Integrated Sachs-Wolfe map reconstruction in the presence of systematic errors

Noah Weaverdyck,[*] Jessica Muir, and Dragan Huterer

Department of Physics, University of Michigan, 450 Church St, Ann Arbor, Michigan 48109-1040, USA

The decay of gravitational potentials in the presence of dark energy leads to an additional, late-time contribution to anisotropies in the cosmic microwave background (CMB) at large angular scales. The imprint of this so-called integrated Sachs-Wolfe (ISW) effect to the CMB angular power spectrum has been detected and studied in detail, but reconstructing its spatial contributions to the CMB *map*, which would offer the tantalizing possibility of separating the early- from the late-time contributions to CMB temperature fluctuations, is more challenging. Here, we study the technique for reconstructing the ISW map based on information from galaxy surveys and focus in particular on how its accuracy is impacted by the presence of photometric calibration errors in input galaxy maps, which were previously found to be a dominant contaminant for ISW signal estimation. We find that both including tomographic information from a single survey and using data from multiple, complementary galaxy surveys improve the reconstruction by mitigating the impact of spurious power contributions from calibration errors. A high-fidelity reconstruction further requires one to account for the contribution of calibration errors to the observed galaxy power spectrum in the model used to construct the ISW estimator. We find that if the photometric calibration errors in galaxy surveys can be independently controlled at the level required to obtain unbiased dark energy constraints, then it is possible to reconstruct ISW maps with excellent accuracy using a combination of maps from two galaxy surveys with properties similar to Euclid and SPHEREx.

I. INTRODUCTION

Cosmic microwave background (CMB) photons undergo a frequency shift as they travel to us from the last scattering surface. On top of the redshift due to the expansion of the Universe, an additional contribution to the temperature anisotropy is introduced whenever the Universe is not matter dominated—for example, right after recombination when radiation contributes non-negligibly or at late times when dark energy becomes important. This so-called integrated Sachs-Wolfe (ISW) effect is given by [1,2]

$$\left.\frac{\Delta T}{\bar{T}}\right|_{\text{ISW}}(\hat{\mathbf{n}}) = \frac{2}{c^2}\int_{t_*}^{t_0} dt\, \frac{\partial \Phi(\mathbf{r}, t)}{\partial t}, \qquad (1)$$

where t_0 is the present time, $t_\star$ is that of recombination, c is the speed of light, $\mathbf{r}$ is the position in comoving coordinates, and Φ is the gravitational potential. The late-time ISW signal (referred to hereafter simply as "ISW") has been statistically detected via measurements of the cross-correlation of CMB temperature maps with galaxy maps [3–18] and, more recently, with maps of CMB-lensing convergence [17,18]. These detections serve as an important consistency test of the

standard model of cosmology and can help constrain the properties of dark energy.

The ISW can provide additional information beyond its power spectrum if its *map* can be reconstructed with sufficient accuracy. Since the total large-angle CMB temperature anisotropy is the sum of early- (hereafter "primordial") and late-time contributions,

$$\left.\frac{\Delta T}{\bar{T}}\right|(\hat{\mathbf{n}}) = \left.\frac{\Delta T}{\bar{T}}\right|_{\text{prim}}(\hat{\mathbf{n}}) + \left.\frac{\Delta T}{\bar{T}}\right|_{\text{ISW}}(\hat{\mathbf{n}}), \qquad (2)$$

reconstructing the ISW map would allow us to isolate the primordial-only anisotropy. This separation of the CMB into early- and late-time contributions can also be useful for a variety of cosmological tests. For example, one could study the temporal origin of the large-angle CMB anomalies reported in, e.g., Ref. [19]. One could also subtract the realization-specific contaminating ISW contribution to estimation of primordial non-Gaussianity [20], something that is currently done using theoretical templates for the ISW-lensing bispectrum [21]. Motivated by these considerations, reconstruction of the ISW map has been the focus of a number of recent analyses [17,18,22–30].

In this paper, we study how ISW map reconstruction is affected by a class of observational and astrophysical systematic errors which we will refer to broadly as

---
[*]Corresponding author.
nweaverd@umich.edu

photometric calibration errors or, for conciseness, calibration errors. These systematics afflict all galaxy surveys at large angular scales, contributing to the significant excess of power at large scales found in many recent surveys, including the Sloan Digital Sky Survey (SDSS) [16,31–37], MegaZ [38], WISE-AGN and WISE-GAL [17], and NVSS [14,16,39,40]. Calibration errors are thus already established as one of the most significant systematics impacting large-angle measurements of galaxy surveys, a fact that has broad implications, such as for measuring scale-dependent bias as a signal of primordial non-Gaussianity. As the statistical power of galaxy surveys continues to grow, the control and understanding of systematics like calibration errors is becoming even more important.

There is a variety of ways in which modern photometric surveys assess and mitigate contamination from systematics, many of which rely on cross-correlating galaxy maps with known systematics templates. This can be used to identify contaminated regions, which are then masked or excluded from the analysis (as in Ref. [41]). Alternatively, one can use these templates to subtract or marginalize over systematics-induced spatial variations in the calculation of, for example, the two-point clustering signal [37,42–47]. Such an approach was taken in Ref. [37] to study the overall detection significance of the ISW effect in SDSS data. The authors found results similar to Ref. [16], the authors of which instead accounted for excess power by adding a low-redshift spike in the source distributions. Most of these correlation corrections are perturbative, however, and therefore require fairly clean maps in which systematic effects are minor to begin with. Additionally, while corrections to the two-point statistics are important for the inference of cosmological parameters, they do not remove the systematics from the maps themselves. Suchyta *et al.* [48] propose an alternative approach, wherein measurement biases are characterized by injecting fake objects into Dark Energy Survey images. This neatly avoids the reliance on having small levels of contamination in the input maps, but it still cannot account for certain systematics, such as dust or flux calibration. Whatever the approach taken, some level of residual calibration error will remain.

Some of us previously showed that at levels of calibration control consistent with current and near-future surveys, residual calibration errors are by far the dominant systematic for ISW signal reconstruction [29]. This motivates us to study their impact in more detail. Namely, we would like to study whether the presence of residual calibration errors can be mitigated by combining information from multiple input maps or through better modeling of the contributions of systematics to observed galaxy power. We also wish to investigate to what extent residual calibration errors similarly impact the signal-to-noise ratio of galaxy-CMB cross-correlation and, in turn, the significance of ISW detection. With this aim, we use ensembles of simulated maps to characterize the performance of ISW reconstruction based on surveys like Euclid and SPHEREx, two proposed wide-angle surveys of which the properties are expected to be good for ISW detection and reconstruction. We also consider the benefits of including Planck-like simulations of CMB intensity in the reconstruction effort.

We begin in Sec. II by describing our model for calibration errors, how we reconstruct the ISW map and evaluate its quality, and which input data sets we use. In Sec. III, we compare the performance of ISW reconstruction when using one versus multiple surveys and investigate the impact different assumptions have on the results. In Sec. IV, we relate map reconstruction to the total signal-to-noise ratio of ISW detection, and we conclude in Sec. V.

II. METHODOLOGY

A. Modeling calibration errors

Photometric calibration is a challenge faced by all photometric galaxy surveys. It refers to the adjustments required to establish a consistent spatial and temporal measurement of flux of the target objects in different observation bands. A number of different systematics must be accounted for in calibration, including but not limited to detector sensitivity variation on the focal plane, variation in observing conditions, the presence of foreground stars (as galaxies near them are less likely to be detected), and extinction by dust. Calibration errors are introduced if these systematics are incompletely or inaccurately accounted for.

Our focus is on how calibration errors affect galaxy number counts. To illustrate this, imagine that a perfectly uniform screen (of, e.g., dust) blocks some light from all galaxies. This pushes the faintest galaxies below the survey's flux limit and leads to observation of fewer galaxies in all directions. A pure monopole change such as this increases shot noise but does not affect the angular clustering signal of galaxies. In contrast, in a more realistic scenario where the opacity of this "screen" depends on direction, it affects the observed galaxy clustering signal by adding spurious power and by coupling different multipoles of the measured power spectrum [49,50]. The presence of unaccounted-for calibration errors can thus introduce biases in cosmological parameter estimates from large scale structure (LSS) surveys. These are particularly severe for the ISW effect and other measurements based on signals that, like calibration errors, enter primarily at large angular scales.

In keeping with the picture of calibration errors as a direction-dependent screen, we model them as a modulation of the true galaxy number counts $N(\hat{\mathbf{n}})$, where $\hat{\mathbf{n}}$ is the direction on the sky. The observed, modulated counts are [49]

$$N_{\text{obs}}(\hat{\mathbf{n}}) = [1 + c(\hat{\mathbf{n}})]N(\hat{\mathbf{n}}), \qquad (3)$$

where the field $c(\hat{\mathbf{n}})$ describes the screening effect of calibration errors. While we will generally refer to this kind

of modulation as the result of "calibration errors," as Eq. (3) makes clear, this formalism can describe any residual effect that modulates a survey's selection function.

Though the expression in Eq. (3) will necessarily couple different multipoles, at low ℓ, the impact of calibration errors on the observed galaxy power spectrum is well approximated by

$$C_\ell^{\text{obs}} \approx C_\ell^{\text{gal}} + C_\ell^{\text{cal}}, \tag{4}$$

neglecting multiplicative terms.

Following Muir and Huterer [29], we model the calibration error field $c(\hat{\mathbf{n}})$ as a Gaussian random field with power spectrum C_ℓ^{cal} and quantify the level of residual calibration errors using its variance,

$$\sigma_{\text{cal}}^2 \equiv \text{Var}(c(\hat{\mathbf{n}})) = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} C_\ell^{\text{cal}}. \tag{5}$$

While the conversion between σ_{cal} and the rms magnitude error depends on the faint-end slope of the luminosity function of tracers in the survey, they can be related roughly as $(\delta m)^{\text{rms}} \simeq \sigma_{\text{cal}}$ [49]. Thus, a survey with residual calibration errors of $\sigma_{\text{cal}}^2 = 10^{-6}$ has been calibrated to roughly a millimagnitude.

Motivated by power spectrum estimates for maps of dust extinction corrections and magnitude limit variations in existing surveys (see Figs. 5 and 6 in Ref. [49]), we choose the fiducial calibration error power spectrum to be

$$C_\ell^{\text{cal}} = \alpha^{\text{cal}} \exp\left[-(\ell/w^{\text{cal}})^2\right], \tag{6}$$

with $w^{\text{cal}} = 10$. The normalization constant α^{cal} is varied to achieve the desired σ_{cal}^2. Figures 1 and 2 show the impact of calibration errors of this form on the angular power
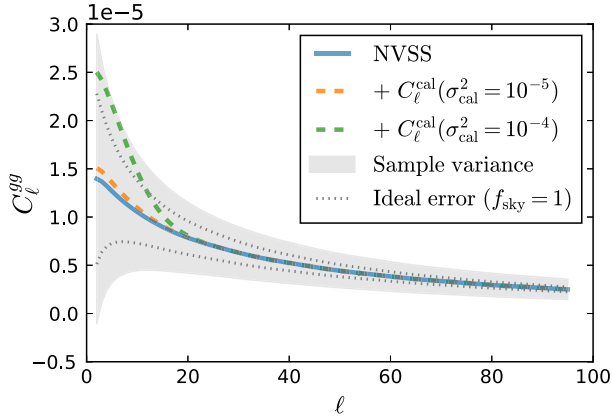

FIG. 1. Effects of calibration errors on galaxy power. The solid curve shows the theoretical angular power spectrum for the NVSS survey [17,18]. The colored dashed curves show the theoretical spectrum with two representative levels of calibration error. The shaded region is the 1σ uncertainty from the survey's sample variance, and the dotted curves indicate the ideal, all-sky cosmic variance.


FIG. 2. Same as Fig. 1, but for the SDSS MphG catalog, following Refs. [17,18]. In this case, the sample variance is driven by sky coverage ($f_{\text{sky}} = 0.22$) as opposed to number density as for NVSS.

spectrum of the NVSS and SDSS MphG galaxy surveys, which have been used to reconstruct the ISW map in previous studies [17,18,27].

For our fiducial model, we assume that calibration error maps for different redshift bins and surveys are uncorrelated with one another. We briefly examine the impact of relaxing such an assumption in Sec. III D.

B. ISW estimator

Similarly to Muir and Huterer [29], we work with the optimal estimator derived by Manzotti and Dodelson [27]. It takes as input n maps, which can include any tracers that carry information about the ISW signal, namely LSS, CMB, or lensing convergence maps. Letting $g_{\ell m}^i$ represent the spherical components of the ith input map, where $i \in \{1, \ldots, n\}$, the optimal estimator for the spherical component of the ISW signal is

$$\hat{a}_{\ell m}^{\text{ISW}} = \sum_{i=1}^{n} R_\ell^i g_{\ell m}^i. \tag{7}$$

The operator

$$R_\ell^i \equiv -N_\ell [D_\ell^{-1}]_{\text{ISW}-i} \tag{8}$$

is a reconstruction filter derived from the covariance matrix,

$$D_\ell = \begin{pmatrix} C_\ell^{\text{ISW,ISW}} & C_\ell^{\text{ISW,1}} & \cdots & C_\ell^{\text{ISW,}n} \\ C_\ell^{1,\text{ISW}} & C_\ell^{1,1} & \cdots & C_\ell^{1,n} \\ \vdots & \vdots & \ddots & \vdots \\ C_\ell^{n,\text{ISW}} & C_\ell^{n,1} & \cdots & C_\ell^{n,n} \end{pmatrix}. \tag{9}$$

In this expression, superscript numbers label the input maps, and $N_\ell \equiv 1/[D_\ell^{-1}]_{\text{ISW}-\text{ISW}}$ estimates the variance of

the reconstruction at multipole ℓ. When a single input map, A, is used, this expression reduces to a simple Wiener filter,

$$\hat{a}_{\ell m}^{\text{ISW}} \xrightarrow{\text{single map}} \frac{C_\ell^{\text{ISW}-A}}{C_\ell^{A-A}} g_{\ell m}^A. \qquad (10)$$

We demonstrate in Appendix A that Eq. (7) is equivalent to the estimator by Manzotti and Dodelson [27], where the CMB temperature map is treated separately from LSS maps, and show that it reduces to the linear-covariance-based filter first proposed by Barreiro et al. [22].

In constructing this ISW estimator, one must make a choice about how to obtain the necessary angular power and cross-power spectra in the covariance matrix. The C_ℓ's can either be extracted from observations (as in Refs. [24,26]) or computed analytically for an assumed cosmology (as in Refs. [17,18,27,29,30]). Analytic calculation is straightforward but introduces a model dependence which can potentially bias results if, for example, calibration error contributions are not modeled correctly [29]. Measuring C_ℓ from observations produces a model-independent estimator and so can help in the case where the theory spectra are inaccurate, but at the expense of limited precision due to sample variance, especially at large scales, scales with low power, or for map combinations that have little correlation.[1] Hybrid methods can also be used, as in Ref. [25], where Barreiro et al. account for observed excesses in the auto-power of NVSS data by using a smoothed fit to data to get the galaxy map's autopower but analytically compute its cross-correlation with the ISW signal.

We therefore consider two limiting cases of constructing the estimator in order to investigate how calibration errors impact the ISW reconstruction:

(1) a "worst" case estimator filter, $R_\ell(C_\ell^{\text{th}})$, where we use the fiducial theory C_ℓ's in the estimator, in which calibration errors' power contributions are not modeled at all, and

(2) a "best" case estimator filter, $R_\ell(C_\ell^{\text{sky}})$, in which calibration error power contributions are modeled perfectly (i.e., the covariance matrix is known). This case may be approximated by, e.g., a smoothed fit of the observed LSS power.

The theoretical spectra are related simply through the expression

$$C_\ell^{\text{sky}} = C_\ell^{\text{th}} + C_\ell^{\text{cal}}, \qquad (11)$$

where C_ℓ^{cal} is the power spectrum of the calibration error field described in Sec. II A. We consider these cases in Secs. III A and III B, respectively.

C. Quality statistic

To quantify the accuracy of a given reconstruction, we use the correlation coefficient between the temperature maps of the true $[T^{\text{ISW}}(\hat{\mathbf{n}})]$ and reconstructed $[T^{\text{rec}}(\hat{\mathbf{n}})]$ ISW signal,

$$\rho = \frac{\frac{1}{N_{\text{pix}}} \sum_k^{N_{\text{pix}}} (T_k^{\text{ISW}} - \bar{T}^{\text{ISW}})(T_k^{\text{rec}} - \bar{T}^{\text{rec}})}{\sigma_{\text{ISW}} \sigma_{\text{rec}}}, \qquad (12)$$

where $\bar{T}^X$ and σ_X^2 are the mean and variance of map $T^X(\hat{\mathbf{n}})$, respectively.[2] We do not include pixel weights in our calculation of ρ, as is done to account for masking effects in Ref. [30]. This is because we work with only full-sky maps, as will be discussed in the next section.

The correlation coefficient can be rewritten in terms of the cross-power between the true ISW map realization and the input tracers,

$$\rho = \frac{\frac{1}{4\pi} \sum_{\ell,i} (2\ell + 1) R_\ell^i \tilde{C}_\ell^{\text{ISW}-i}}{\sigma_{\text{ISW}} \sigma_{\text{rec}}}, \qquad (13)$$

where the tilde denotes pseudo-C_ℓ measured from a given map realization, and we have used Eq. (7) to write

$$\tilde{C}_\ell^{\text{ISW}-\text{rec}} = \sum_i \frac{1}{2\ell + 1} \sum_m [a_{\ell m}^{\text{ISW}}]^* R_\ell^i g_{\ell m}^i \qquad (14)$$

$$= \sum_i R_\ell^i \tilde{C}_\ell^{\text{ISW}-i}. \qquad (15)$$

Because the measured correlation coefficient depends on the specific realization, we assess reconstruction accuracy for a given set of input map properties as follows. We simulate a large number of realizations of correlated maps, then apply the ISW estimator to obtain associated reconstructed ISW maps, and by comparing these with the true ISW maps, we obtain a sample distribution for ρ. Its mean value $\bar{\rho}$, which in the limit of an infinitely large ensemble will approach an expectation value $\langle \rho \rangle$, provides a statistical measure of how accurately the estimator can reproduce the true ISW signal. Studying how $\bar{\rho}$ changes in response to variations in survey properties and modeling choices therefore allows us to understand which factors are most important for obtaining an accurate ISW reconstruction.

We can avoid the computational cost of generating many simulation ensembles by noting that we can obtain a good estimate for the expectation value of ρ if we make the approximation

---

[1]Using the observed spectra also violates an assumption in the maximum likelihood derivation of the estimator, in which the covariance is assumed to be known (i.e., independent of the measured signal).

[2]We also considered s, which measures the rms error between true and reconstructed ISW maps as a complementary quality statistic, but found that for the cases studied here, the information it provided was largely redundant to that given by ρ.

$$\langle\rho\rangle = \left\langle\frac{\frac{1}{4\pi}\sum_{\ell,i}(2\ell+1)R_\ell^i \tilde{C}_\ell^{\mathrm{ISW}-i}}{\sigma_{\mathrm{ISW}}\sigma_{\mathrm{rec}}}\right\rangle \qquad (16)$$

$$\approx \frac{\frac{1}{4\pi}\sum_{\ell,i}(2\ell+1)R_\ell^i C_\ell^{\mathrm{ISW}-i}}{\hat{\sigma}_{\mathrm{ISW}}\hat{\sigma}_{\mathrm{rec}}}, \qquad (17)$$

that is, we replace the pseudo-C_ℓ's with their expectation value across realizations, $\tilde{C}_\ell \to C_\ell$. We will refer to the quantity in Eq. (17) as $\hat{\rho}$, defining

$$\hat{\sigma}_{\mathrm{ISW}} = \sqrt{\frac{1}{4\pi}\sum_\ell (2\ell+1)C_\ell^{\mathrm{ISW}}} \qquad (18)$$

$$\hat{\sigma}_{\mathrm{rec}} = \sqrt{\frac{1}{4\pi}\sum_{\ell,i,j}(2\ell+1)R_\ell^i R_\ell^j C_\ell^{ij}}, \qquad (19)$$

to approximate the rms fluctuations in the true and reconstructed ISW maps. Here, the indices i and j label the input tracer maps and the sum over ℓ runs over the multipoles $\ell \in [2, 95]$, a range chosen to conservatively to include all scales where the ISW signal is important.

We have tested the approximation $\hat{\rho} \approx \langle\rho\rangle$ in Eq. (17) extensively and found it works well when the estimator filter R_ℓ is built from analytically computed spectra but can break down if R_ℓ is composed of $\tilde{C}_\ell$'s extracted from map realizations. This behavior is related to the way in which using measured C_ℓ's makes ρ depend on $\tilde{C}_\ell$, such that $\bar{\rho} = \rho(\langle\tilde{C}_\ell\rangle)$ is no longer a good approximation of $\langle\rho(\tilde{C}_\ell)\rangle$. Appendix B discusses this in more detail.

D. Simulated surveys

By working with simulated maps, we are able to study in detail how calibration error levels and modeling choices affect ISW signal reconstruction.

Since we are concerned only with large scales, we model the ISW signal, total CMB temperature anisotropy, and galaxy number density fluctuations as correlated Gaussian fields. We use HEALPY [51] to generate map realizations based on input auto- and cross-power spectra which we compute analytically following the standard expressions given, e.g., in Ref. [29]. We use the Limber approximation for $\ell \geq 20$, having verified that this affects ρ at the level of 0.1% or less for the surveys and range of σ_{cal}^2 considered here. We compute C_ℓ for multipoles with $\ell \leq 95$, as this range contains almost all of the ISW signal [52]. Accordingly, our simulations are sets of HEALPIX maps of resolution $\mathrm{NSIDE} = (\ell_{\max}+1)/3 = 32$. We refer the reader to Ref. [29] for a more detailed description of the reconstruction pipeline.

Because our goal is to study the impact of calibration errors and not survey geometry, we assume full-sky coverage in all of our analyses. Bonavera *et al.* [30] found that in overlapping regions of partial sky LSS surveys, ISW

reconstruction quality degrades only slightly compared to the full-sky case. Therefore, the performance of a given estimator using full-sky maps should be indicative of its performance using maps with only partial sky coverage.

Our fiducial cosmological model is ΛCDM, with the best-fit cosmological parameter values from Planck 2015, $\{\Omega_c h^2, \Omega_b h^2, \Omega_\nu h^2, h, n_s\} = \{0.1188, 0.0223, 0, 0.6774, 0.9667\}$. Unless otherwise stated, ISW reconstructions are performed on 2000 map realizations for each analysis and include multipole information down to $\ell_{\min} = 2$.

Within this framework, four pieces of information are required to model a LSS survey: the distribution of its sources along the line of sight $n(z)$, a prescription for how they are binned in redshift, their linear bias $b(z)$, and their projected number density per steradian $\bar{n}$. Below, we describe how our choices for these characteristics are based on the properties of promising future probes of the ISW effect.

1. Euclid-like LSS survey

Our fiducial survey is modeled on Euclid, a future LSS survey with large sky coverage and a deep redshift distribution [53], which is expected to be an excellent probe of the ISW effect [52,54]. We assume the redshift distribution used by Martinet *et al.* [55],

$$\frac{dn}{dz} = \frac{3}{2z_0^3}z^2 \exp\left[-(z/z_0)^{1.5}\right], \qquad (20)$$

which has a maximum at $z_{\mathrm{peak}} \simeq 1.21 z_0$. We choose $z_0 = 0.7$ and $\bar{n} = 3.5 \times 10^8$, with a photo-$z$ redshift uncertainty of $\sigma(z) = 0.05(1+z)$ which smoothes the edges of redshift bins. For simplicity, we assume a constant galaxy bias of $b(z) = 1$. Our results are qualitatively insensitive to this choice as long as the bias is reasonably well approximated for the input maps. This is because the bias term cancels in the estimate of the ISW signal, so that fractional differences between true and modeled bias have little impact on ρ. We refer the reader to Ref. [29] for further details on both fitting for bias and the impact that mismodeling can have on reconstruction.

In Sec. III A, we investigate the improvement in ISW map reconstruction when the fiducial Euclid-like survey is split into six redshift bins with edges at $z \in \{0.01, 0.4, 0.8, 1.2, 1.6, 2, 3.5\}$ (see inset of Fig. 3), as compared to the unbinned case. We subsequently use the six-binned Euclid survey as our fiducial case.

2. SPHEREx-like LSS survey

We model a second survey on the SPHEREx All-Sky Spectral Survey (SPHEREx), a proposed survey that has been optimized to study LSS in the low-redshift universe. One of its goals is to place stringent limits on primordial non-Gaussianity [56], which will require rigorous control
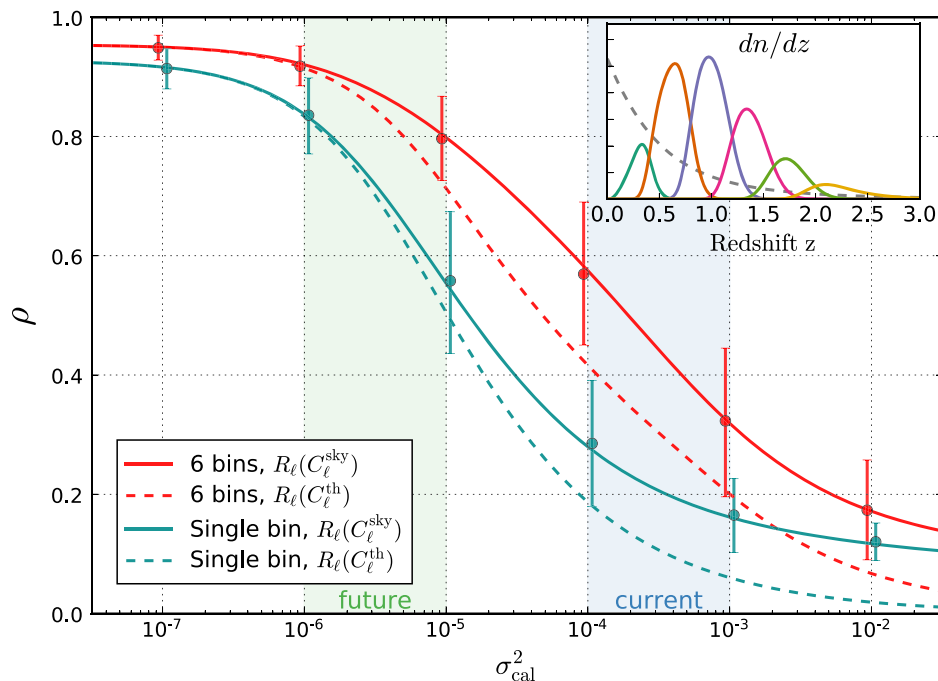
FIG. 3. Quality of map reconstruction ρ vs the calibration error variance σ_{cal}^2 for our fiducial, Euclid-like survey. The colors of the lines indicate how tomographic information is handled, showing that splitting the survey into six redshift bins (red) improves the reconstruction compared to the single-bin case (blue). Solid curves indicate cases when the calibration error is included in the ISW estimator $[R_\ell(C_\ell^{sky})]$, while the dashed curves show the reconstructions in which the effects of the calibration errors are not included $[R_\ell(C_\ell^{th})]$ (see Sec. II B for details). Points (offset horizontally for clarity) show the mean ($\bar{\rho}$) of 2,000 realizations, with error bars indicating the 68% spread across realizations. The corresponding smooth curves are $\hat{\rho}$, the analytical estimate of $\bar{\rho}$ from Eq. (17). The inset illustrates the redshift distribution across bins overlaid with the ISW kernel in gray (reproduced from Ref. [29]). The vertical, shaded regions show the approximate current and projected levels of control over residual calibration errors. Calibration errors between redshift bins are modeled as uncorrelated.

of calibration errors. Given this, SPHEREx will provide excellent input map candidates for ISW map reconstruction. Its shallower reach makes it complementary to the deeper mapping of the LSS provided by Euclid.

SPHEREx will identify galaxies with varying levels of redshift uncertainty, ranging from $\sigma_z < 0.003(1 + z)$ up to $\sigma_z > 0.1(1 + z)$. Grouping these into catalogs with different levels of precision provides collections of galaxies useful for different science goals. The $\sigma_z < 0.1(1 + z)$ catalog with a projected $\sim$300 million galaxies was identified in Ref. [56] as the best subsample for f_{NL}^{loc} detection. Our investigations confirm it to be the best for ISW detection as well. We therefore fit its projected redshift distribution given in Ref. [56] to the functional form for dn/dz given in Eq. (20). We select $z_0 = 0.46$, which results in a peak dn/dz of $z_{peak} \simeq 0.56$. We have confirmed that our results are not strongly sensitive to changes in this redshift distribution, in agreement with the findings of Ref. [29].

We use a projected number density of $\bar{n} = 6.6 \times 10^7$ and consider the case where the survey is split into six redshift bins. We choose their edges by scaling the Euclid-like survey's binned redshift distribution to the SPHEREx median redshift, resulting in redshift bin edges at $z \in \{0.01, 0.26, 0.53, 0.79, 1.05, 1.31, 2.30\}$. This still

provides sufficient sampling of the field in each bin to ensure that shot noise is subdominant to the galaxy signal power.

3. Planck-like CMB survey

CMB data have frequently been used in conjunction with LSS data for ISW map reconstruction. Recent examples include Ref. [17], which used NVSS radio data, the Planck-lensing convergence map, and Planck temperature data. That analysis was subsequently extended to include more LSS tracers in Ref. [18]. However, in both of these cases, residual systematics limit the usefulness of lensing data to scales of $\ell \geq 10$ and $\ell \geq 8$, respectively. Bonavera *et al.* [30] investigated the usefulness of CMB data for ISW reconstruction using a simulation pipeline similar to ours, finding that both CMB temperature and polarization data only modestly improve reconstruction quality but carry a greater benefit when the LSS tracers themselves contain less information (due to, e.g., noise or other properties of the survey).

It is then natural to ask whether CMB data can help mitigate the impact of calibration errors in LSS maps. We therefore consider CMB temperature as an additional input map. To compute the total CMB temperature power

spectrum, C_ℓ^{TT}, we compute the primordial-only contributions using a modified version of CAMB [57] and add them to our calculations for C_ℓ^{ISW}. As the CMB power spectrum is determined within the limits of cosmic variance at low ℓ and the ISW signal is already dominated by the primary (that is, non-ISW) CMB anisotropies, we do not include calibration errors in the generation of CMB temperature maps. Though CMB polarization and lensing could provide additional information, residual systematics remain at large scales for each (see Refs. [58,59], respectively), so for simplicity, we do not include them in this analysis.

III. RESULTS

To characterize the impact of calibration errors in LSS surveys on the ISW map reconstruction, and the potential to mitigate these impacts, we look at multiple combinations of input maps with different properties. Specifically, we consider the impact of binning in redshift, of adding CMB intensity data, and of including additional LSS information from another survey. For each of these studies, we examine two limiting cases for the estimator. The best case scenario, which we will reference as $R_\ell(C_\ell^{\text{sky}})$, is when one perfectly models all contributions to the galaxy power, including residual calibration errors. The worst case, referenced by $R_\ell(C_\ell^{\text{th}})$, is when the estimator is built out of theoretical spectra with no power from calibration errors. The power spectra in these two cases are related by Eq. (11).

We use the analytical $\hat{\rho}$ to estimate the mean reconstruction quality across a wide range of σ_{cal}^2, while performing reconstruction on simulated maps for selected values, to both verify the accuracy of $\hat{\rho}$ and to generate error bars for the spread of ρ across simulations.

A. One survey: Binning in redshift

We first consider the Euclid-like survey alone and investigate the impact of binning in redshift on the quality of reconstructions in the presence of calibration errors. We model calibration errors in the binned case by adding the calibration field's power to the autopower spectrum of each bin i per Eq. (4): $C_\ell^{i,i} \to C_\ell^{i,i} + C_\ell^{\text{cal}}$. We do not add any power to the cross-spectra, though we test the impact of contamination in the cross-spectra in Sec. III D.[3]

The results of this study are shown in Fig. 3. For reference, we use a vertical shaded band to mark the level of calibration corresponding to current surveys, defined roughly as the range bounded by the residual SDSS DR8 limiting magnitude variations [60] and the SDSS

---

[3]In reality, the power contribution from calibration errors will also vary somewhat across bins, depending on the redshift dependence of the faint-end slope of the luminosity function for the tracer population. We have assumed here for simplicity that the power contribution is independent of redshift.

übercalibration [61]. ("Future" levels are defined roughly as those between that required to limit bias on cosmological parameters to below their projected uncertainties and an intermediate level before bias becomes unacceptable; see Refs. [49,29] for details.)

As shown in Fig. 3, splitting the survey into six redshift bins results in significant improvement in reconstruction at all levels of calibration error. This improvement is comparable to reducing σ_{cal}^2 of the single-bin case by a factor of 10 at "current" levels.

Tomographic information is useful because it allows galaxy counts to be weighted more optimally, taking advantage of the fact that the ISW effect becomes stronger at lower redshift as dark energy becomes more dominant and structure growth slows. For instance, considering the expected ISW reconstruction power from each bin when using optimal weights [i.e., the squared contribution of each term in Eq. (7), using $R_\ell(C_\ell^{\text{sky}})$], we find that with no calibration error, the first redshift bin contributes 87% as much power as the second bin, with subsequent bins contributing 58%, 31%, 15%, and 10% as much power, respectively. There is an additional benefit to binning when calibration errors are considered. Since the low-redshift bins have a higher clustering signal than the high-redshift bins, they are less impacted by the same level of calibration error. Thus, the optimal weighting changes depending on the level of calibration error. When calibration errors are increased to $\sigma_{\text{cal}}^2 = 10^{-4}$, for example, the first bin contributes the most power to the reconstruction, with bins 2–6 only contributing 39%, 12%, 4%, 2%, and 1% as much power. As we will show later, this error-level-dependent weighting will mean adding information from a shallower survey such as SPHEREx makes reconstruction more robust against calibration errors.

The importance of accounting for calibration errors in the estimator is apparent in the difference between the dashed and solid curves, where doing so improves $\bar{\rho}$ for $\sigma_{\text{cal}}^2 \gtrsim 10^{-6}$, with $\Delta\bar{\rho} \approx 0.1$–0.2 at current levels of calibration. This improvement is roughly comparable to the improvement seen from binning in redshift.

Though for clarity we do not include this case in the figure, we additionally studied the effect of using the observed, unsmoothed galaxy-galaxy power in the estimator (that is, $\tilde{C}_\ell$, power spectra extracted from map realizations rather than computed analytically). We find that in this case, $\bar{\rho}$ converges to the same value as the $R_\ell(C_\ell^{\text{sky}})$ case when calibration errors are very large but is greatly reduced from $\bar{\rho}$ found using either $R_\ell(C_\ell^{\text{sky}})$ or $R_\ell(C_\ell^{\text{th}})$ when calibration errors are small ($\sigma_{\text{cal}}^2 \ll 10^{-5}$). For example, for a single input map in the limit of no calibration errors, quality reduces from $\bar{\rho} = 0.93$ to 0.83 when we switch to using observed $\tilde{C}_\ell$'s. If we also use the observed (unsmoothed) *cross-correlation* between the LSS map and the CMB for the galaxy-ISW term in the estimator,

TABLE I. Mean reconstruction quality coefficients $\bar{\rho}$ of ISW map reconstructions for various combinations of input maps and select levels of calibration error. The second column indicates $\bar{\rho}$ for the case of zero calibration error. The following columns show the reconstruction quality for two nonzero values of the calibration error variance; here, $R_\ell(C_\ell^{th})$ [$R_\ell(C_\ell^{sky})$] indicates the case where calibration errors are unaccounted [accounted] for in the estimator. Note, when $\sigma_{cal}^2 = 0$, $C_\ell^{th} = C_\ell^{sky}$.

	R_ℓ	$R_\ell(C_\ell^{th})$		$R_\ell(C_\ell^{sky})$	
σ_{cal}^2	0	10^{-6}	10^{-4}	10^{-6}	10^{-4}
TT	0.46	$\cdots$	$\cdots$	$\cdots$	$\cdots$
Euclid (1 bin)	0.92	0.83	0.19	0.84	0.29
Euclid (6 bin)	0.95	0.91	0.41	0.92	0.57
SPHEREx (6 bin)	0.89	0.88	0.52	0.88	0.62
Euc + Spx + TT	0.96	0.92	0.47	0.93	0.73

reconstruction quality is further degraded to $\bar{\rho} = 0.74$ in the absence of calibration errors. This is because primary CMB anisotropies are large compared to ISW contributions, causing the measured galaxy-CMB correlation to receive relatively large noise contributions from chance correlations between LSS maps and the primordial CMB.

Given the significant improvement in reconstruction that binning provides, from here forward, we adopt the configuration with six tomographic bins as our fiducial Euclid-like survey.

B. Effect of adding Planck *TT* data

We now consider adding information from the Planck-like CMB temperature map described in Sec. II D. When used as the only input map, the reconstruction is considerably worse than that found using the ideal Euclid-like survey (Table I). We include it in our study, however, because any realistic study attempting to reconstruct the ISW signal will likely include CMB temperature data. Additionally, the reconstruction quality attainable with CMB temperature data alone provides a useful baseline against which to compare the performance of estimators based on LSS maps.

With CMB temperature data alone, we find an average reconstruction quality of $\bar{\rho} = 0.46$, in good agreement with Ref. [30]. To put this into proper context, however, it is important to note that there is a large scatter around that mean; while the average reconstruction quality is indicative of performance, any single realization, such as that of our own Universe, can vary substantially in fidelity. The purple band in Fig. 4 shows the extent of this scatter for ISW reconstruction based on just the CMB map.

When CMB temperature information is combined with that from LSS maps, it significantly improves reconstruction quality, but only if the true galaxy power spectrum C_ℓ^{sky} (including calibration error contributions) is used in the estimator, as can been seen by the behavior of


FIG. 4. Reconstruction quality when using binned Euclid-like survey and CMB intensity data separately and in combination. The purple curve and shaded band show the mean and 68% spread of ρ from simulations. As in Fig. 3, red curves are results when using the binned Euclid-like survey alone, whereas blue curves are the result of using both the fiducial survey and CMB intensity map. Like in Fig. 3, solid curves are for the case where calibration error power is correctly modeled in the estimator, and dashed curves are for when they are not modeled at all. Neglecting the presence of calibration errors in a LSS map can actually degrade the quality of the ISW reconstruction compared to using the CMB temperature alone.

the solid curves on the right-hand side of Fig. 4. The blue $\rho(\sigma_{cal}^2)$ curve describing the CMB + LSS reconstruction tracks the maximum of the curves corresponding to reconstructions using the CMB and LSS input maps separately, shown by the purple and red curves, respectively. This occurs because the estimator down-weights the LSS survey the more it is affected by calibration errors, converging to the *TT*-only reconstruction quality in the limit of large calibration errors. If one does not model calibration error power contributions, however, then any improvement from combining multiple input maps is marginal at best and can in fact result in a *worse* reconstruction than just using the CMB data alone. This demonstrates the importance of ensuring that the LSS C_ℓ's used in the ISW estimator are a good fit to the observed spectra.

C. Effect of an additional LSS survey: SPHEREx-like

We now consider the addition of our fiducial six-bin SPHEREx-like survey described in Sec. II D 2, assuming for simplicity that it has the same level of calibration errors as the Euclid-like survey. Results are shown in Fig. 5.

In the limit of no calibration errors, the SPHEREx-like survey offers little additional information. In fact, adding both SPHEREx and CMB *TT* results in negligible improvement over the Euclid-like only case ($\Delta\bar{\rho} < 0.003$ compared to a spread of $\sigma_{Euc+Spx+TT} = 0.019$).

However, by comparing the black and blue curves, we see that including the SPHEREx-like survey does make the reconstruction somewhat more robust against calibration errors. The reason for this is similar to why binning in redshift is helpful: recall that, in the case of binning, having narrow, low-redshift bins means having some bins with higher galaxy autopower than the unbinned case, which then have less susceptibility to a given level of calibration error. Similarly, SPHEREx has a shallower redshift distribution, and thus an intrinsically higher clustering signal, so that it can actually provide a better reconstruction than the Euclid-like survey at moderate levels of calibration error. We would expect to see similarly increased robustness to calibration errors for any tracer with a larger clustering signal, including tracers with a larger bias.

Finally, just as for Euclid, we find that if calibration errors are not accounted for in the estimator, then adding LSS data can actually result in a worse reconstruction than that from using CMB temperature data alone.

D. Effect of varying calibration error properties

We now test how sensitive the results in the previous sections are to our assumptions about calibration errors, showing the results in Fig. 6.

First, the left panel shows what happens when we vary the level of cross-correlation between the calibration errors of different LSS maps. It is conceivable that residual
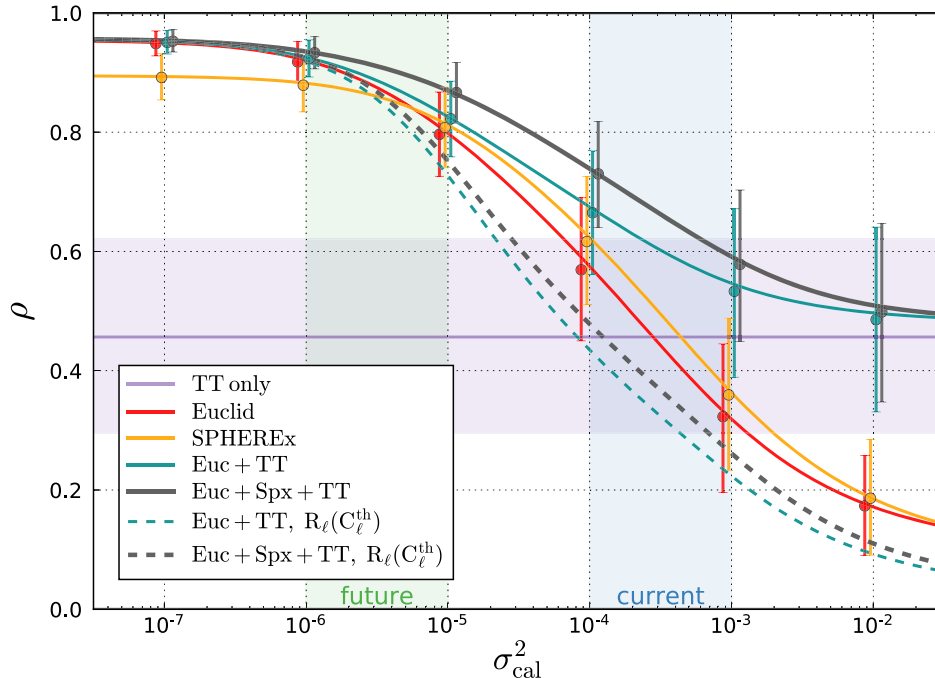

FIG. 5.   Comparison of ISW reconstruction quality using the LSS surveys and CMB temperature individually and in combination, for various levels of calibration error in the Euclid-like and SPHEREx-like surveys. Colors are the same as those of Fig. 4. Both of the LSS surveys are split into six redshift bins (see Sec. III A), with calibration errors uncorrelated between bins and surveys. The dashed curve shows the combined reconstruction if calibration errors are not included in the estimator. Using LSS surveys to improve the ISW map reconstruction from the CMB temperature-only case requires calibration errors to be controlled to $\sigma_{cal}^2 \lesssim 10^{-4}$.
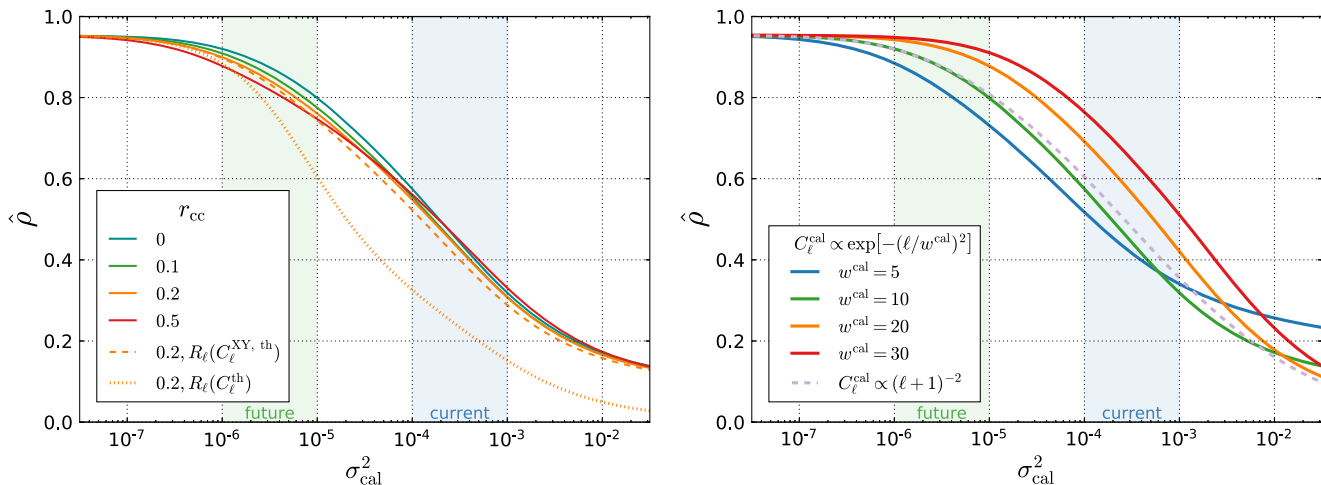
FIG. 6.   *Left:* Effect of cross-correlation between calibration errors in different bins of the fiducial Euclid-like survey, given by $C_\ell^{\mathrm{cal},ij} = r_{\mathrm{cc}}\sqrt{C_\ell^{\mathrm{cal},ii}C_\ell^{\mathrm{cal},jj}}$, for bins $i \neq j$. Solid curves have calibration errors accounted for in the estimator $[R_\ell(C_\ell^{\mathrm{sky}})]$. The dashed curve indicates the case where only the autopower contributions of the calibration errors are accounted for in the estimator $[R_\ell(C_\ell^{\mathrm{XY,th}})]$, and the dotted curve indicates the case where calibration errors are not accounted for at all $[R_\ell(C_\ell^{\mathrm{th}})]$. Cross-correlation of the errors results in mild degradation of the reconstruction for $\sigma_{\mathrm{cal}}^2 \lesssim 10^{-4}$ but otherwise has little effect as long as the autocorrelation is correctly modeled in the estimator. *Right:* Dependence of ρ on the shape of C_ℓ^{cal}. Solid curves indicate C_ℓ^{cal} of the same form as Eq. (6) but with width w^{cal} varied. The dashed curve indicates the case where the error spectrum takes the form $C_\ell^{\mathrm{cal}} \propto (\ell+1)^{-2}$. Reconstruction fares worse when calibration error power contributions are more concentrated at the largest angles, where the ISW kernel is largest. In all cases, the estimator uses the true observed LSS power spectrum (C_ℓ^{sky}).

calibration errors can be correlated across different bins of a single survey, or even across different surveys, especially if the error has an astronomical origin. To model such correlation, we set the level of cross-correlation between the calibration errors of maps i and j using a parameter r_{cc}, where

$$C_\ell^{\mathrm{cal},ij} = r_{\mathrm{cc}}\sqrt{C_\ell^{\mathrm{cal},ii}C_\ell^{\mathrm{cal},jj}}, \quad \text{for } i \neq j \qquad (21)$$

As we only consider cases where calibration errors in all maps are characterized by the same C_ℓ^{cal}, this reduces to

$$C_\ell^{\mathrm{cal},ij} = r_{\mathrm{cc}}C_\ell^{\mathrm{cal}}, \quad \text{for } i \neq j. \qquad (22)$$

We consider the six-bin fiducial Euclid-like survey and find that this kind of correlated error results in mild degradation of the reconstruction for $\sigma_{\mathrm{cal}}^2 \lesssim 10^{-4}$, but otherwise it has little effect as long as calibration errors are correctly modeled in the estimator [that is, $R_\ell(C_\ell^{\mathrm{sky}})$ is used].

If calibration errors are *not* accounted for [$R_\ell(C_\ell^{\mathrm{th}})$ is used], reconstruction suffers considerably, as shown by the dotted curve. We also use a dashed curve [labeled $R_\ell(C_\ell^{\mathrm{XY,th}})$] for the case where the estimator filter correctly accounts for the autopower contributions of calibration errors but neglects the cross-power contributions. As seen by comparing the solid, dashed, and dotted orange curves in Fig. 6, reconstruction quality is far more sensitive to

accurate modeling of the calibration error contribution to the autopower than to the cross-power. Thus, fitting the observed autopower for each map but using theoretical cross-powers, as is done in Ref. [25], should harm the reconstruction relatively little, depending on the fitting scheme; we find $\Delta\hat\rho \approx -0.03$ at $\sigma_{\mathrm{cal}}^2 = 10^{-4}$ for $r_{\mathrm{cc}} = 0.2$, far less than the typical variation over realizations shown in Fig. 3.

Additionally, we study the impact of changing the shape of the calibration error power spectrum C_ℓ^{cal}, showing the results in the right panel of Fig. 6. We first vary the width parameter w^{cal} of the calibration error power spectrum C_ℓ^{cal} given in Eq. (6). Results for different values of w^{cal} are qualitatively similar, though for fixed σ_{cal}^2, the reconstruction is less sensitive to calibration errors when w^{cal} is larger. The reason for this is that ρ is most sensitive to contamination at the lowest multipoles, as will be discussed in Sec. IV. Using a power law $C_\ell^{\mathrm{cal}} \propto (\ell+1)^{-2}$ gives results similar to our fiducial Gaussian form with $w^{\mathrm{cal}} = 10$.

IV. IMPACT OF CALIBRATION ERRORS ON S/N ESTIMATES

Given the extent to which calibration errors degrade the quality of ISW signal reconstruction, it makes sense to ask whether they also impact the signal-to-noise ratio (S/N) of ISW detection. Detection of the ISW signal via the cross-correlation between the CMB and LSS maps has been the

subject of considerable study, as it serves as an important consistency test for the presence of dark energy. The significance of detection varies considerably depending on the LSS tracer sample and the statistical methods used [62], as well as how systematics in the LSS data are treated [16,37,40]. Reference [37] used systematics templates to correct the observed power spectra for SDSS galaxies, finding a S/N loss of ~0.5 if such corrections are neglected. Reference [16] introduced a low-redshift spike in the source distributions in order to reproduce the observed excess autopower in NVSS and SDSS catalogs and estimate that such systematics result in an uncertainty of ΔS/N $\pm$ 0.4. The most recent results come from the Planck Collaboration, which found ~4σ evidence for the ISW effect, with most of the signal coming from cross-correlation of the CMB temperature with the NVSS radio catalog and CMB lensing [18].

The maximum achievable signal-to-noise ratio can be obtained by considering an ideal survey that perfectly traces the ISW (i.e., $C_\ell^{gg} = C_\ell^{Tg} = C_\ell^{ISW}$), resulting in a maximum S/N ~ 6–10 for ΛCDM cosmology [13,14,35,52,54,62,63].

Our goal is to study how calibration errors impact the significance of ISW detection. There are multiple ways one can quantify detection of the ISW effect, including correlation detection between LSS and the CMB, template matching to an assumed model, or model comparison. Each of these methods relies on different assumptions and tests different statistical questions (see Ref. [62] for a detailed review). Here, we adopt the simple correlation detection statistic which quantifies the expected deviation from a null hypothesis of no correlation between LSS (g) and CMB temperature (T). In this formalism, the S/N for ISW detection is

$$\left(\frac{S}{N}\right)^2 \simeq \sum_\ell (\mathbf{C}_\ell^{Tg})^* (\mathbf{C}_\ell^{cov})^{-1} \mathbf{C}_\ell^{Tg}, \qquad (23)$$

where we have assumed the multipoles contribute independently to the S/N. Here, $\mathbf{C}_\ell^{Tg}$ is a vector of the ISW-LSS cross-spectra, and the covariance matrix elements corresponding to LSS maps i and j can be written as

$$C_\ell^{cov,ij} = \langle \Delta C_\ell^{Ti} \Delta C_\ell^{Tj} \rangle \qquad (24)$$

$$\simeq \frac{C_\ell^{Ti} C_\ell^{Tj} + C_\ell^{TT}(C_\ell^{ij} + C_\ell^{cal,ij} + \delta_{ij}\frac{1}{\bar{n}_{ij}})}{f_{sky}(2\ell+1)}, \qquad (25)$$

where the last term in the numerator is due to shot noise and δ_{ij} is the Kronecker delta.[4] Equations (23) and (25)

---

[4]Strictly speaking, this will result in a slight underestimate of the significance, as technically the null hypothesis covariance, with $C_\ell^{Ti} = C_\ell^{Tj} = 0$ in Eq. (25), should be used. However, as the galaxy-ISW cross-power terms are small compared to the galaxy autopower, we follow the practice in most of the literature of keeping them in the S/N calculation.

demonstrate that all cosmological tests using LSS-CMB cross-correlation are limited in their constraining power due to sample variance and the relatively large amplitude of the primordial CMB fluctuations. They also make it clear that calibration errors will reduce the significance of ISW detection.

We assume calibration errors to be uncorrelated between maps, so $C_\ell^{cal,ij} \to \delta_{ij} C_\ell^{cal,ij}$. For a single LSS map, Eq. (23) reduces to the form

$$\left(\frac{S}{N}\right)^2 \simeq f_{sky} \sum_\ell \frac{(2\ell+1)(C_\ell^{Tg})^2}{(C_\ell^{Tg})^2 + C_\ell^{TT}(C_\ell^{gg} + C_\ell^{cal} + 1/\bar{n}_g)}. \qquad (26)$$

If there are no calibration errors, we find S/N = 6.6 for our Euclid-like survey, which is near the maximum[5] for this cosmology, S/N = 6.7. As σ_{cal}^2 increases from 0 to current levels, the total S/N reduces to 4.9–5.7, a drop of only ~15%–30%. This can be seen in the S/N values listed for various σ_{cal}^2 in the legend of Fig. 7. In contrast, for the same level of error, average reconstruction quality $\hat{\rho}$ is reduced by 40%–60%. Clearly, ISW signal reconstruction is substantially more affected by calibration errors than is ISW detection significance.

The greater robustness of the total S/N to calibration errors is due to the fact that it has support at higher multipoles. This is most easily illustrated in the single-map case, where the contribution per multipole to the total signal-to-noise ratio is

$$\left(\frac{S}{N}\right)^2_\ell \equiv \frac{d(S/N)^2}{d\ell} = (2\ell+1)\frac{(C_\ell^{Tg})^2}{C_\ell^{TT} C_\ell^{gg} + (C_\ell^{Tg})^2}. \qquad (27)$$

Figure 7 shows how the contribution per multipole responds to different levels of calibration error.

As σ_{cal}^2 increases, the signal-to-noise ratio decreases at lower multipoles, but contributions at higher multipoles remain unchanged. These higher-multipole contributions are thus still available to contribute to the overall S/N.

Map reconstruction is more sensitive to the largest scales. For the single-map case, this can be illustrated analytically as follows. Using the single-map estimator from Eq. (10), we can write the estimated reconstruction quality statistic as

---

[5]This limit can in principle be increased, e.g., through the method of Frommert et al. [64] in which the observed LSS map is used to reduce the local variance and which in our case brings the maximum possible S/N to 7.2, or through the inclusion of polarization data as in the work by Liu et al. [65].

FIG. 7. Contribution to squared signal-to-noise ratio per multipole for our fiducial Euclid-like survey with varying levels of calibration error. Total combined S/N for each level of calibration error $\sigma^2_{\rm cal}$ is given in the legend.

$$\hat{\rho} = \frac{\frac{1}{4\pi}\sum_\ell (2\ell+1)(\frac{C_\ell^{Tg}}{C_\ell^{gg}})C_\ell^{Tg}}{\sigma_{\rm ISW}\sqrt{\frac{1}{4\pi}\sum_\ell(2\ell+1)(\frac{C_\ell^{Tg}}{C_\ell^{gg}})^2 C_\ell^{gg}}}$$

$$= \frac{1}{\sigma_{\rm ISW}}\sqrt{\frac{1}{4\pi}\sum_\ell(2\ell+1)\frac{(C_\ell^{Tg})^2}{C_\ell^{gg}}}$$

$$= \frac{1}{\sigma_{\rm ISW}}\sqrt{\frac{1}{4\pi}\sum_\ell\left(\frac{S}{N}\right)^2_\ell C_\ell^{TT}\left(1+\frac{(C_\ell^{Tg})^2}{C_\ell^{gg}C_\ell^{TT}}\right)}$$

$$\approx \frac{1}{\sigma_{\rm ISW}}\sqrt{\frac{1}{4\pi}\sum_\ell\left(\frac{S}{N}\right)^2_\ell C_\ell^{TT}}. \qquad (28)$$

Here, $(S/N)^2_\ell$ is the quantity given by Eq. (27), which, when summed over ℓ, gives $(S/N)^2$. Thus, we see from Eq. (28) that $\hat{\rho}$ is proportional to a total (S/N) of which the terms are weighted by C_ℓ^{TT}. Since C_ℓ^{TT} drops sharply as $\sim\ell^{-2}$, the quality of map reconstruction $\hat{\rho}$ is more impacted by large-angle calibration errors than the overall S/N is. This is also a primary cause for the degradation in reconstruction quality seen when $\sigma^2_{\rm cal}$ was concentrated at lower multipoles in Sec. III D.

V. CONCLUSIONS

Reconstruction of the integrated Sachs-Wolfe signal would allow, for the first time, a clean separation of the CMB temperature anisotropies into contributions from 300,000 years after the big bang and those from some ~10 billion years later. This, in turn, would allow for a more informed assessment of the origin of the "large-angle CMB anomalies" and a more complete elimination of ISW contaminants to CMB-based measurements of primordial non-Gaussianity. Accurate ISW reconstruction requires wide-angle large-scale structure maps from which the

gravitational potential evolution can be inferred, but in practice, these maps are plagued by photometric calibration errors due to a host of atmospheric, instrumental, and selection-induced systematics. In our previous work, we found that the realistic levels of residual calibration error severely degrade the accuracy of the reconstructed ISW map.

In this paper, we investigated how the effects of residual photometric calibration errors on the ISW map reconstruction can be mitigated by using tomographic information and by combining data from multiple surveys. To quantify the amount of residual calibration errors, we use their variance $\sigma^2_{\rm cal}$, the square root of which is roughly equal to the rms *magnitude* fluctuations induced by these systematics.

We find that for a Euclid-like survey with a single redshift bin, to achieve a reconstruction comparable in quality to that derived from the CMB temperature map alone (with an average correlation between the true and reconstructed ISW maps of only $\bar{\rho}\approx 0.46$), one must limit the variance of the calibration error field to $\sigma^2_{\rm cal}\lesssim 10^{-5}$.

This can be improved significantly if we exploit the tomographic information available by binning the LSS data in redshift (Fig. 3). We also show that if the model spectra in the estimator differ substantially from those used to generate the input maps, e.g., by using theoretical power spectra that do not account for the observed excess auto-power in the LSS survey from calibration errors, then reconstruction quality is substantially degraded. It is therefore crucial to verify that the theoretical spectra in the estimator are a good fit to those observed or to otherwise use smoothed fits.

We next consider how using multiple input maps, probing different tracers and redshift ranges, improves ISW signal reconstruction. We find that, as long as the excess power contributed by calibration errors is adequately modeled in the estimator, the resulting reconstruction is always better than that from either of the input maps individually. If the excess power from calibration errors is not modeled, however, adding a map can actually degrade reconstruction (Fig. 4). The CMB temperature map adds information to the reconstruction at all levels of calibration error but is especially valuable if the LSS maps are subject to calibration errors with $\sigma^2_{\rm cal}\gtrsim 10^{-5}$. Using a six-bin SPHEREx-like survey provides qualitatively similar results to the Euclid-like one, but because it is shallower, the reconstruction is less accurate in the limit of no calibration errors ($\rho_{\rm Spx}=0.89\pm 0.04$ vs $\rho_{\rm Euc}=0.95\pm 0.02$, where errors indicate the 68% spread across realizations). However, SPHEREx's shallower depth and therefore intrinsically higher clustering signal becomes an asset in the presence of calibration errors, making the survey more robust against calibration errors and leading to a better reconstruction for $\sigma^2_{\rm cal}\gtrsim 6\times 10^{-6}$. (Similarly, we would expect the increased clustering of tracers with larger bias to

help mitigate the effects of calibration error as well.) Therefore, a combination of a shallower and a deeper LSS survey provides complementarity useful for separation of calibration errors from the ISW signal and necessary for a good ISW map reconstruction in the presence of such errors.

Using all three simulated surveys as input—Euclid, SPHEREx, and CMB temperature—significantly improves reconstruction compared to using a single survey with current levels of residual calibration errors, or CMB temperature data alone. We find that if the calibration errors can be controlled to the level of $\sigma_{cal}^2 \lesssim 10^{-6}$ ($\sigma_{cal}^2 \lesssim 10^{-5}$), then the combination of Euclid, SPHEREx, and CMB temperature maps can produce the ISW map reconstruction to an excellent accuracy of $\rho = 0.93 \pm 0.03$ ($\rho = 0.87 \pm 0.05$). This is roughly the same level of calibration control required for future LSS surveys to avoid biasing measurements of cosmological parameters like the non-Gaussianity parameter f_{NL} and the dark energy equation of state. Thus, high-quality ISW reconstructions will, in a sense, "come for free" with the developments needed for cosmology constraints with next-generation surveys.

We additionally test the robustness of our results against changes to the properties of the calibration error field, looking at cross-correlations between calibration errors in different maps as well as the shape of their spectrum. We found that cross-correlation between the calibration errors of different tracer maps degrades the reconstruction most for $10^{-6} \lesssim \sigma_{cal}^2 \lesssim 10^{-4}$ but that this effect is relatively minor, provided the autocorrelation contribution is accounted for in the estimator (Fig. 6, left).

We also find that, compared to map reconstruction accuracy, the overall significance of ISW detection is less strongly affected by calibration errors. This is because its signal is distributed more broadly in multipole space than that of the map reconstruction quality statistic. To clarify this, we analytically relate $\hat{\rho}$ to the commonly used ISW detection S/N statistic in the case of a single-input LSS map and show that $\hat{\rho}$ is weighted by an additional factor of C_ℓ^{TT}, demonstrating map reconstruction's greater sensitivity to the largest scales [Eq. (28)].

As an extension to this work, one could study how the inclusion of CMB lensing and polarization data can improve ISW map reconstruction, provided the systematics present in those data sets could be sufficiently accurately modeled. The results of Bonavera et al. [30] indicate that the use of lensing as input can contribute significantly to reconstruction quality, but they also show that current noise levels limit its effectiveness. Notably, the residual lensing systematics at $\ell \leq 8$ present a challenge, as this is where the ISW signal is strongest, and we expect these systematics to affect reconstruction with CMB lensing and polarization in a manner broadly similar to unaccounted for calibration errors in LSS maps at those scales.

Further work could also be performed to concretely explore how to best approximate the best case reconstruction scenario, wherein calibration errors are fully accounted for, by using real input data. Here, we have only characterized the limiting cases where the residual calibration error contribution to the LSS power is fully known or fully unknown and have not addressed intermediate cases where they are partially accounted for.

Finally, we have only worked in the full-sky case, whereas real data will necessarily have only partial sky coverage. Others have already shown that incomplete sky coverage only very minorly degrades reconstruction quality for areas covered by the input data sets [30], and we do not expect this to change in the presence of calibration errors. Nevertheless, a full analysis that attempts to make predictions for real surveys should take the actual sky coverage and survey-specific systematics into account.

Even with these considerations, it is clear that accounting for the presence of residual calibration errors in LSS surveys is a critical step for any reconstruction of the ISW map, as their presence and treatment impact both the survey characteristics and set of input maps that produce the optimal reconstruction.

ACKNOWLEDGMENTS

APPENDIX A: EQUIVALENCE WITH OTHER ESTIMATORS

We now demonstrate the equivalence between our estimator for the ISW map coefficients $\hat{a}_{\ell m}^{ISW}$ and the estimators proposed by Manzotti and Dodelson [27] and Barreiro et al. [22].

Our estimator in Eq. (7) is based on a version of the likelihood from Manzotti and Dodelson [27] that has been reformulated to handle observed CMB maps like any other input map. Manzotti and Dodelson [27] derive their estimator using the likelihood

$$
\mathcal{L}(T^{ISW}) \propto \frac{1}{\sqrt{\det(CD)}} \times \exp\left\{ -\frac{1}{2} d^T D^{-1} d \right\}
$$
$$
\times \exp\left\{ -\frac{1}{2}(T^{obs} - T^{ISW}) C^{-1} (T^{obs} - T^{ISW}) \right\},
$$

(A1)

where $C \equiv C^p + C^n$ is the angular power spectrum of the primordial C^p and noise C^n contributions to CMB temperature fluctuations, d is a vector of ISW and LSS tracer maps, and D is the covariance matrix between the ISW and

LSS tracers [see Ref. [27], Eqs. (4)–(6)], with ISW maps associated with the first (1) index. This likelihood is a product of the independent likelihoods for $(T^{\text{obs}} - T^{\text{ISW}})$ and for the input maps in d.

Instead of explicitly including independent terms for the primordial CMB and LSS tracers (which are assumed to have no cross-correlation), we include the total observed CMB temperature,

$$T^{\text{obs}} = T^p + T^{\text{ISW}}, \tag{A2}$$

where T^p includes both the primordial CMB temperature as well as any instrumental noise terms. We then expand the data vector to include T^{obs},

$$\mathbf{d}_{\ell m} = (a_{\ell m}^{\text{ISW}}, g_{\ell m}^1, ..., g_{\ell m}^n) \to (a_{\ell m}^{\text{ISW}}, g_{\ell m}^1, ..., g_{\ell m}^n, a_{\ell m}^{\text{obs}}),$$

with $a_{\ell m}$ indicating spherical components of ISW and CMB temperature fluctuations and $g_{\ell m}$ indicating components of LSS overdensity. The covariance matrix is similarly expanded to account for the cross-correlation of T^{obs} with the ISW and LSS tracers

$$D_\ell \to \begin{pmatrix} C_\ell^{\text{ISW,ISW}} & C_\ell^{\text{ISW},1} & \cdots & C_\ell^{\text{ISW},n} & C_\ell^{\text{ISW,obs}} \\ C_\ell^{1,\text{ISW}} & C_\ell^{1,1} & \cdots & C_\ell^{1,n} & C_\ell^{1,\text{obs}} \\ \vdots & \vdots & \ddots & \vdots & \\ C_\ell^{n,\text{ISW}} & C_\ell^{n,1} & \cdots & C_\ell^{n,n} & C_\ell^{n,\text{obs}} \\ C_\ell^{\text{obs,ISW}} & C_\ell^{1,\text{obs}} & \cdots & C_\ell^{n,\text{obs}} & C_\ell^{\text{obs,obs}} \end{pmatrix}. \tag{A3}$$

Assuming that at the scales we consider the observed CMB is cross-correlated with other LSS tracers only through the ISW, we have

$$\begin{aligned} C_\ell^{\text{obs,ISW}} &= C_\ell^{\text{ISW,ISW}} \\ C_\ell^{\text{obs,LSS}_i} &= C_\ell^{\text{ISW,LSS}_i}, \\ C_\ell^{\text{obs,obs}} &= C_\ell^{p,p} + C_\ell^{\text{ISW,ISW}}, \end{aligned} \tag{A4}$$

assuming there is no residual cross-correlation between the primordial and late-time CMB. Maximizing the resulting likelihood

$$\mathcal{L}(T^{\text{ISW}}) \propto \frac{1}{\sqrt{\det(D)}} \times \exp\left\{-\frac{1}{2} d^T D^{-1} d\right\} \tag{A5}$$

gives the optimal estimator given in Sec. II B.

To show that this is equivalent to the estimator derived from Eq. (A1), we focus on the case of using CMB temperature and a single LSS tracer as input maps. For compactness, and to make the connections with other ISW estimators in the literature more apparent, we adopt

notation from Barreiro $et\ al.$ [22], where s, g, and T indicate the ISW, LSS tracer, and observed CMB temperature signals, respectively. We then have

$$\mathbf{d}_{\ell m} = (a_{\ell m}^s, g_{\ell m}^g, a_{\ell m}^T), \tag{A6}$$

and

$$D_\ell = \begin{pmatrix} C_\ell^{ss} & C_\ell^{sg} & C_\ell^{ss} \\ C_\ell^{sg} & C_\ell^{gg} & C_\ell^{sg} \\ C_\ell^{ss} & C_\ell^{sg} & C_\ell^{TT} \end{pmatrix}. \tag{A7}$$

From Eqs. (7) and (8), our estimator gives

$$\begin{aligned} \hat{a}_{\ell m}^s &= \frac{-1}{[D_\ell^{-1}]_{11}} ([D_\ell^{-1}]_{12} g_{\ell m}^g + [D_\ell^{-1}]_{13} a_{\ell m}^T) \\ &= \left(\frac{C_\ell^{sg}(C_\ell^{TT} - C_\ell^{ss})}{C_\ell^{gg} C_\ell^{TT} - (C_\ell^{sg})^2}\right) g_{\ell m}^g \\ &\quad + \left(\frac{C_\ell^{ss} C_\ell^{gg} - (C_\ell^{sg})^2}{C_\ell^{gg} C_\ell^{TT} - (C_\ell^{sg})^2}\right) a_{\ell m}^T. \end{aligned} \tag{A8}$$

We now calculate the estimator of Manzotti and Dodelson [27]. Denoting their covariance matrix without the CMB as D_ℓ', we use their Eq. (9):

$$\begin{aligned} \hat{a}_{\ell m}^{s,\text{MD}} &= ([C_\ell^{pp}]^{-1} + [D'^{-1}]_{11})^{-1} \\ &\quad \times (-[D_{\ell'}^{-1}]_{12} g_{\ell m}^g + [C_\ell^{pp}]^{-1} a_{\ell m}^T) \\ &= \left(\frac{1}{C_\ell^{pp}} + \frac{C_\ell^{gg}}{\det|D_\ell'|}\right)^{-1} \left[\frac{C_\ell^{sg}}{\det|D_\ell'|} g_{\ell m}^g + \frac{1}{C_\ell^{pp}} a_{\ell m}^T\right] \\ &= \left(\frac{1}{\det|D_\ell'| + C_\ell^{gg} C_\ell^{pp}}\right) [C_\ell^{pp} C_\ell^{sg} g_{\ell m}^g + (\det|D_\ell'|) a_{\ell m}^T]. \end{aligned} \tag{A9}$$

Expanding the determinant, we find

$$\begin{aligned} \hat{a}_{\ell m}^{s,\text{MD}} &= \left(\frac{C_\ell^{sg} C_\ell^{pp}}{C_\ell^{gg}(C_\ell^{ss} + C_\ell^{pp}) - (C_\ell^{sg})^2}\right) g_{\ell m}^g \\ &\quad + \left(\frac{C_\ell^{ss} C_\ell^{gg} - (C_\ell^{sg})^2}{C_\ell^{gg}(C_\ell^{ss} + C_\ell^{pp}) - (C_\ell^{sg})^2}\right) a_{\ell m}^T, \end{aligned} \tag{A10}$$

which, using the relation $C_\ell^{TT} = C_\ell^{pp} + C_\ell^{ss}$, is equivalent to the estimator given by Eq. (A9).

This is also equivalent to the estimator proposed by Barreiro $et\ al.$ [22], which uses the Cholesky decomposition (L) of the covariance matrix D_ℓ'' (denoted C_ℓ in Ref. [22]),

$$D_\ell'' = \begin{bmatrix} C_\ell^{gg} & C_\ell^{sg} \\ C_\ell^{sg} & C_\ell^{ss} \end{bmatrix} = L_\ell L_\ell^T. \tag{A11}$$

Note that here the ISW index is last instead of first, in contrast to the covariance matrices D_ℓ and D'_ℓ used previously. The Cholesky decomposition is written as

$$L_\ell = \begin{bmatrix} \sqrt{C_\ell^{gg}} & 0 \\ \dfrac{C_\ell^{sg}}{\sqrt{C_\ell^{gg}}} & \sqrt{C_\ell^{ss} - \dfrac{(C_\ell^{sg})^2}{C_\ell^{gg}}} \end{bmatrix}. \qquad \text{(A12)}$$

Equations. (8) and (9) from Ref. [22] give the ISW estimate (their $\hat{s}_{\ell m}$) is

$$\hat{a}_{\ell m}^{s,\text{B08}} = \frac{L_{12}}{L_{11}}\left(1 - \frac{L_{22}^2}{L_{22}^2 + C_\ell^{pp}}\right)g_{\ell m}^g + \frac{L_{22}^2}{L_{22}^2 + C_\ell^{pp}}a_{\ell m}^T,$$

where we have suppressed the ℓ-dependence of L and combined their observed ISW signal ($s_{\ell m}$) and noise ($n_{\ell m}$) terms into the single term $a_{\ell m}^T$. We use C_ℓ^{pp} to denote the combined power of noise and the primordial CMB, in keeping with the notation above. Plugging this into Eq. (A12), we obtain

$$\begin{aligned}
\hat{a}_{\ell m}^{s,\text{B08}} &= \frac{C_\ell^{sg}}{C_\ell^{gg}}\left(\frac{C_\ell^{pp}}{(C_\ell^{ss} + C_\ell^{pp}) - (C_\ell^{sg})^2/C_\ell^{gg}}\right)g_{\ell m}^g \\
&\quad + \frac{C_\ell^{ss} - (C_\ell^{sg})^2/C_\ell^{gg}}{(C_\ell^{ss} + C_\ell^{pp}) - (C_\ell^{sg})^2/C_\ell^{gg}}a_{\ell m}^T \\
&= \left(\frac{C_\ell^{sg}(C_\ell^{TT} - C_\ell^{ss})}{C_\ell^{gg}C_\ell^{TT} - (C_\ell^{sg})^2}\right)g_{\ell m}^g \\
&\quad + \left(\frac{C_\ell^{ss}C_\ell^{gg} - (C_\ell^{sg})^2}{C_\ell^{gg}C_\ell^{TT} - (C_\ell^{sg})^2}\right)a_{\ell m}^T, \qquad \text{(A13)}
\end{aligned}$$

which is the same as Eq. (A9).

APPENDIX B: ESTIMATING ρ WITH $R_\ell(\tilde{C}_\ell)$

Here, we show why using raw pseudo-C_ℓ's ($\tilde{C}_\ell$) in the estimator results in a degraded reconstruction, for which $\bar{\rho}$ is not well approximated by $\hat{\rho}$ [Eq. (17)].

For a given realization, ρ is constructed from the covariance between the true and reconstructed ISW maps [Cov($T^{\text{ISW}}, T^{\text{rec}}$), i.e., the numerator in Eq. (13)] normalized by the square root of the individual variances of the true and reconstructed ISW maps (σ_{True}^2 and σ_{rec}^2, respectively). We therefore focus on how using realization-specific $\tilde{C}_\ell$'s in the estimator filter R_ℓ affects the individual C_ℓ contributions to σ_{rec}^2 and Cov($T^{\text{ISW}}, T^{\text{rec}}$) ($\sigma_{\text{ISW}}^2$ is unaffected by our choice of R_ℓ). For simplicity, we work with a single input map.

If the ISW estimator filter R_ℓ is constructed from analytically computed model C_ℓ's, the power spectrum of the ISW map for a given realization will be

$$\begin{aligned}
\tilde{C}_\ell^{\text{rec-rec,th}} &= R_\ell^2(C_\ell^{\text{th}})\tilde{C}_\ell^{\text{gal-gal}}, \\
&= \left(\frac{C_\ell^{\text{ISW-gal}}}{C_\ell^{\text{gal-gal}}}\right)^2 \tilde{C}_\ell^{\text{gal-gal}}. \qquad \text{(B1)}
\end{aligned}$$

We add the superscript "th" to distinguish this reconstructed ISW power spectrum from the one where the filter R_ℓ is built from $\tilde{C}_\ell$'s, which will be discussed shortly. The expectation value for this over many realizations is

$$\begin{aligned}
\langle\tilde{C}_\ell^{\text{rec-rec,th}}\rangle &= \left(\frac{C_\ell^{\text{ISW-gal}}}{C_\ell^{\text{gal-gal}}}\right)^2 \langle\tilde{C}_\ell^{\text{gal-gal}}\rangle, \\
&= \frac{(C_\ell^{\text{ISW-gal}})^2}{C_\ell^{\text{gal-gal}}}. \qquad \text{(B2)}
\end{aligned}$$

Now, let us look at the behavior of the reconstructed ISW power when the galaxy autopower spectra in the estimator filter are extracted from the observed maps. Denoting this version of the filter by

$$\tilde{R}_\ell \equiv \frac{C_\ell^{\text{ISW-gal}}}{\tilde{C}_\ell^{\text{gal-gal}}}, \qquad \text{(B3)}$$

we write

$$\begin{aligned}
\tilde{C}_\ell^{\text{rec-rec}} &= \tilde{R}_\ell^2 \tilde{C}_\ell^{\text{gal-gal}}, \\
&= (C_\ell^{\text{ISW-gal}})^2\left(\frac{1}{\tilde{C}_\ell^{\text{gal-gal}}}\right). \qquad \text{(B4)}
\end{aligned}$$

Because the measured $\tilde{C}_\ell^{\text{gal-gal}}$ appears in the denominator of this expression, taking its expectation value over many realizations is somewhat more complicated. To do so, we use the fact that $(2\ell+1)\tilde{C}_\ell^{\text{gal-gal}}$ is χ^2-distributed with $2\ell+1$ degrees of freedom. This means $\tilde{C}_\ell^{\text{rec-rec}}/(2\ell+1)$ follows an inverse-χ^2 distribution, with an expectation value[6]

$$\left\langle\frac{\tilde{C}_\ell^{\text{rec-rec}}}{2\ell+1}\right\rangle = \frac{1}{2\ell-1}\frac{(C_\ell^{\text{ISW-gal}})^2}{C_\ell^{\text{gal-gal}}}. \qquad \text{(B5)}$$

Therefore, the average reconstructed power is

$$\begin{aligned}
\langle\tilde{C}_\ell^{\text{rec-rec}}\rangle &= \frac{2\ell+1}{2\ell-1}\langle\tilde{C}_\ell^{\text{rec-rec,th}}\rangle \\
&= \langle\tilde{C}_\ell^{\text{rec-rec,th}}\rangle\left(1 + \frac{2}{2\ell-1}\right). \qquad \text{(B6)}
\end{aligned}$$

---

[6]We refer the reader to Refs. [66,67] for discussions of the bias introduced when inverting an estimator, with implications specifically for estimating the inverse covariance matrix.

Because $\langle \tilde{C}_\ell^{\mathrm{rec-rec}} \rangle$ is strictly positive, this increased power results in an increase in the total variance of the reconstruction map $\tilde{\sigma}_{\mathrm{rec}}^2$ compared to that from the theory-only filter reconstruction $\sigma_{\mathrm{rec,th}}^2$,

$$\langle \tilde{\sigma}_{\mathrm{rec}}^2 \rangle = \frac{1}{4\pi} \sum_\ell (2\ell+1) \langle \tilde{C}_\ell^{\mathrm{rec-rec}} \rangle$$

$$= \sigma_{\mathrm{rec,th}}^2 + \frac{1}{4\pi} \sum_\ell (2\ell+1) \frac{\langle \tilde{C}_\ell^{\mathrm{rec-rec,th}} \rangle}{\ell - 1/2}. \quad \text{(B7)}$$

In contrast, we find the average cross-power $\langle \tilde{C}_\ell^{\mathrm{ISW-rec}} \rangle$ between reconstructed and true ISW maps remains unchanged. The increased power of the reconstruction thus results in a net decrease in $\langle \rho \rangle$, per Eq. (13), and hence is not well approximated by simply substituting the theory C_ℓ, as is done to compute $\hat{\rho}$. Additionally, this suggests that a simple scaling of R_ℓ in order to "debias" the reconstruction will not improve ρ.

To understand why the cross-power does not increase, we again use the observed galaxy autopower in the estimator and approximate its expectation value. The cross-power is given by

$$\tilde{C}_\ell^{\mathrm{ISW-rec}} = \tilde{R}_\ell \tilde{C}_\ell^{\mathrm{ISW-gal}}$$

$$= \left(\frac{C_\ell^{\mathrm{ISW-gal}}}{\tilde{C}_\ell^{\mathrm{gal-gal}}} \right) \tilde{C}_\ell^{\mathrm{ISW-gal}}. \quad \text{(B8)}$$

Here, we have a quotient of two *nonindependent* χ^2 random variables. Generically, we can approximate the average of a function of two random variables X and Y through a second-order Taylor expansion about the mean of each (μ_X, μ_Y),

$$\langle f(X,Y) \rangle \approx f(\mu_X, \mu_Y) + \frac{1}{2} f''_{XX}(\mu_X, \mu_Y) \langle (X-\mu_X)^2 \rangle$$
$$+ f''_{XY}(\mu_X, \mu_Y) \langle (X-\mu_X)(Y-\mu_Y) \rangle$$
$$+ \frac{1}{2} f''_{YY}(\mu_X, \mu_Y) \langle (Y-\mu_Y)^2 \rangle,$$

where a prime indicates a derivative with respect to the respective subscripted variable. By taking $f(X,Y)$ to be $\tilde{C}_\ell^{\mathrm{ISW-rec}}$, of the form X/Y, then from Eq. (B8), we can approximate the mean cross-power to be

$$\langle \tilde{C}_\ell^{\mathrm{ISW-rec}} \rangle \approx C_\ell^{\mathrm{ISW-gal}} \frac{\langle \tilde{C}_\ell^{\mathrm{ISW-gal}} \rangle}{\langle \tilde{C}_\ell^{\mathrm{gal-gal}} \rangle}$$

$$\times \left(1 - \frac{\mathrm{Cov}(\tilde{C}_\ell^{\mathrm{ISW-gal}}, \tilde{C}_\ell^{\mathrm{gal-gal}})}{\langle \tilde{C}_\ell^{\mathrm{ISW-gal}} \rangle \langle \tilde{C}_\ell^{\mathrm{gal-gal}} \rangle} + \frac{\mathrm{Var}(\tilde{C}_\ell^{\mathrm{gal-gal}})}{\langle \tilde{C}_\ell^{\mathrm{gal-gal}} \rangle^2} \right)$$

$$= C_\ell^{\mathrm{ISW-gal}} \frac{\langle \tilde{C}_\ell^{\mathrm{ISW-gal}} \rangle}{\langle \tilde{C}_\ell^{\mathrm{gal-gal}} \rangle} \left(1 - \frac{2}{2\ell+1} + \frac{2}{2\ell+1} \right),$$

where we used

$$\mathrm{Var}(\tilde{C}_\ell^{\mathrm{gal-gal}}) = \frac{2}{2\ell+1} C_\ell^{\mathrm{gal-gal}},$$

$$\mathrm{Cov}(\tilde{C}_\ell^{\mathrm{gal-gal}}, \tilde{C}_\ell^{\mathrm{ISW-gal}}) = \frac{2}{2\ell+1} C_\ell^{\mathrm{gal-gal}} C_\ell^{\mathrm{ISW-gal}}.$$

The corrective terms vanish, and we find

$$\langle \tilde{C}_\ell^{\mathrm{ISW-rec}} \rangle \approx C_\ell^{\mathrm{ISW-gal}} \frac{\langle \tilde{C}_\ell^{\mathrm{ISW-gal}} \rangle}{\langle \tilde{C}_\ell^{\mathrm{gal-gal}} \rangle}$$

$$= \langle \tilde{C}_\ell^{\mathrm{ISW-rec,th}} \rangle. \quad \text{(B9)}$$

Then, on average, the cross-power between the true and reconstructed ISW maps is unchanged from the theory case. Since the multipoles are independent, this means the total covariance between the true and reconstructed ISW maps is unchanged as well:

$$\langle \widetilde{\mathrm{Cov}}(T^{\mathrm{ISW}}, T^{\mathrm{rec}}) \rangle = \frac{1}{4\pi} \sum_\ell (2\ell+1) \langle \tilde{C}_\ell^{\mathrm{ISW-rec}} \rangle$$

$$= \frac{1}{4\pi} \sum_\ell (2\ell+1) \langle \tilde{C}_\ell^{\mathrm{ISW-rec,th}} \rangle.$$

While for the autopower $\tilde{C}_\ell^{\mathrm{rec-rec}}$ we were able to derive an analytical result, a similar Taylor expansion treatment to the same order as the cross-power results in an additive correction of $2/(2\ell+1)$, or

$$\langle \tilde{C}_\ell^{\mathrm{rec-rec}} \rangle \approx \left(\frac{2\ell+3}{2\ell+1} \right) \langle C_\ell^{\mathrm{rec-rec,th}} \rangle, \quad \text{(B10)}$$

which is a good approximation to the analytical result found in Eq. (B6).

[1] R. Sachs and A. Wolfe, Astrophys. J. **147,** 73 (1967).
[2] W. Hu and N. Sugiyama, Phys. Rev. D **50,** 627 (1994).
[3] S. Boughn and R. Crittenden, Nature (London) **427,** 45 (2004).
[4] P. Fosalba, E. Gaztanaga, and F. Castander, Astrophys. J. **597,** L89 (2003).
[5] M. R. Nolta et al. (WMAP Collaboration), Astrophys. J. **608,** 10 (2004).

[6] P.-S. Corasaniti, T. Giannantonio, and A. Melchiorri, Phys. Rev. D **71**, 123521 (2005).

[7] N. Padmanabhan, C. M. Hirata, U. Seljak, D. J. Schlegel, J. Brinkmann, and D. P. Schneider, Phys. Rev. D **72**, 043525 (2005).

[8] P. Vielva, E. Martinez-Gonzalez, and M. Tucci, Mon. Not. R. Astron. Soc. **365**, 891 (2006).

[9] J. D. McEwen, P. Vielva, M. P. Hobson, E. Martinez-Gonzalez, and A. N. Lasenby, Mon. Not. R. Astron. Soc. **376**, 1211 (2007).

[10] T. Giannantonio, R. G. Crittenden, R. C. Nichol, R. Scranton, G. T. Richards, A. D. Myers, R. J. Brunner, A. G. Gray, A. J. Connolly, and D. P. Schneider, Phys. Rev. D **74**, 063520 (2006).

[11] A. Cabre, P. Fosalba, E. Gaztanaga, and M. Manera, Mon. Not. R. Astron. Soc. **381**, 1347 (2007).

[12] A. Rassat, K. Land, O. Lahav, and F. B. Abdalla, Mon. Not. R. Astron. Soc. **377**, 1085 (2007).

[13] T. Giannantonio, R. Scranton, R. G. Crittenden, R. C. Nichol, S. P. Boughn, A. D. Myers, and G. T. Richards, Phys. Rev. D **77**, 123520 (2008).

[14] S. Ho, C. Hirata, N. Padmanabhan, U. Seljak, and N. Bahcall, Phys. Rev. D **78**, 043519 (2008).

[15] J.-Q. Xia, M. Viel, C. Baccigalupi, and S. Matarrese, J. Cosmol. Astropart. Phys. 09 (2009) 003.

[16] T. Giannantonio, R. Crittenden, R. Nichol, and A. J. Ross, Mon. Not. R. Astron. Soc. **426**, 2581 (2012).

[17] P. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **571**, A19 (2014).

[18] P. A. R. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A21 (2016).

[19] D. J. Schwarz, C. J. Copi, D. Huterer, and G. D. Starkman, Classical Quantum Gravity **33**, 184001 (2016).

[20] J. Kim, A. Rotti, and E. Komatsu, J. Cosmol. Astropart. Phys. 04 (2013) 021.

[21] P. A. R. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A17 (2016).

[22] R. Barreiro, P. Vielva, C. Hernandez-Monteagudo, and E. Martinez-Gonzalez, IEEE J. Sel. Top. Signal Process. **2**, 747 (2008).

[23] B. R. Granett, M. C. Neyrinck, and I. Szapudi, Astrophys. J. **701**, 414 (2009).

[24] C. Francis and J. Peacock, Mon. Not. R. Astron. Soc. **406**, 14 (2010).

[25] R. B. Barreiro, P. Vielva, A. Marcos-Caballero, and E. Martinez-Gonzalez, Mon. Not. R. Astron. Soc. **430**, 259 (2013).

[26] A. Rassat, J. L. Starck, and F. X. Dupe, Astron. Astrophys. **557**, A32 (2013).

[27] A. Manzotti and S. Dodelson, Phys. Rev. D **90**, 123009 (2014).

[28] B. R. Granett, A. Kovcs, and A. J. Hawken, Mon. Not. R. Astron. Soc. **454**, 2804 (2015).

[29] J. Muir and D. Huterer, Phys. Rev. D **94**, 043503 (2016).

[30] L. Bonavera, R. B. Barreiro, A. Marcos-Caballero, and P. Vielva, Mon. Not. R. Astron. Soc. **459**, 657 (2016).

[31] A. R. Pullen and C. M. Hirata, Publ. Astron. Soc. Pac. **125**, 705 (2013).

[32] S. Ho *et al.*, Astrophys. J. **761**, 14 (2012).

[33] S. Ho *et al.*, J. Cosmol. Astropart. Phys. 05 (2015) 040.

[34] N. Agarwal, S. Ho, and S. Shandera, J. Cosmol. Astropart. Phys. 02 (2014) 038.

[35] T. Giannantonio, A. J. Ross, W. J. Percival, R. Crittenden, D. Bacher, M. Kilbinger, R. Nichol, and J. Weller, Phys. Rev. D **89**, 023511 (2014).

[36] N. Agarwal *et al.*, J. Cosmol. Astropart. Phys. 04 (2014) 007.

[37] C. Hernndez-Monteagudo *et al.*, Mon. Not. R. Astron. Soc. **438**, 1724 (2014).

[38] S. A. Thomas, F. B. Abdalla, and O. Lahav, Phys. Rev. Lett. **106**, 241301 (2011).

[39] C. Blake and J. Wall, Mon. Not. R. Astron. Soc. **329**, L37 (2002).

[40] C. Hernandez-Monteagudo, Astron. Astrophys. **520**, A101 (2010).

[41] R. Scranton *et al.* (SDSS Collaboration), Astrophys. J. **579**, 48 (2002).

[42] A. J. Ross *et al.*, Mon. Not. R. Astron. Soc. **417**, 1350 (2011).

[43] B. Leistedt, H. V. Peiris, D. J. Mortlock, A. Benoit-Lévy, and A. Pontzen, Mon. Not. R. Astron. Soc. **435**, 1857 (2013).

[44] B. Leistedt and H. V. Peiris, Mon. Not. R. Astron. Soc. **444**, 2 (2014).

[45] F. Elsner, B. Leistedt, and H. V. Peiris, Mon. Not. R. Astron. Soc. **456**, 2095 (2016).

[46] H. Awan, E. Gawiser, P. Kurczynski, R. L. Jones, H. Zhan, N. D. Padilla, A. M. Muñnoz Arancibia, A. Orsi, S. A. Cora, and P. Yoachim, Astrophys. J. **829**, 50 (2016).

[47] F. Elsner, B. Leistedt, and H. V. Peiris, Mon. Not. R. Astron. Soc. **465**, 1847 (2017).

[48] E. Suchyta *et al.* (DES Collaboration), Mon. Not. R. Astron. Soc. **457**, 786 (2016).

[49] D. Huterer, C. E. Cunha, and W. Fang, Mon. Not. R. Astron. Soc. **432**, 2945 (2013).

[50] D. L. Shafer and D. Huterer, Mon. Not. R. Astron. Soc. **447**, 2961 (2015).

[51] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, Astrophys. J. **622**, 759 (2005).

[52] N. Afshordi, Phys. Rev. D **70**, 083536 (2004).

[53] R. Laureijs *et al.* (EUCLID Collaboration), arXiv:1110.3193.

[54] M. Douspis, P. G. Castro, C. Caprini, and N. Aghanim, Astron. Astrophys. **485**, 395 (2008).

[55] N. Martinet, J. G. Bartlett, A. Kiessling, and B. Sartoris, Astron. Astrophys. **581**, A101 (2015).

[56] O. Doré *et al.*, arXiv:1412.4872.

[57] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000).

[58] N. Aghanim *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A11 (2016).

[59] P. A. R. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A15 (2016).

[60] E. S. Rykoff, E. Rozo, and R. Keisler, arXiv:1509.00870.

[61] N. Padmanabhan *et al.*, Astrophys. J. **674**, 1217 (2008).

[62] F.-X. Dupe, A. Rassat, J.-L. Starck, and M. Fadili, Astron. Astrophys. **534**, A51 (2011).

[63] W. Hu and R. Scranton, Phys. Rev. D **70**, 123002 (2004).

[64] M. Frommert, T. A. Ensslin, and F. S. Kitaura, Mon. Not. R. Astron. Soc. **391**, 1315 (2008).

[65] G.-C. Liu, K.-W. Ng, and U.-L. Pen, Phys. Rev. D **83**, 063001 (2011).

[66] J. Hartlap, P. Simon, and P. Schneider, Astron. Astrophys. **464**, 399 (2007).

[67] A. Taylor, B. Joachimi, and T. Kitching, Mon. Not. R. Astron. Soc. **432**, 1928 (2013).