that all the mass falls within this then we have $\sigma = 1.6'/3$ i.e. angular surface density profile is a Gaussian with $\sigma \approx 0.53'$. In order to determine the normalization factor of this profile, we use the fact that $\int d^2\vec{r} \Sigma(\vec{r}) = N_{\text{gal}} M$ where N_{gal} is the total number of galaxies, M is their mass, and $d^2\vec{r} = \text{pixel area (in steradians)} \times d_\chi^2$. Here the pixel area corresponds to the area of a single pixel of the map and d_χ is the average comoving distance to the galaxy sample. After getting this normalization, we have $\Sigma(\vec{\theta})$ which gives us $\kappa(\vec{\theta})$ using Eq. (F2). Since the τ map is created using the same galaxy map as described in Appendix G, the κ map and τ map are correlated.

2. Simulating a tSZ map correlated with the τ map

We make some simple assumptions to simulate a tSZ map correlated with the τ map. For example, we see from Fig. 8 of [11] that the amplitude of the tSZ profile of CMASS galaxies is around $-10 \mu\text{K arcmin}^2$ at 150 GHz. We then assume that the tSZ profile is Gaussian in nature centered at each galaxy location. The profile from Fig. 8 of [11] extends out to $6'$ and is a beam convolved profile where the beam full width half maximum (FWHM) is $1.4'$. Therefore, deconvolved FWHM for our simulated tSZ profiles is $\sqrt{6^2 - 1.4^2} \approx 5.8'$. Similar to lensing, in order to determine the correct normalization for the Gaussian profile, we use the fact that $\int d^2\vec{\theta} T = N_{\text{gal}} \times -10 \mu\text{K arcmin}^2$ where T is the tSZ profile in μK and $\int d^2\vec{\theta} T = (\text{pixel area in arcmin}^2) \sum_i T_i$ where $\sum_i T_i$ is the temperature profile summed over all the pixels.

The normalization is therefore determined such that this relation holds and we get our simulated tSZ map.

We then calculate the power spectrum of our tSZ map and compare it with the expected power from galaxies with halo mass $5 \times 10^{12} M_\odot < M < 10^{14} M_\odot$ and redshift $0.4 < z < 0.7$, i.e. mass and redshift range similar to that of CMASS galaxies. We use the halo model predictions from [59] for this purpose. This prediction provides an upper limit on the expected power from the CMASS galaxies. As expected, we find that the tSZ power from our simulated map is below the halo model predictions and it also falls down much quicker on small scales. We call this map as 'tSZ low'. In order to determine the upper limit for the tSZ bias to the τ QE, we thus simulate another tSZ map whose power spectrum roughly matches with the halo model predictions. For this, we simply convolve our galaxy map by a Gaussian beam with $\sigma = 0.88'$ instead of $\sigma = 2.46'$ for our first tSZ map, such that the shape of its power spectrum roughly matches with that of halo model predictions. Then we multiply it by a factor of 1.3 on map level to match the amplitude of the power spectrum with predictions. Please note that these two numbers, namely the beam with $\sigma = 0.88'$ and extra normalization factor of 1.3 have been determined through trial and error to match the halo model prediction. This gives us a second tSZ map which we use to determine the highest expected bias from the tSZ emission coming from CMASS galaxies to cross-correlation of τ reconstructed with a QE with galaxies. We call this map as 'tSZ high'.

APPENDIX G: SIMULATION METHODS FOR TI SNR FORECAST

1. Signal maps

First, we create a τ map correlated with CMASS galaxy positions. We choose map dimensions of $10^\circ \times 10^\circ$ as a balance between sufficient statistics and computational efficiency. To make the map, we create a *plate carrée* (CAR) projection map of the $200^\circ < \text{RA} < 210^\circ$, $10^\circ < \text{Dec} < 20^\circ$ region with a pixel scale of $0.5'$. We choose this region to avoid any large masked regions in the CMASS sample footprint. At each galaxy position, we draw a 2D isotropic Gaussian with FWHM= $5'$ and amplitude of 0.000245. These values are chosen based on the fit to the signed estimator stacked profile using ACT×unWISE data [38].

As shown in Fig. 6, stacking on the unWISE (and potentially LSST) positions causes an excess signal from the two-halo term. To avoid artificially boosting the signal for unWISE and LSST during stacking, we use the measured τ signal from the CMASS map for all forecasts.

Finally, we create a separate signal map for each of the ACT, SPT, SO and CMB-S4 forecasts by convolving the map with a Gaussian kernel with a FWHM of $1.6'$, $1.6'$,

1.4' and 1.0', respectively, to match the beam of the corresponding experiment's noise map (see below).

2. Noise maps

The noise maps for the forecast consist of Gaussian random field realizations with power spectra corresponding to each CMB experiment: ACT, SPT, SO, and CMB-S4. The spectra are plotted in Fig. 4. For ACT, we use the total power spectrum measured on the *Planck* + ACT needlet internal linear combination (ILC) map, which includes all the CMB anisotropies and is convolved with a 1.6' Gaussian beam [60]. For SPT, we use the 5-year projected ILC residual power spectrum curve including the noise and foreground contributions [49,61]. For SO, we use forecast ILC power spectrum with the "goal" sensitivity specifications from [51]. For CMB-S4, we use detector noise-only (i.e. no foregrounds) spectrum assuming 1 μ K-arcmin sensitivity. We add an analytic lensed CMB contribution to the SPT, SO and CMB-S4 power spectra before convolving with a Gaussian beam of 1.6', 1.4' and 1.0', respectively. The maps have the same dimensions, pixel scale and CAR projection as the signal maps.

3. Filtering and stacking

We give a basic overview of the filtering and stacking algorithm here but refer the reader to [11] for further details. To obtain unbiased T_L and T_S estimates, we apply scale-dependent filtering to our maps. We find the compensated aperture photometry filter used in [11] results in a significant noise bias. Instead, we apply low- and high-pass filtering to the maps to decorrelate the large- and small-scale temperature estimates:

$$f_\ell^{\text{low-pass}} = \begin{cases} 1 & \text{if } \ell < 2000 \\ \cos\left(\frac{(\ell-2000)\pi}{300}\right) & \text{otherwise} \\ 0 & \text{if } \ell > 2150, \end{cases} \quad (\text{G1})$$

$$f_\ell^{\text{high-pass}} = \begin{cases} 0 & \text{if } \ell < 2350 \\ \sin\left(\frac{(\ell-2350)\pi}{300}\right) & \text{otherwise} \\ 1 & \text{if } \ell > 2500. \end{cases} \quad (\text{G2})$$

For the stacked analysis we need catalogs for each galaxy sample. For the CMASS and unWISE samples, the catalog consists of the galaxies that lie in the region $200^\circ < \text{RA} < 210^\circ$, $10^\circ < \text{Dec} < 20^\circ$. As the DC2 footprint does not cover this region, we simply take a $10 \times 10 \text{ deg}^2$ patch from the DC2 footprint (specifically $55^\circ < \text{RA} < 65^\circ$, $-40^\circ < \text{Dec} < -30^\circ$) and shift these positions to the above region. Using real (or realistically simulated) galaxy positions ensures the effects of clustering are taken into account in our forecast. The CMASS,

unWISE, and DC2 samples have 7659, 515,095 and 11,904,577 galaxies in this region, respectively.

After applying the low- and high-pass filtering, we use bilinear interpolation to make small $\sim 10' \times 10'$ cutouts of the temperature maps centered on each galaxy position. We estimate T_L for a galaxy position by measuring the mean temperature of the low-pass filtered map in a $\sim 9'$ disk centered on the galaxy. The T_S radial profile is estimated using nine annular ring aperture photometry filters centered on the galaxy with outer radii equally-spaced between 1' and 6'. In each aperture i , we measure the mean temperature of the high-pass filtered map, $T_{S,i}$. The estimated mean τ in each aperture is calculated as

$$\hat{\tau}_i = -\frac{T_L T_{S,i}}{(T_L)^2}. \quad (\text{G3})$$

The optimal stacked estimator is inverse-variance weighted, as shown in Sec. IV C. In the forecast, we use uniform weighting for simplicity. Thus, the stacked τ profile simplifies to

$$\hat{\tau} = -\frac{\sum_{j=1}^{N_{\text{gal}}} T_{L,j} T_{S,j}}{\sum_{j=1}^{N_{\text{gal}}} (T_{L,j})^2}. \quad (\text{G4})$$

4. Covariance estimation and SNR calculation

To estimate the covariance matrix, we simulate 128 independent noise map realizations, filter and stack in the same manner as above and measure a mean τ profile for each realization. We calculate the covariance as the sample covariance over the 128 mean τ profiles. The Hartlap correction [58] to this covariance matrix is $(n_{\text{sims}} - n_{\text{bins}} - 2)/(n_{\text{sims}} - 1) = 1.08$ for 128 simulations and nine profile bins.

Using simulations to estimate the covariance matrix is superior to bootstrap resampling as it accounts for the covariance added by the clustering of the galaxies. We find the difference between the covariance estimated from simulations versus from 10,000 bootstrap samples is negligible for the CMASS sample. However, bootstrap resampling underestimates the covariance by $\sim 20\%$ for the unWISE sample and by more than 50% for the LSST sample, as shown in Fig. 5.

For a given experiment with noise covariance matrix, C , and mean τ profile, \mathbf{v} , we measure the detection SNR of the signal profile over the $10^\circ \times 10^\circ$ map as

$$\left(\frac{S}{N}\right)_{\text{map}}^2 = \mathbf{v}^T C^{-1} \mathbf{v}. \quad (\text{G5})$$

We rescale this result to the full survey sky fraction and galaxy number density as shown in Table II.

APPENDIX H: SNR CALCULATION FOR THE τ QUADRATIC ESTIMATOR

For the QE of τ , the SNR for the cross-power spectrum of the τ map reconstructed using QE and galaxy map g is calculated analytically as

$$\left(\frac{S}{N}\right)^2 = f_{\text{sky}} \sum_{\ell_{\text{bmin}}}^{\ell_{\text{bmax}}} \frac{(2\ell_b + 1) \Delta\ell (C_{\ell_b}^{\tau g})^2}{(C_{\ell_b}^{\hat{\tau}\hat{\tau}})^2 + C_{\ell_b}^{\hat{\tau}\hat{\tau}} \times C_{\ell_b}^{gg}}, \quad (\text{H1})$$

where

$$C_{\ell_b} = \frac{1}{\Delta\ell} \sum_{\ell \in [\ell_1, \ell_2]} C_{\ell}, \quad (\text{H2})$$

and $\Delta\ell$ is the bin width. Here, $C_{\ell_b}^{\tau g}$ is the cross-power spectrum between input τ map and g map, $C_{\ell_b}^{\hat{\tau}\hat{\tau}} = C_{\ell_b}^{\tau\tau} + N_{\ell_b}^{\tau\tau}$ where $C_{\ell_b}^{\tau\tau}$ is the power spectrum of the input τ map and $N_{\ell_b}^{\tau\tau}$ is the reconstruction noise for QE given by

Eq. (12). $C_{\ell_b}^{\hat{\tau}\hat{\tau}}$ is therefore the power spectrum of the reconstructed τ map. $C_{\ell_b}^{gg}$ is the total galaxy power spectrum including the clustering and shot noise components, and f_{sky} is the common sky fraction between the CMB map and galaxy survey. We start with the τ map and CMASS galaxy map as described in Appendix G, and cross-correlate them to calculate the signal $C_{\ell_b}^{\tau g}$. For the noise part, we take into account different clustering and shot noise levels in the CMASS, unWISE, and LSST galaxy samples. For CMASS sample and LSST galaxies, we calculate the galaxy clustering power spectrum using a halo model approach where we make sure that our model matches with the clustering power spectrum of the CMASS maps used in cross-correlations. For unWISE galaxies, we use the measured power spectra for the ‘blue’ sample [62,63]. The shot noise levels for each sample are calculated using the number density values provided in Table II.

For different CMB experiments, the reconstruction noise for the quadratic estimator for τ is different and is calculated using Eq. (12).

-
- [1] N. Aghanim, S. Majumdar, and J. Silk, *Rep. Prog. Phys.* **71**, 066902 (2008).
- [2] Y. B. Zeldovich and R. A. Sunyaev, *Astrophys. Space Sci.* **4**, 301 (1969).
- [3] R. A. Sunyaev and Y. B. Zeldovich, *Comments Astrophys. Space Phys.* **4**, 173 (1972), <https://ui.adsabs.harvard.edu/abs/1972CoASP...4..173S/abstract>.
- [4] R. A. Sunyaev and I. B. Zeldovich, *Mon. Not. R. Astron. Soc.* **190**, 413 (1980).
- [5] L. Salvati, M. Douspis, and N. Aghanim, *Astron. Astrophys.* **614**, A13 (2018).
- [6] B. Bolliet, B. Comis, E. Komatsu, and J. F. Macías-Pérez, *Mon. Not. R. Astron. Soc.* **477**, 4957 (2018).
- [7] B. Soergel, A. Saro, T. Giannantonio, G. Efstathiou, and K. Dolag, *Mon. Not. R. Astron. Soc.* **478**, 5320 (2018).
- [8] A. Amon and G. Efstathiou, *Mon. Not. R. Astron. Soc.* **516**, 5355 (2022).
- [9] G. Aricò, R. E. Angulo, M. Zennaro, S. Contreras, A. Chen, and C. Hernández-Monteagudo, *Astron. Astrophys.* **678**, A109 (2023).
- [10] E. Schaan *et al.*, *Phys. Rev. D* **93**, 082002 (2016).
- [11] E. Schaan *et al.*, *Phys. Rev. D* **103**, 063513 (2021).
- [12] S. Amodeo *et al.*, *Phys. Rev. D* **103**, 063514 (2021).
- [13] E. Schiappucci *et al.*, *Phys. Rev. D* **107**, 042004 (2023).
- [14] J. P. Ostriker and E. T. Vishniac, *Astrophys. J. Lett.* **306**, L51 (1986).
- [15] E. T. Vishniac, *Astrophys. J.* **322**, 597 (1987).
- [16] C. Dvorkin and K. M. Smith, *Phys. Rev. D* **79**, 043003 (2009).
- [17] C. Hernández-Monteagudo and R. A. Sunyaev, *Astron. Astrophys.* **509**, A82 (2010).
- [18] C. Feng and G. Holder, *Phys. Rev. D* **97**, 123523 (2018).
- [19] O. Doré, J. F. Hennawi, and D. N. Spergel, *Astrophys. J.* **606**, 46 (2004).
- [20] J. C. Hill, S. Ferraro, N. Battaglia, J. Liu, and D. N. Spergel, *Phys. Rev. Lett.* **117**, 051301 (2016).
- [21] A. Kusiak *et al.*, *Phys. Rev. D* **104**, 043518 (2021).
- [22] E.-M. Mueller, F. de Bernardis, R. Bean, and M. D. Niemack, *Phys. Rev. D* **92**, 063501 (2015).
- [23] N. S. Sugiyama, T. Okumura, and D. N. Spergel, *J. Cosmol. Astropart. Phys.* **01** (2017) 057.
- [24] E.-M. Mueller, F. de Bernardis, R. Bean, and M. D. Niemack, *Astrophys. J.* **808**, 47 (2015).
- [25] M. Münchmeyer, M. S. Madhavacheril, S. Ferraro, M. C. Johnson, and K. M. Smith, *Phys. Rev. D* **100**, 083508 (2019).
- [26] W. Hu, *Astrophys. J.* **529**, 12 (2000).
- [27] A. Natarajan, N. Battaglia, H. Trac, U.-L. Pen, and A. Loeb, *Astrophys. J.* **776**, 82 (2013).
- [28] A. Roy, A. Lapi, D. Spergel, and C. Baccigalupi, *J. Cosmol. Astropart. Phys.* **05** (2018) 014.
- [29] F. Bianchini and M. Millea, *Phys. Rev. D* **107**, 043521 (2023).
- [30] B. Horowitz, S. Ferraro, and B. D. Sherwin, *Mon. Not. R. Astron. Soc.* **485**, 3919 (2019).
- [31] B. Hadzhiyska, B. D. Sherwin, M. Madhavacheril, and S. Ferraro, *Phys. Rev. D* **100**, 023547 (2019).
- [32] W. Hu and T. Okamoto, *Astrophys. J.* **574**, 566 (2002).
- [33] M. Su *et al.*, arXiv:1106.4313.
- [34] R. A. Sunyaev and I. B. Zeldovich, *Annu. Rev. Astron. Astrophys.* **18**, 537 (1980).

- [35] S. Y. Sazonov and R. A. Sunyaev, *Mon. Not. R. Astron. Soc.* **310**, 765 (1999).
- [36] M. Shehzad Emritte, S. Colafrancesco, and P. Marchegiani, *J. Cosmol. Astropart. Phys.* **07** (2016) 031.
- [37] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020).
- [38] W. R. Coulton *et al.*, arXiv:2401.13033.
- [39] T. Namikawa *et al.*, *Phys. Rev. D* **104**, 063514 (2021).
- [40] T. Namikawa, D. Hanson, and R. Takahashi, *Mon. Not. R. Astron. Soc.* **431**, 609 (2013).
- [41] A. Roy, A. van Engelen, V. Gluscevic, and N. Battaglia, *Astrophys. J.* **951**, 50 (2023).
- [42] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai, *Astrophys. J.* **786**, 13 (2014).
- [43] S. J. Osborne, D. Hanson, and O. Doré, *J. Cosmol. Astropart. Phys.* **03** (2014) 024.
- [44] N. MacCrann *et al.*, arXiv:2304.05196.
- [45] E. Schaan and S. Ferraro, *Phys. Rev. Lett.* **122**, 181301 (2019).
- [46] N. Sailer, E. Schaan, and S. Ferraro, *Phys. Rev. D* **102**, 063517 (2020).
- [47] B. A. Benson *et al.*, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, edited by W. S. Holland and J. Zmuidzinas (SPIE, Bellingham, Washington, 2014).
- [48] A. N. Bender *et al.*, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, Vol. 10708 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, edited by J. Zmuidzinas and J.-R. Gao (SPIE, Bellingham, Washington, 2018), p. 1070803.
- [49] J. A. Sobrin *et al.*, *Astrophys. J. Suppl. Ser.* **258**, 42 (2022).
- [50] S. W. Henderson *et al.*, *J. Low Temp. Phys.* **184**, 772 (2016).
- [51] P. Ade *et al.* (Simons Observatory Collaboration), *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [52] K. Abazajian *et al.*, arXiv:1907.04473.
- [53] K. S. Dawson *et al.*, *Astron. J.* **145**, 10 (2013).
- [54] B. Reid *et al.*, *Mon. Not. R. Astron. Soc.* **455**, 1553 (2016).
- [55] E. L. Wright *et al.*, *Astron. J.* **140**, 1868 (2010).
- [56] A. M. Meisner, D. Lang, and D. J. Schlegel, *Res. Notes Am. Astron. Soc.* **2**, 1 (2018).
- [57] B. Abolfathi *et al.* (LSST Dark Energy Science (LSST DESC) Collaboration), *Astrophys. J. Suppl. Ser.* **253**, 31 (2021).
- [58] J. Hartlap, P. Simon, and P. Schneider, *Astron. Astrophys.* **464**, 399 (2007).
- [59] A. Maniyar, M. Béthermin, and G. Lagache, *Astron. Astrophys.* **645**, A40 (2021).
- [60] W. R. Coulton *et al.*, *Phys. Rev. D* **109**, 063530 (2024).
- [61] S. Raghunathan and Y. Omori, *Astrophys. J.* **954**, 83 (2023).
- [62] A. Krolewski, S. Ferraro, E. F. Schlafly, and M. White, *J. Cosmol. Astropart. Phys.* **05** (2020) 047.
- [63] A. Krolewski, S. Ferraro, and M. White, *J. Cosmol. Astropart. Phys.* **12** (2021) 028.