Preparing for the cosmic shear data flood: Optimal data extraction and simulation requirements for stage IV dark energy experiments

Peter L. Taylor,[*] Thomas D. Kitching, and Jason D. McEwen

Mullard Space Science Laboratory, University College London,
Holmbury St. Mary, Dorking, Surrey RH5 6NT, United Kingdom

Upcoming photometric lensing surveys will considerably tighten constraints on the neutrino mass and the dark energy equation of state. Nevertheless it remains an open question of how to optimally extract the information and how well the matter power spectrum must be known to obtain unbiased cosmological parameter estimates. By performing a principal component analysis (PCA), we quantify the sensitivity of 3D cosmic shear and tomography with different binning strategies to different regions of the lensing kernel and matter power spectrum, and hence the background geometry and growth of structure in the Universe. We find that a large number of equally spaced tomographic bins in redshift can extract nearly all the cosmological information without the additional computational expense of 3D cosmic shear. Meanwhile a large fraction of the information comes from small poorly understood scales in the matter power spectrum, that can lead to biases on measurements of cosmological parameters. In light of this, we define and compute a cosmology-independent measure of the bias due to imperfect knowledge of the power spectrum. For a Euclid-like survey, we find that the power spectrum must be known to an accuracy of less than 1% on scales with $k \leq 7h$ Mpc^{-1}. This requirement is not absolute since the bias depends on the magnitude of modeling errors, where they occur in k-z space, and the correlation between them, all of which are specific to any particular model. We therefore compute the bias in several of the most likely modeling scenarios and introduce a general formalism and public code, `RequiSim`, to compute the expected bias from any nonlinear model.

I. INTRODUCTION

As photons travel from distant galaxies, their paths are gravitationally distorted by inhomogeneities in the gravitational field. This changes the ellipticities of observed galaxies, which on the largest scales is referred to as cosmic shear. Since this shear signal is sensitive to both the growth of structure and the background geometry of the Universe, by measuring the statistics of these distortions over a large number of galaxies it is possible to constrain cosmological parameters [1–4].

As the number of source galaxies in most surveys peaks in the relatively low redshift Universe, below $z = 2$, cosmic shear experiments are primarily sensitive to physics encoded by cosmological parameters that affect the late Universe. This makes weak lensing an ideal probe to distinguish between models of dark energy and determine the sum of the masses of neutrinos [5].

We are entering a golden age of Stage IV weak lensing experiments [6] as data from Euclid[1] [7], WFIRST[2] [8], and

LSST[3] [9] will be available within the next decade. Before these data sets arrive, it is important to determine the optimal method to extract cosmological information.

At present there are two proposed methods that use shear information in different ways: so-called "tomography" and "3D cosmic shear", proposed in [10,11], respectively. The former refers to identically weighting galaxies at all redshifts and compresses the data into tomographic bins, where all galaxies inside a certain redshift range are assigned the same redshift; using this approach, the expected errors on the dark energy equation of state parameters converge for approximately 10–20 bins [12]. Meanwhile the latter technique refers to a weighting scheme only, where data along the line of sight are given the spherical-Bessel weight that depends on radial and angular wave numbers within a single redshift bin. We discuss the motivation for this weight in Appendix A.

Comparing these two techniques is the first of objective of this paper. We compare tomography and 3D cosmic shear using the new Generalized Lensing and Shear Spectra (`GLaSS`) code, soon to be made publicly available [13]

[*]peterllewelyntaylor@gmail.com

as part of the Cosmosis [14] modular cosmological software package.

As no data compression takes place, it has been suggested that spherical-Bessel weighted lensing is more sensitive to radial information [11,15,16]. However, in tomography, compressing the data by coarsely binning in the radial direction is not strictly necessary. While it is true that for a very large number of redshift bins, shot noise will dominate in the intrabin power spectra ("autocorrelation") of bins, the interbin power spectra ("cross-correlation") between bins is free of shot noise and will contain the majority of the information. We investigate these issues in this paper.

In any case, stage IV experiments offer at least an order of magnitude improvement in *precision* over existing surveys. With increasing statistical precision it is important to keep systematics in check to ensure that measurements remain unbiased, to avoid far-reaching but incorrect conclusions. Potential sources of bias include photometric redshift errors [17], inaccurate intrinsic alignment models [18], and instrumental inaccuracies and uncertainties [19].

An additional uncertainty comes from the difficulty in modeling nonlinearities [20] and baryonic effects [21] at small scales (high k-modes), leading to inaccurate matter power spectrum models. There is no easy way to separate out the signal contributions from these small scales because small scale perturbations in the matter power spectrum at low redshift contribute to the same modes in the signal as larger scale perturbations at higher redshift. To avoid contamination from these poorly understood scales, it is imperative to understand how bias in power spectrum modeling propagates through to bias on the lensing signal itself. This is the second objective of this work.

In Sec. II we derive expressions for the signal and the noise for the most general weighted two-point statistic and show how tomography and 3D cosmic shear are both just special cases. Next we present a heuristic guide to the principal component analysis (PCA) formalism that we will use to assess the information content of these statistics. A more detailed exposition can be found in Appendix C. We also show how biases in the matter power spectrum modeling can be summarized in terms of a "knowledge matrix" (that encodes the level of knowledge one has about the matter power spectrum model, including correlations between k-modes and redshift) and present a cosmological parameter-independent measure of the bias in the lensing signal due to matter power spectrum modeling errors. In Sec. IV, we determine how 3D cosmic shear and tomographic lensing are sensitive to the matter power spectrum and the lensing kernel, using the formalism presented in Sec. II. We show the bias in the lensing statistics induced by biases in the matter power spectrum for a variety of knowledge matrices corresponding to realistic cases.

Our code, RequiSim, used in this last part is made publicly available and can be used to assess whether any given matter power spectrum simulation or model is accurate enough for a Euclid-like lensing survey. Finally in the Appendices we outline our modeling choices, discuss motivations for the spherical-Bessel weight in 3D cosmic shear, provide details about PCA formalism, address the challenges of removing sensitivity to small scales in the matter power spectrum, and provide information about RequiSim.

II. COSMIC SHEAR FORMALISM

A. The generalized spherical transform

The shear field is defined everywhere, but it can only be sampled at the position of galaxies. We can transform the sampled shear field into the spherical-Bessel basis. This is commonly referred to as "3D cosmic shear" or "3D weak lensing." In this case the shear field is given by

$$\gamma_{\ell m}(\eta) = \sqrt{\frac{2}{\pi}} \sum_g \gamma_g(r_g, \boldsymbol{\theta}_g) j_\ell(\eta r_g) {}_2Y_{\ell m}(\boldsymbol{\theta}_g), \quad (1)$$

where the sum is over all galaxies g, $j_\ell(\eta r_g)$ are spherical Bessel functions and ${}_2Y_{\ell m}(\boldsymbol{\theta}_g)$ are spin-2 spherical harmonics. Motivations for the Bessel weight are discussed in Appendix A. Here we explicitly write the harmonic variable as η, so that it is not confused with k, which is used to denote the wave number in the matter power spectrum only.

As discussed in [22], we could also weight the data by an arbitrary weight function, W. In [22] this is taken to be a function of the comoving distance r only, but in general the weight function can also depend on a radial and angular wave number η and ℓ: $W_\ell(\eta, r)$. Weights were also considered in [23,24], but we consider a more general formalism. Replacing the Bessel functions with a general weight, we define the generalized spherical transform given by

$$\gamma_{\ell m}(\eta) = \sqrt{\frac{2}{\pi}} \sum_g \gamma_g(r_g, \boldsymbol{\theta}_g) W_\ell(\eta, r_g) {}_2Y_{\ell m}(\boldsymbol{\theta}_g), \quad (2)$$

where η is a label that can be a wave number or a real-space quantity. The expression for the lensing matter power spectrum becomes

$$C_\ell^{\gamma\gamma}(\eta_1, \eta_2) = \frac{9\Omega_m^2 H_0^4}{16\pi^4 c^4} \frac{(\ell+2)!}{(\ell-2)!} \int \frac{dk}{k^2} G_\ell^\gamma(\eta_1, k) G_\ell^\gamma(\eta_2, k),$$

$$(3)$$

where Ω_m is the fractional energy density of matter, c is the speed of light in a vacuum, and H_0 is the present day Hubble constant. The G-matrix is defined as

$$G_\ell^\gamma(\eta, k) \equiv \int dz_p dz' n(z_p) p(z'|z_p)$$
$$\times W_\ell(\eta, r[z']) U_\ell(r[z'], k) \quad (4)$$

and the U-matrix, which contains all the cosmological information, is given by

$$U_\ell(r[z],k) \equiv \int_0^r dr' \frac{F_K(r,r')}{a(r')} j_\ell(kr') P^{1/2}(k;r').$$ (5)

In the above expressions $n(z)$ is the radial distribution of galaxies; photometric uncertainty $p(z|z')$ gives the probability that a galaxy has a redshift z, given a photometric redshift measurement z'; $P(k,r)$ is the matter power spectrum; r is the comoving distance; and the lensing kernel, $F_K(r,r')$, is defined as

$$F_K(r,r') \equiv \frac{r-r'}{rr'},$$ (6)

for a flat cosmology. We are implicitly assuming an equal-time power spectrum throughout. This has been found to be a good approximation [25].

The weights also propagate to the shot noise [22], which becomes

$$N_\ell^{ee}(\eta_1,\eta_2) = \frac{\sigma_e^2}{2\pi^2} \int dz\, n(z) W_\ell(\eta_1,r) W_\ell(\eta_2,r),$$ (7)

where σ_e^2 is the variance of the intrinsic ellipticities in galaxies. We take $\sigma_e = 0.3$ throughout [26].

When taking the weight function, $W_\ell(\eta, r(z)) = j_\ell(\eta r)$ in Eqs. (4) and (7), the normal 3D cosmic shear equations are recovered. For the bin associated with redshift region I, in the tomographic case, we take the weight function, W^I, in Eq. (4) as a top hat function in redshift only, so that

$$W^I(z) \equiv \begin{cases} 1 & \text{if } z \in I \\ 0 & \text{if } z \notin I \end{cases}$$ (8)

and the shot noise reduces to

$$N_{ij} = \frac{\sigma_e^2}{2\pi^2} N_i \delta_{ij},$$ (9)

where i and j label the bin numbers and N_i is the number of galaxies in bin i.

The expressions above can be simplified further. As found in [27], the extended Limber approximation, given in [28], is sufficiently accurate for $\ell > 100$. Then the U-matrix becomes

$$U_\ell(r,k) = \frac{F_k(r,\nu(k))}{ka(\nu(k))} \sqrt{\frac{\pi}{2(\ell+1/2)}} P^{1/2}(k,\nu(k)),$$ (10)

where $\nu(k) \equiv \frac{\ell+1/2}{k}$. This implies that for fixed ℓ-modes above $\ell \approx 100$, the signal is sensitive to the power spectrum only along the curve $\ell + 1/2 = kr$ [28]. We refer to these curves as "Limber lines" in the $(k, r[z])$ plane (they are plotted in Fig. 9 for different ℓ-modes).

The prefactors used in the equations above for the signal and noise do not follow the standard notation. However once we have made the Limber approximation [given in Eq. (10)], our expressions reduce to the usual ones given in [10,29]. Our convention is taken to be consistent with the one used in [11].

B. A review of cosmic shear likelihood analyses

We review the details of the likelihood analysis used to extract cosmological information from cosmic shear data. This has been used successfully to constrain cosmological parameters in [2,4,30].

Assuming the likelihood is Gaussian, the likelihood for a set of parameters, p, is

$$\ln \mathcal{L}(p) = -\frac{1}{2} \sum_{a,b} [D_a - T_a(p)] C_{ab}^{-1} [D_b - T_b(p)],$$ (11)

where D_a is a data vector, T_a is a theory vector, and C_{ab}^{-1} is the inverse covariance matrix between elements of the data vector.

The covariance can be computed from ray-tracing simulations [2], bootstrapped directly from the data [31], computed from fast lognormal simulations [4,32,33], or by approximating the shear field as a Gaussian random field [34]. A thorough discussion of the relative merits of each of these approaches is given in [2].

In a power spectrum analysis, the data vector is the spectrum of the observed generalized spherical-transformed shear field defined in Eq. (2). Meanwhile the theory vector is computed from Eq. (3).

In practice it is difficult to account for the impact of mask on the theory vector. For this reason, most analyses take the data and theory vectors as two-point correlation functions $\xi_\pm$. This data vector is readily computed from a shear catalogue and the theory vector is found by first computing the raw lensing spectrum and then applying a filter. More details can be found in [27]. The filter preferentially weights the sensitivity to certain modes in the lensing spectrum. We ignore this complication in this work and focus on the information contained in the raw lensing power spectrum.

Once the covariance matrix is known, we can sample to compute the posterior of the cosmological parameters. The computation time is dominated by the calculation of the lensing spectrum and in particular the computation of the matter power spectrum. It is infeasible to run a full N-body simulation at each point in parameter space so we must resort to fast emulator simulations to compute the matter power spectrum. These suffer from a loss of accuracy at small scales.

In the remainder of this paper we try to answer the following two questions:
  (i) Assuming that the covariance matrix can be computed to sufficient accuracy, what is the optimal

transform weight in Eq. (2)? If we further restrict our attention just to the tomographic case, what is the optimal binning strategy?

(ii) Since the theory vector depends on the matter power spectrum, to what accuracy does this have to be known to obtain an unbiased cosmological measurement?

C. Tomographic binning strategy

In this paper we investigate two different tomographic binning schemes: equal galaxies per bin and equally spaced redshift bins. These are shown for 10-bin tomography in Fig. 1. We have overlaid the bin boundaries onto the predicted Euclid wide-field differential galaxy distribution, defined by $\frac{\partial N(z)}{\partial z} \equiv N_g \frac{\partial n(z)}{\partial z}$, where N_g is the number of galaxies in the survey. An equal number of galaxies per bin is the convention; and for less than 20 bins, it captures more information than equally spaced bins in redshift (see Fig. 4 and the discussion in Secs. IV A and IV B). However, to completely capture the 3D information, around 100 bins must be used (see Secs. IV A and IV B). This presents two challenges for the equal number of galaxies per bin scheme:

(i) Since the distribution of galaxies, $n(z)$, falls sharply for high redshifts, the last bin will span a large z-range. Thus high redshift information will be lost unless an incredibly large number of bins are used.

(ii) Meanwhile if more than 20 bins are used, the width of the tomographic bins near the peak of $n(z)$ becomes very small. Placing an equal number of galaxies into each bin would require globally increasing the resolution, slowing the computation time.

For the later reason, we will only consider tomography with an equal number of galaxies per bin, up to 20 bins.

While it is unconventional to use over 20 bins in a likelihood analysis, we stress that it barely increases the computation time of the theory vector. Using our algorithm, it takes 1.55 seconds to compute the lensing spectrum once the matter power spectrum has been computed for 100 bins, compared to 1.40 seconds for 10 bins on a single 2.7 GHz Intel i5 Core on a 2015 Macbook Pro with 8 GB of RAM. We stress that although it is easy to compute the theory vector with a large number of bins, it may be difficult to compute a significantly accurate inverse covariance due to numerical noise [2].

III. PRINCIPLE COMPONENTS AND BIAS

A. A heuristic guide to principle component analysis

Assuming a fiducial cosmology it is possible to compute the lensing spectra given in Eq. (3). Then for any set of parameters, $\{\theta_i\}$, we can estimate the constraining power of a lensing survey using the Fisher matrix formalism. Details are given in Appendix B.

The parameters are often degenerate. For example, the purple oval in Fig. 2 represents the 2σ confidence limit on
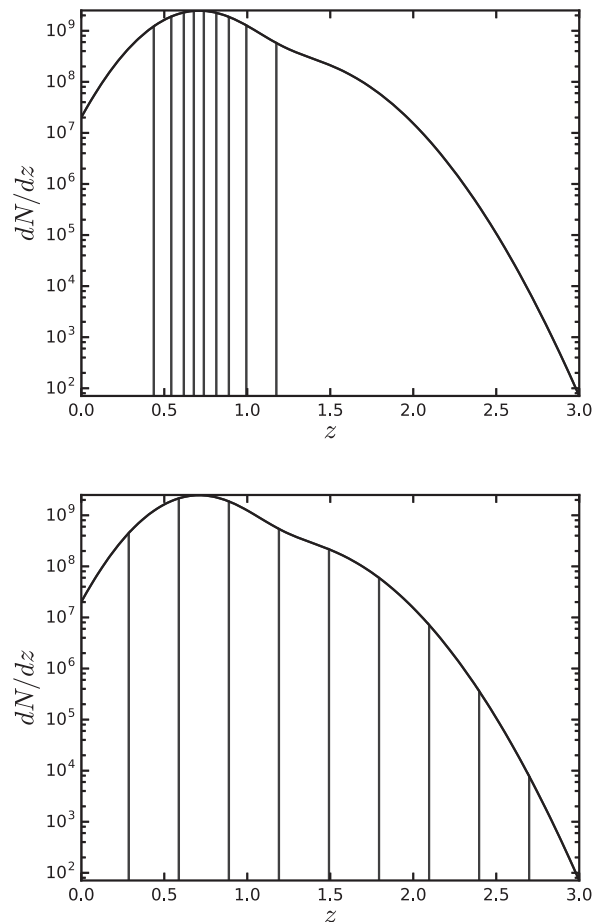

FIG. 1. Differential number of galaxies as a function of redshift for two different tomographic binning strategies. Top: 10 bins with an equal number of galaxies per bin. Equal galaxy binning captures most of the information with a small number of bins (see Fig. 4), but an extremely large number of bins would be needed to capture the high-redshift information since the highest redshift bin is so wide. For a larger number of bins, near the peak of the distribution the bin width would fall below the resolution of the computation grid. Bottom: 10 bins spaced equally in redshift. This binning strategy does not capture as much information for a small number of bins (see Fig. 4), but it is easy to increase the number of bins.

parameters θ_1 and θ_2 estimated from a Fisher matrix analysis. Nevertheless, it is possible to rotate into a new basis of parameters, $\{A_i\}$, which are independent, as in the second panel of Fig. 2.

In this paper we will divide the comoving distance $r(z)$ and the temporally evolving power spectrum $P(k, z)$ into cells in z and z-k, respectively, closely following the analysis of [35]. We estimate constraints on the amplitude of these cells from cosmic shear using the Fisher matrix formalism. As in the example above, the amplitudes of these cells are expected to be degenerate, so we find new linear combinations of these cells called components, whose amplitude are independent. Thus the components extract independent information.
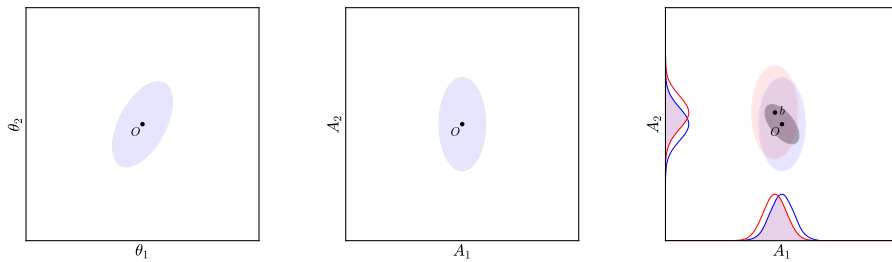
FIG. 2. This figure illustrates the procedure of calculating the principle components (see Sec. III A) and the variance weighted overlap (see Sec. III C), which measures the bias in higher-dimensional spaces. Left: Parameter constraints for two degenerate parameters θ_1 and θ_2 from a cosmic shear experiment. Middle: Rotation into a new basis of parameters that are independent. Right: A bias, b, is drawn from the dark hashed oval representing the covariance on our modeling uncertainty. The new shifted covariance, centered on b, is represented by the red oval. The overlap integral between the shifted and unshifted distribution marginalized onto each direction is shown. Weighting these by the inverse variance on parameters A_1 and A_2 and marginalizing over all shifts defines the variance weighted overlap $\mathrm{R}(P_{\mathrm{VWO}})$.

The components that are most tightly constrained are called the principle components (PCs). The sum of the inverse variance of each of the components, which we write as I_{tot}, is called the "total information content" and is a figure of merit for the constraining power of an experiment.

We investigate how many PCs are needed to extract the majority of the information. Ordering the components from the most to the least tightly constrained, we compute the information fraction, extracted as a function of the number of principle components.

Since we are interested in where information from the lensing kernel [Eq. (6)] and power spectrum come from in k-z space, we compute the Cramer-Rao bound on the amplitude inside each cell [$\sigma_{\mathrm{CRB}}(r(z))$ and $\sigma_{\mathrm{CRB}}(P(k,z))$ for comoving distance and power spectrum cells, respectively]. We also compute the variance weighted sum (S_{vw}) on the amplitude of each cell. The first measures the inverse conditional error on each cell, while the second measures how tightly different regions of the comoving distance and power spectrum are constrained, while accounting for correlations between cells. More details can be found in the Appendices C and D.

Since the information about the comoving distance is contained in the lensing kernel, we refer to the comoving distance PCs as lensing kernel PCs for the remainder of the text.

B. Review of the overlap integral formalism

As well as understanding where lensing information originates, we must also understand the bias. The statistical error expected from upcoming surveys is fixed by the survey volume and the number of observed galaxies. To obtain unbiased results without increasing the statistical error, contributions from all sources of bias must be kept below a certain threshold. We review the formalism in [19] (hereafter M12), which defines this threshold.

M12 considered an experiment which measured a parameter with a Gaussian likelihood and statistical

uncertainty σ. The bias, b, shifted the likelihood distribution, but did not change its shape. The distance between the two distributions was quantified by the overlap integral between the shifted and unshifted distributions:

$$p_{\mathrm{overlap}}(b) = 1 - \mathrm{erf}\left(\frac{1}{2\sqrt{2}}\frac{|b|}{\sigma}\right). \qquad (12)$$

If this was greater than 0.95 (or less conservatively 0.90), then the results are said to be *unbiased*.

However, if the exact value of the bias was known, it could be subtracted off, so the authors of M12 reinterpreted b as a 95% confidence limit on the true bias. If the knowledge of the bias is also normally distributed, its standard deviation is then $\sigma_b = b/2$. Marginalizing over $p_{\mathrm{overlap}}(b)$ by drawing b values from this distribution and using the same overlap criteria as before defines requirements on the magnitude of the bias $|b|$. These are $|b| < 0.31\sigma(0.62\sigma)$ for a 95%(90%) overlap.

M12 used this formalism to place requirements on the total systematic bias needed for an unbiased measurement of the dark energy equation of state. However, stage IV cosmic shear experiments will place new constraints on other interesting parameters like the sum of neutrino masses and w_a in the Chevalier-Polarski-Linder parametrization [36,37]. We need to ensure that these, and parameters in any other cosmological parametrization, will not be biased from inaccurate models of the power spectrum.

Perturbations to the lensing kernel and matter power spectrum will not have the same impact on the lensing spectrum. Thus we assume that these two varieties of PCs are only weakly correlated. Then, to check whether inaccurate power spectrum models induce bias, it is sufficient to ensure that power spectrum PCs are unbiased, and we can ignore bias propagating into the lensing kernel PCs. This is done by generalizing the 1D overlap integral formalism to higher dimensions in the next section.

C. The variance weighted overlap

In analogy with the 1D case, we compute a higher-dimensional overlap integral between a biased and unbiased distribution. We envision a high-dimensional parameter space of PC amplitudes and we will measure a generalization of the overlap between biased and unbiased probability distributions of these amplitudes.

Since the parameters of interest are the power spectrum PCs, the distribution of biases in the unbiased case is taken to be the multivariate Gaussian with a mean of zero and a covariance calculated from inverting the diagonal Fisher matrix in PCA space.

The distribution in the biased case has a shift in the mean (the multivariate equivalent of $|b|$) that is drawn from a Gaussian with mean zero and covariance K. We refer to this covariance as the "knowledge matrix", since it describes our confidence in our knowledge of the power spectrum. The 2σ confidence region from this covariance is represented by the dark hashed ellipse in Fig. 2. We interpret b_i as a 95% confidence limit on the true bias on principal component i.[4] The elements of K are then defined by

$$K_{ij} \equiv \frac{b_i}{2}\frac{b_j}{2}. \qquad (13)$$

The values of the resulting marginalized overlap integral are uninformative because they depend on the number of PCs included beyond those which contain 99% of the information (intuition about hypervolumes in higher-dimensional spaces is often wrong [38]).

We instead define a new measure called the variance weighted overlap ($P_{\rm VWO}$), which reduces to the marginalized overlap integral in 1D used in [19]. We draw shifts in the mean from a multivariate Gaussian with the covariance given by the knowledge matrix. Then, instead of computing the hypervolume of the overlap region, we marginalize to compute the 1D overlaps in each direction, forming a set of 1D overlap volumes $\{p^i\}$ for each principal component i.

Since not all PCs are equally important, we take an inverse variance weighted average over this set to form the variance weighted probability overlap given by

$$P_{\rm VWO} \equiv \frac{\sum_i w_i p^i}{\sum_i w_i}, \qquad (14)$$

where w_i is the inverse variance on principal component i. Keeping consistency with M12, if $P_{\rm VWO} > 95\%$, then the result is said to be unbiased. The requirements using this formalism are described in Sec. IV C.

This procedure is illustrated in the right panel of Fig. 2. The blue oval represents the unshifted measured covariance

---

[4]In practice the knowledge matrix should be found in k-z space, where it is known for a given simulation, and rotated into PCA space.

centered on the origin O. A bias, b, is drawn from the dark hashed oval representing the covariance on our modeling uncertainty. The new shifted covariance, centered on b, is represented by the red oval. Marginalizing onto each axis, we compute the overlap integral between the shifted and unshifted distribution on each parameter. After weighting these by the inverse variance on each component and marginalizing over all shifts, we find the variance weighted overlap ($P_{\rm VWO}$).

IV. RESULTS

A. Lensing kernel PCs

We determine the lensing kernel PCs for tomography with an equal number of galaxies per bin, tomography with equally spaced z bins and 3D cosmic shear with $\ell_{\max} = 3000$. We will refer to equally spaced z-bin tomography with 100 bins as "supertomography". The PCAs for supertomography with smaller ℓ-mode cuts are also found. Our modeling choices (see Appendix E) mimic the Euclid wide survey. The PCs are found in 40 redshift slices with $z \in (0., 3.0]$.

The supertomographic PCs are shown in Fig. 3, along with the fiducial comoving distance which is actually being constrained, the Cramer-Rao bound, the variance weighted sum, and a plot of the fraction of the information content captured by the first N PCs. From the Cramer-Rao bound and variance weighted sum, it is clear that Euclid will primarily be sensitive to the lensing kernel for redshifts in the range $z \in (0.1, 1.0)$. Also, only three PCs are needed to capture 92% of the information.

The PCs for 10-bin tomography with an equal number of galaxies per bin and 3D cosmic shear look nearly identical to the supertomographic ones plotted in Fig. 3. However, slightly less information is captured for both of the other analyses. This is summarized in Table I, Case I.

3D cosmic shear should capture as much information as supertomography since no radial data compression takes place. This is not the case in our analysis, where 3D cosmic shear captures 6% less information than supertomography. This is because in order to calculate the lensing spectra quickly, we use a low-resolution computation grid. In Appendix F, we investigate using a higher-resolution computation grid using lower-resolution PCs. At the highest resolution we considered, only 3% of the information is lost to numerical noise (Table I, Case II).

In 10-bin tomography, 6% of the lensing kernel information is lost due to inherent data compression. Using more bins would help reduce this number and we examine the convergence of the total information content for different binning strategies in Fig. 4. Initially an equal number of galaxies per bin converges more quickly, but it is infeasible for a large number of bins and information is lost at high redshifts (see Sec. II C). Meanwhile an equal
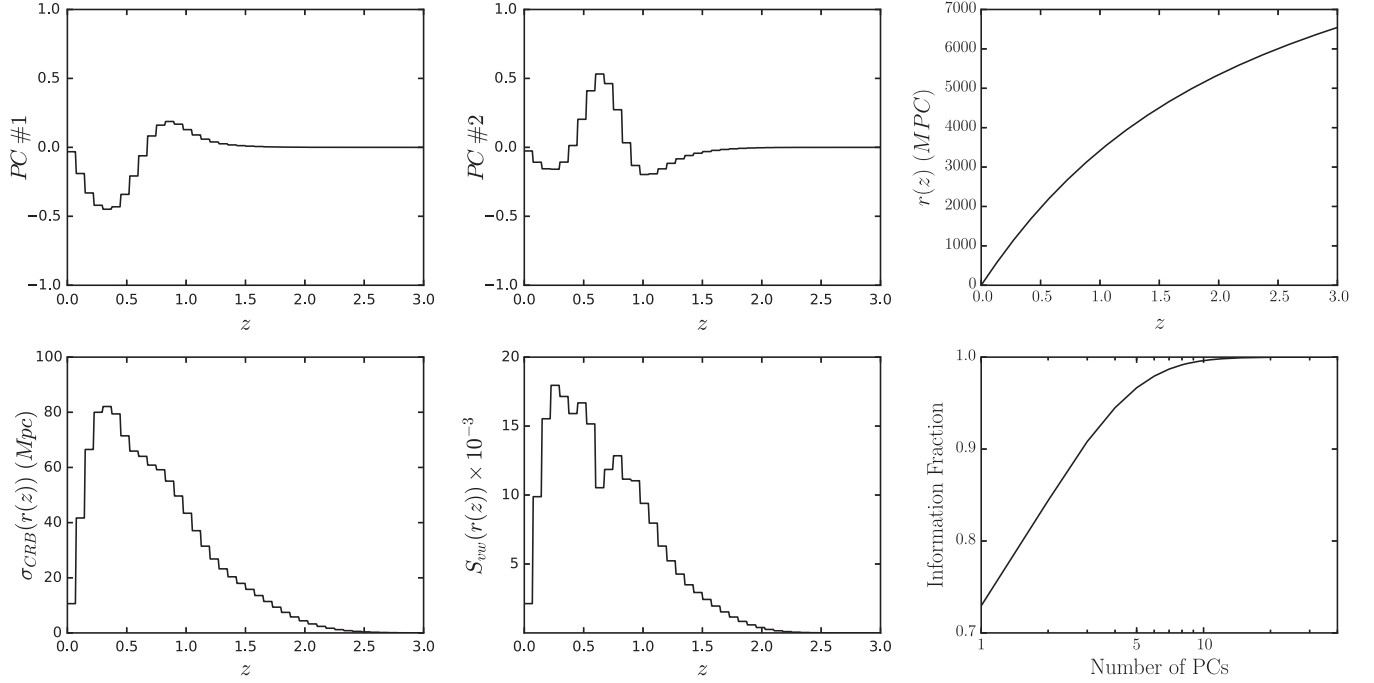
FIG. 3.   The first two supertomographic lensing kernel principal components, the transverse comoving distance for the fiducial cosmology, the Cramer-Rao bound, the variance weighted sum, and the information fraction for the number of principal components. See Sec. III A for an explanation of the terminology. From S_{vw} and the Cramer-Rao bound, this method is primarily sensitive to $z \in [0., 1.5]$. Three PCs capture the majority of the information.

redshift spacing binning strategy captures 99%(99.9%) with 50(90) bins.

The relative difference in the Cramer-Rao bound between tomography and supertomography is shown in Fig. 5. This confirms that 10-bin tomography loses information at high redshifts, beyond $z = 1$, as expected. This difference is unimportant for constraining the dark energy equation of state parameters—that only become relevant below $z \sim 1$, explaining the quick convergence of dark energy constraints with the number of tomographic bins found in [12]. However, these higher redshifts are precisely

where we expect a cross-correlation signal with CMB lensing. Cross-correlating with the CMB lensing signal will help bring lensing systematics under control [39] to substantially improve the dark energy and neutrino mass figure of merit [15].

Finally we assess the impact of angular scale cuts. Taking ℓ-cuts reduces the sensitivity to the lensing kernel. This is shown in Table II. When taking ℓ-cuts, information is primarily lost at intermediate redshifts, near $z = 0.5$. This can be seen from Fig. 6, where the Cramer-Rao bounds on the lensing kernel for different ℓ-cuts are plotted.

TABLE I.   The total information content, I_{tot} [see Eq. (C3)], in the lensing kernel and the matter power spectrum. Specifically we compare 3D cosmic shear, supertomography (100 equally space bins in redshift), and 10-bin tomography in *GLaSS*. Case I: High-resolution PCA grid and low-resolution computation grid. Supertomography outperforms 10-bin tomography because there is no data compression in the former case. It also outperforms 3D cosmic shear due, but only due to slow convergence of the latter technique (see Case II). Case II: Low-resolution PCA grid and high-resolution computation gird. Numerically, 3D cosmic shear has converged to within 3% of supertomography.

	Case I			Case II	
PCA run statistic	Supertomography	3D[a]	Supertomography	3D[a]	10-bin tomography
I_{tot} (lensing kernel)	1.00	0.94	0.94	1.00	0.97
I_{tot} (power spectrum)	1.00	0.77	0.99	1.00	0.97

[a]Due to the slow numerical convergence of 3D cosmic shear, the total information content of is not fully converged for 3D lensing case (see Appendix F and Fig. 13).
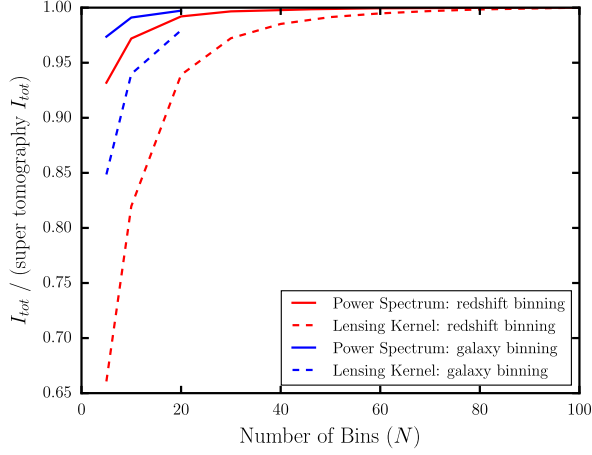
FIG. 4.   Fraction of the total information content, I_{tot}, of super-tomography for different binning strategies as a function of the number of tomographic bins. Red: Equally spaced bins in redshift. Blue: Equal number of galaxies per bin. Solid lines: Power spectrum information fraction. Dashed lines: Lensing kernel information fraction. An equal number of galaxies per bin initially converges faster, but this binning scheme loses information at high z and is intractable for a large number of bins (see Sec. II C). A total of 99%(99.9%) of the information is captured from both the lensing kernel and the power spectrum with 50(90) bins.

B. Power spectrum PCs

We compute the power spectrum PCs for tomography with an equal number of galaxies per bin, tomography with equally spaced z bins, and 3D cosmic shear with
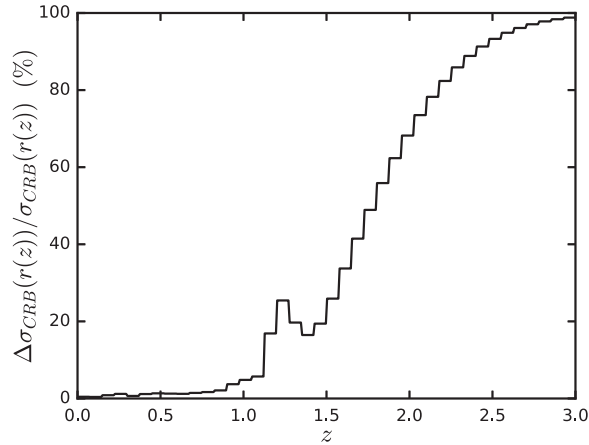

FIG. 5.   The relative difference in the Cramer-Rao bounds, on the lensing kernel, between supertomography and 10-bin tomography with an equal number of galaxies per bin. 10-bin tomography primarily loses information at high z (see Sec. II C). At redshifts above $z = 2$, more than half the lensing kernel information is lost in the 10 bin case. The difference in the total information content extracted remains small since most of the information comes from redshifts below $z = 1.5$. The small peak in the plot at $z = 1.3$ is due to the coarseness of 10-bin tomography, and we have checked that this feature disappears when more bins are used.

TABLE II.   Relative information content, I_{tot} [see Eq. (C3)], for the lensing kernel and the power spectrum for 10-bin tomography with different ℓ-cuts. The maximum values of ℓ are denoted as ℓ_{max}. For cuts below $\ell = 1000$, $\sim 50\%$ of the information is lost.

ℓ_{max}	3000	2500	2000	1500	1000	500
I_{tot} (lensing kernel)	1.00	0.94	0.86	0.75	0.60	0.38
I_{tot} (power spectrum)	1.00	0.94	0.84	0.71	0.51	0.26

$\ell_{max} = 3000$. PCs for supertomography with lower ℓ-cuts are also found. To save time, we compute these on a coarse grid before zooming in on the region of primary sensitivity by first perturbing the matter power spectrum on a 5×5 grid logarithmically spaced in k and linearly spaced in z. More than 95% of the signal is contained in the first two z bins and the last two k bins. The PCs are then computed on a 10×10 grid just inside this region.

The resulting PCs for supertomography are shown in Fig. 7 and the total information content is displayed in Table I, Case I. A high-resolution supertomography PCA run, on a 20×20 grid, is plotted in Fig. 9. Due to memory constraints, a high-resolution run is not done for 3D cosmic shear.

The supertomographic PCs look very similar to those of 3D cosmic shear (not shown), but since the eigenvalues are larger, more information is extracted (see Table I) in the former case, but just like in the lensing kernel case, this is due to the slow numerical convergence of 3D cosmic shear. The ratio of the information contents of 3D cosmic shear and supertomography is plotted in Fig. 13 on a two-by-two PCA grid. By a resolution of $N = 5000$, the relative information content of 3D cosmic shear is within 3% of supertomography. This is displayed in Table I, Case II.
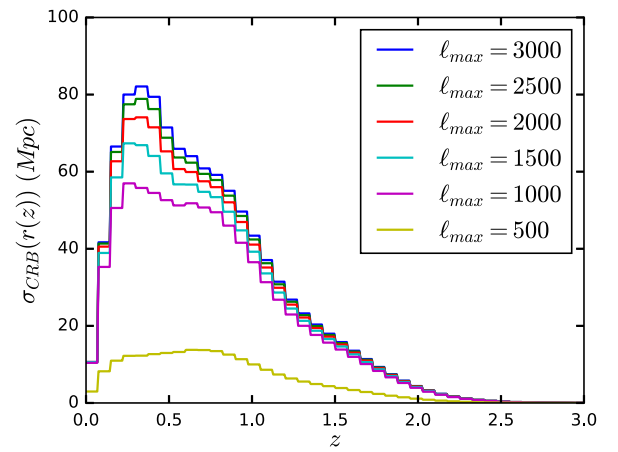

FIG. 6.   The supertomographic Cramer-Rao bound on the lensing kernel with different ℓ-cuts. Information is primarily lost at intermediate redshifts, near $z = 0.5$, and catastrophically lost below $\ell_{max} < 1000$.
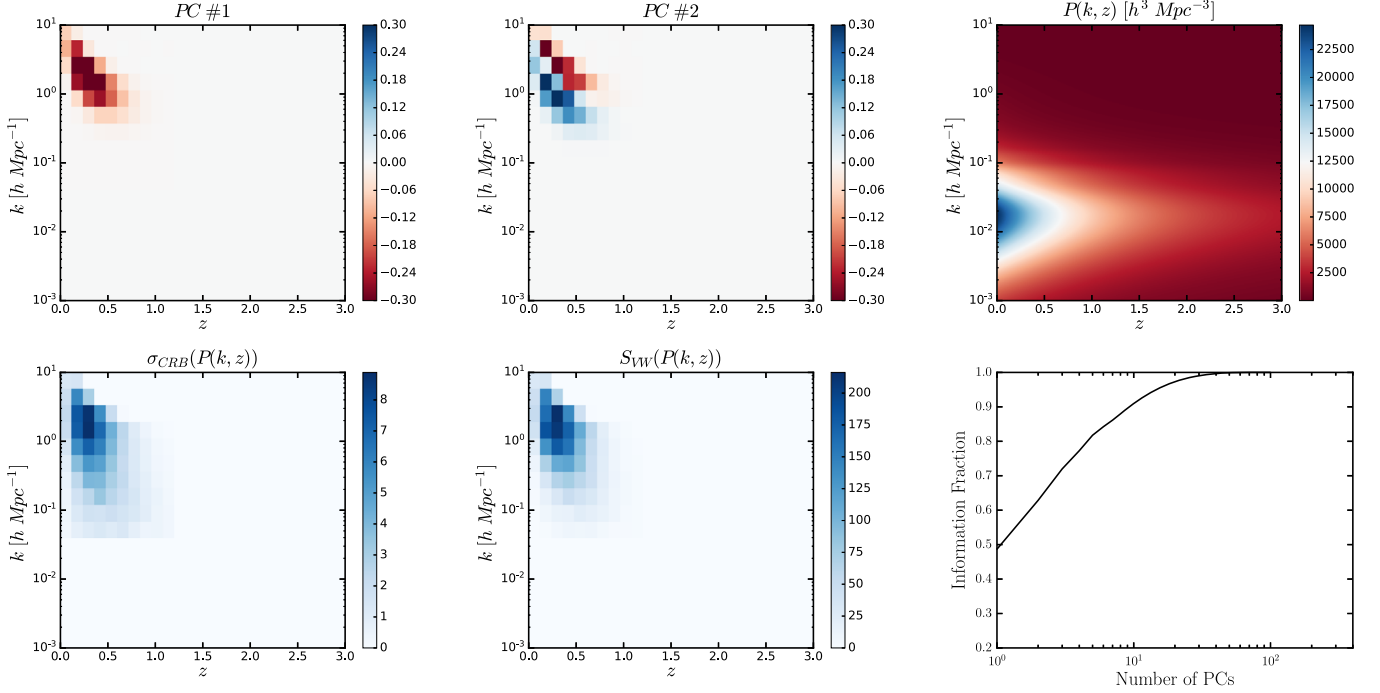
FIG. 7. The first two supertomographic lensing power spectrum PCs, the power spectrum of the fiducial cosmology, the Cramer-Rao bound, the variance weighted sum, and the information fraction. See Sec. III A for an explanation of the terminology. Cosmic shear is primarily sensitive to the power spectrum at large scales above $0.5h$ Mpc^{-1} and at low redshift. Nearly half the signal lies above $k = 1.5$ hMpc^{-1}, where the power spectrum becomes hard to model. Within principal components, bins are strongly correlated across this cut so power spectrum modeling errors at high k induces bias at low k.

The shape of the PCs is not surprising. The first principal component is a relatively broad feature following the Limber line $\ell \sim 1500$. Meanwhile the higher PCs show multiple broad features tracing multiple Limber lines. This is particularly noticeable for the higher-resolution PCs in Fig. 9. In the absence of shot noise, each ℓ-mode is independent and sensitive to the power spectrum along the Limber lines. However, shot noise induces correlation between Limber lines (ℓ-modes), as the uncertainty on neighboring bins in the power spectrum is correlated, causing the broad features in the PCs. We have performed a test without shot noise, and the principal components trace separate Limber lines exactly, without broadening, as expected.

The lensing kernel ensures that cosmic shear experiments are most sensitive to regions of the matter power spectrum that are at half the comoving distance of the bulk of the source galaxies in the survey. Along with the temporally evolving amplitude of the power spectrum, this ensures that lensing is primarily sensitive to the power spectrum in the redshift range $z \in [0.1, 0.6]$ (see Fig. 7).

The convergence of total information content with different binning strategies is plotted in Fig. 4. For an equal redshift spacing binning strategy, 99%(99.9%) of the power spectrum information is captured in 20(60) bins.

Meanwhile, around half of the signal to the power spectrum lies above $k = 1.5$ hMpc^{-1}. At such small scales, the power spectrum is difficult to model, and there is usually a modeling error of around 10% [40]. High- and low-k bins are correlated along the PCs, so a modeling error on the former will induce a bias in the later. We quantify the bias for different ℓ-mode cuts in Sec. IV D.

Taking an ℓ-mode cut removes sensitivity above a given k-cut. This is because low-ℓ Limber lines only lie above the k-cut at low redshift, where the sensitivity is suppressed by the lensing kernel, which peaks at half the distance to the peak of the galaxy distribution near $z = 0.7$. This is shown in Fig. 9, where we project the inverse error onto the z- and k-axes, taking the average. Cuts below $\ell = 1000$ significantly reduce the sensitivity to scales smaller than $k = 1.5h$ Mpc^{-1}, but around half the sensitivity to the power spectrum is lost (see also Table II). Meanwhile, the amount of information gained by including more ℓ-modes slows rapidly above $\ell = 2500$.

C. Correlations in power spectrum bias

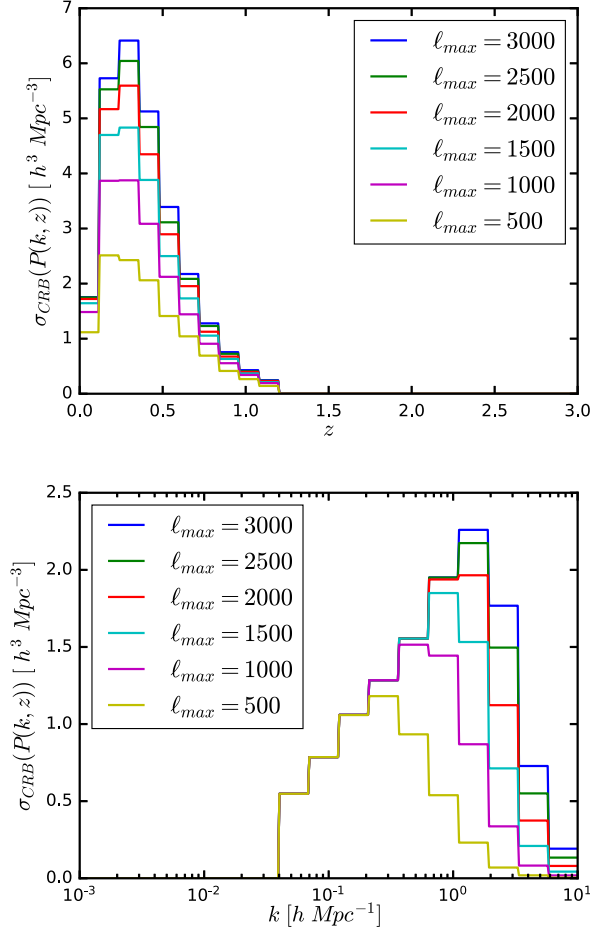We use PCs for supertomography for the remainder of Sec. IV, since these fully capture the 3D information.

FIG. 8. The average Cramer-Rao projected onto the z-axis (above) and k-axis (below) for different ℓ-mode cuts in 10-bin tomography. Cuts below $\ell \sim 1000$ significantly reduce the sensitivity to k-modes above $1.5h$ Mpc^{-1}.

It is clear that for cosmic shear studies not all regions of the matter power spectrum are equally important. Since lensing is barely sensitive to the power spectrum above $k = 10h$ Mpc^{-1}, modeling errors at smaller scales can be large without inducing bias on parameters inferred from cosmic shear. Meanwhile, accurately modeling the power spectrum near $k = 1.5h$ Mpc^{-1}, where the sensitivity peaks, is extremely important. However there is a more subtle effect that can dramatically change the requirement on the accuracy of power spectrum models. This is the degeneracy in modeling errors between regions in the power spectrum.

In Sec. III C we defined the variance weighted overlap ($P_{\rm VWO}$) as a measure of bias, which depends on the correlation between biases in different regions. These correlations describe what we call the "shape of the knowledge matrix". This is illustrated in two dimensions by the ellipticity and orientation of the dark hashed ellipse in Fig. 2. We plot the effect of having differently shaped knowledge matrices in Fig. 10. Here the $P_{\rm VWO}$ as a function

of the variance is plotted as a function of the bias normalized against the statistical variance: $\langle b^2 \rangle / \langle \sigma^2 \rangle$.[5]

Four cases are considered. The first two are extremes, while the second two are more realistic:

(i) K has the same shape as the PC covariance C, appropriately normalized. $P_{\rm VWO} > 0.95$ requires $\langle b^2 \rangle / \langle \sigma^2 \rangle < 0.15$.

(ii) K has the same shape as the inverse PC covariance, C^{-1}. $P_{\rm VWO} > 0.95$ requires $\langle b^2 \rangle / \langle \sigma^2 \rangle < 0.005$.

(iii) Lensing PCs are completely independent of the techniques used to generate power spectra, so it is reasonable to assume that the biases on individual PCs are uncorrelated, in which case K is proportional to the identity. $P_{\rm VWO} > 0.95$ requires $\langle b^2 \rangle / \langle \sigma^2 \rangle < 0.02$.

(iv) In the limit of linear growth, all k-modes are independent and the power spectrum grows according to a growth factor. It is then reasonable to assume that the biases are uncorrelated for different k bins, but maximally correlated at different redshifts for fixed k. K is computed in $k - z$ and then rotated to PCA space. We refer to this as the "fiducial shape" and use this shape in all that follows. $P_{\rm VWO} > 0.95$ requires $\langle b^2 \rangle / \langle \sigma^2 \rangle < 0.01$.

Even in the last two most realistic cases, the degeneracy between modeling errors in different regions changes the requirements on the modeling bias by a factor of 2.

D. Simulation requirements

We now put constraints on the magnitude of the bias, $|b|$, for a Euclid-like survey, assuming K has the fiducial shape (modeling assumptions are listed in Appendix E). In practice a knowledge matrix should be estimated from a real simulation, but in this section we examine a few test cases which are motivated above. For applications to real simulations, our code, RequiSim, is made public.

Figure 11 shows the $P_{\rm VWO}$ as a function of $|b|$, assumed to be constant everywhere, for different ℓ-mode cuts. Ensuring that $P_{\rm VWO} > 0.95$ requires $|b| < 1\%(3\%)$ for $\ell_{\rm cut} = 3000(500)$. In practice, there are other contributions to the error budget beyond modeling the power spectrum, so the requirement on $P_{\rm VWO}$ should be made more stringent.

Finally in Fig. 12, we consider the more realistic case in which the bias varies across some transition wave number: $k = k_b$. We assume 10% bias above k_b and 1% below. If the power spectrum can be modeled to within 1% down to scales $k_b = 7h$ Mpc^{-1}, then $P_{\rm VWO} > 0.95$ for all ℓ-cuts.

---

[5]These are computed by rescaling each element of C by a constant factor so that $\mathrm{Tr}(C) = d$, where d is the number of dimensions in the PCA basis so that $\langle \sigma^2 \rangle = 1$. Then K is normalized so that $\mathrm{Tr}(K) = db^2/4$ in the basis where K is diagonal. In 1D, this reduces to the normalization convention used in M12.
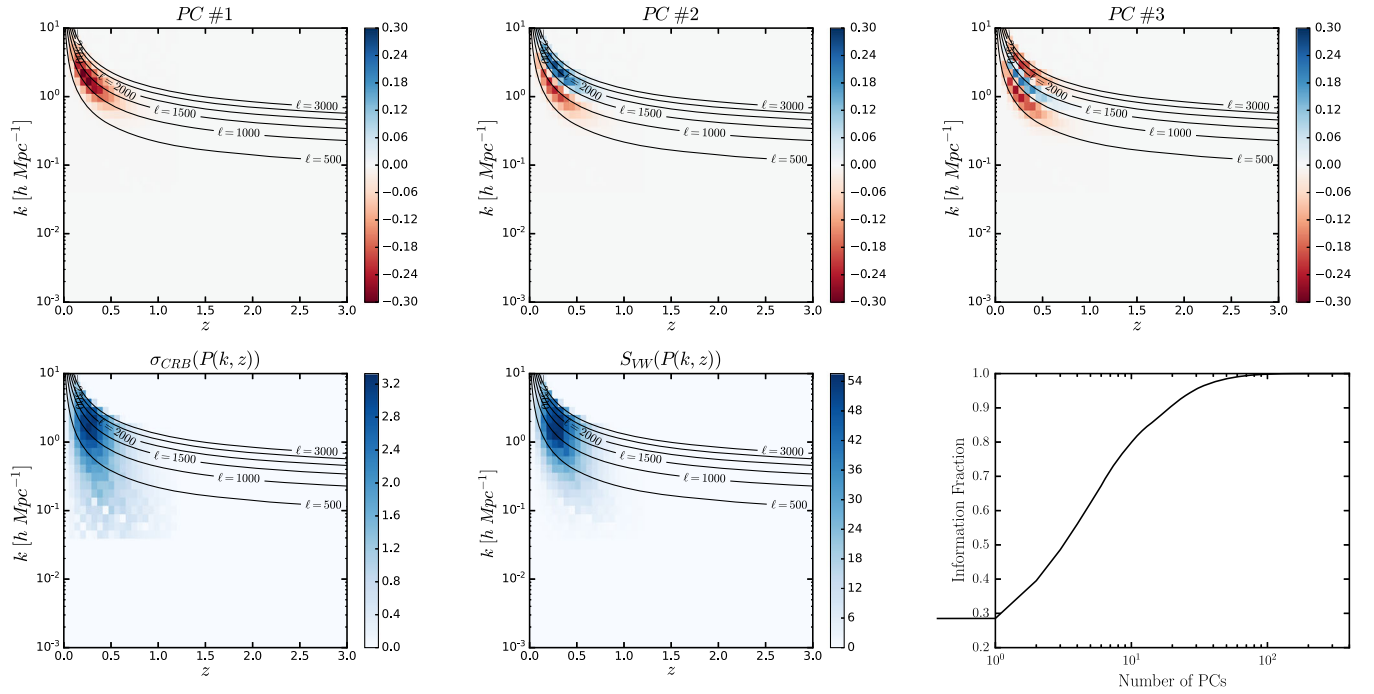
FIG. 9.   The first three high-resolution supertomographic power spectrum PCs, the Cramer-Rao bound, the variance weighted sum, and the information fraction. The PCs follow Limber lines [see discussion below Eq. (10)] plotted for different ℓ-modes in black.

However, as k_b is decreased corresponding to lowering the k-mode above which one feels confident in the accuracy of a simulation, $P_{\rm VWO}$ starts to fall off. This is particularly noticeable if we use a large $\ell_{\rm max}$. In fact if there is a 10% bias above $k_b = 1.1h$ Mpc^{-1} and only 1% below, then we would need to take an $\ell_{\rm cut} < 1000$ to get an unbiased result.

Meanwhile, since the lensing signal is extremely insensitive to large k-modes, models of the power spectrum at large k can be extremely inaccurate without biasing results. For example, using an $\ell_{\rm cut} = 3000$ assuming a bias of 50% above $k_b = 7h$ Mpc^{-1} and 0.5% below still yields $P_{\rm VWO} > 0.95$.

With these constraints, accurately modeling baryonic effects that affect large scales below $k = 1h$ Mpc^{-1}, like AGN feedback [41], must take priority. Accurately modeling small scale effects, like radiative gas cooling [41],
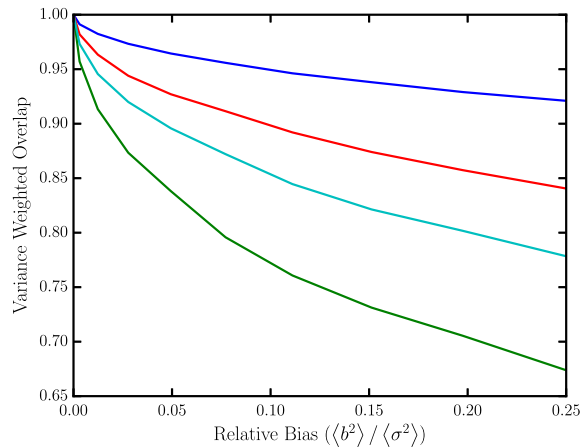

FIG. 10. $P_{\rm VWO}$ for different covariance K. Correlation between bias in different k-z regions has a large effect on simulation requirements. Factor of 3 difference on $\langle b^2 \rangle / \langle \sigma^2 \rangle$ requirements for $P_{\rm VWO} > 0.95$. For extreme K, K is correlated with lensing PCs covariance C (blue) and anticorrelated with C (green). For realistic K, K is diagonal and constrains all PCs equally well (red). Cyan shows the fiducial K-shape (see Fig. 11).
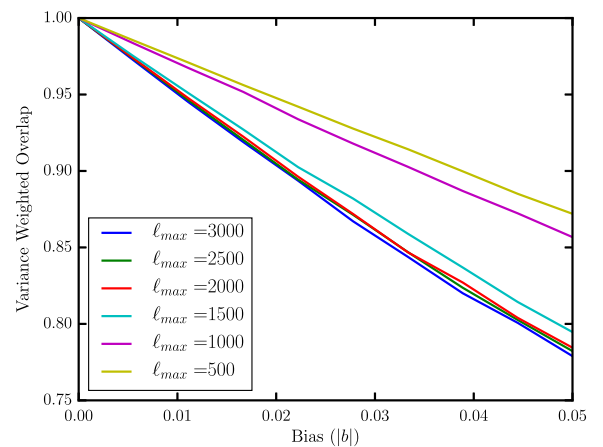

FIG. 11. $P_{\rm VWO}$ as a function of $|b|$ for different ℓ-mode cuts. Bias is naively assumed to be maximally correlated in k and constant everywhere. To ensure $P_{\rm VWO} > 0.95$ for $\ell_{\rm max} = 3000(500)$ requires the bias to be less than 1%(3%).
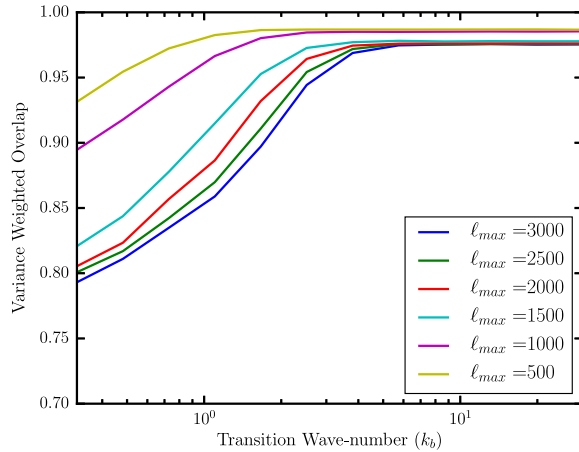
FIG. 12. P_{VWO} assuming a 10% error above the transition wave number $k = k_b$ and 1% below, for different ℓ-cuts. In this scenario, to avoid taking an ℓ-cut, the power spectrum should be known up to $k \sim 4h$ Mpc^{-1} for unbiased results. Increasing our knowledge of the power spectrum at scales smaller than $k \sim 7h$ Mpc^{-1} does not significantly change the bias on the lensing signal.

which only becomes important beyond $k = 5h$ Mpc^{-1}, is less important for cosmic shear studies.

V. CONCLUSION

We have compared the sensitivity of tomographic lensing with different binning strategies and 3D cosmic shear to the power spectrum and lensing kernel, independently from any assumptions about the underlying cosmological model. We draw the following conclusions:

  (i) While 3D cosmic shear captures the full 3D information content, it is slow and difficult to compute.

 (ii) While tomography using an equal number of galaxies per bin captures the majority of the information with very few bins and is computational straightforward, this technique loses information at high redshifts. This is where the cross-correlations with CMB lensing will be strongest.

(iii) Tomography using equally spaced redshift bins does not capture as much information as tomography using an equal number of galaxies per bin with very few bins; however, it incurs little computational cost to increase the number of bins to capture all the 3D information. We estimate that 99%(99.9%) of the information from both the lensing kernel and the power spectrum will be captured with 50(90) bins.

Nevertheless, the large covariance produced with a large number of bins may pose a challenge for a full likelihood analysis. Using our generalized formalism it should be possible to construct a weight function that retains the speed advantage of tomography while

capturing the majority of information in just a few modes. This is left to a future work. Meanwhile, the majority of the structure growth information is extracted from the power spectrum in the region $k \in [1h$ Mpc$^{-1}, 7h$ Mpc$^{-1}]$ and $z \in [0.1, 1.0]$.

Sensitivity to such high-k modes poses a problem. Nonlinear and baryonic physics which are hard to model become important at these scales. We have investigated how bias from incorrectly modeling these scales propagates to bias in the signal.

Generalizing the analysis in [19] to higher dimensions, we have shown that requirements depend not only on the magnitude of the bias but where they occur in $k - z$ space and on the correlation between biases at different scales and redshifts.

Assuming that the biases are maximally correlated in redshift along fixed k, and uncorrelated for different k-modes, as they would be in the limit of linear growth, and that the bias is the same everywhere, we find the power spectrum must be modeled to at least 1% accuracy for $k \leq 7h$ Mpc^{-1}. There are also other sources of bias, so the power spectrum should be modeled more accurately than this so that it does not subsume all of its error budget allocation. This will depend on the extent to which other systematics are brought under control.

Unless correlations between errors in different regions of the power spectrum are extremely anticorrelated with the lensing PCs, then current simulations are not at the stage where they can be used without taking an ℓ-cut. The stated accuracy of HALOFIT [42] is 5% for $k \leq 1h$ Mpc^{-1} and 10% for $k \leq 10h$ Mpc^{-1}. Meanwhile, COSMIC EMU [43] reports 4% accuracy for $k \in [0.1h$ Mpc$^{-1}, 10h$ Mpc$^{-1}]$ and HMCode [40] reports 5% accuracy for $k \in [0.1h$ Mpc$^{-1}, 10h$ Mpc$^{-1}]$.

Our assumptions are likely oversimplistic, so we provide the public code, RequiSim, to compute the bias on the lensing signal from inaccurate power spectrum models produced by any simulation.

Although we have not computed the bias on the lensing signal for existing power spectrum codes, we can provide a qualitative road map forward for simulators. Since the power spectrum is largely insensitive to scales $k > 7h$ Mpc^{-1}, simulators should focus on accurately modeling scales of $k < 7h$ Mpc^{-1} first.

ACKNOWLEDGMENTS

APPENDIX A: MOTIVATION FOR THE BESSEL WEIGHT

There are in general three reasons to write a signal in spherical-Bessel space:

(1) Spherical-Bessel functions follow an orthogonality relation [44].

(2) Spherical-Bessel functions and spherical harmonics are eigenfunctions of the Laplace operator in spherical coordinates. This ostensibly comes from the Laplacian used to relate the lensing potential and hence the shear in terms of the density field through the power spectrum.

(3) If the observed signal traces the cosmological density field directly, e.g., as is the case in galaxy clustering, then at a redshift z, the projected power spectrum of the signal in a spherical-Bessel representation, $C_\ell(k; z)$, is related to the matter power spectrum, $P(k; z)$, by an equality [44]:

$$C_\ell(k; z) = P(k; z). \qquad (A1)$$

This implies that cutting out high k-modes from the observed projected spectrum should cleanly remove sensitivity to small, poorly understood, high-k mode scales in the matter power spectrum [45].

The first consideration is a valid reason to use a spherical-Bessel basis set for weak lensing as it ensures that the shot noise [see Eq. (7)] is uncorrelated and does not become too large. However, *any* orthogonal set of functions will suffice in Eq. (1), and the Bessel functions are needlessly expensive to compute compared to other choices.

The second consideration is not relevant here because only the Newtonian potential must be expressed in this basis to relate it to the cosmological density field, through a Poisson equation; this is where the Bessel functions in Eq. (5) originate [44].

Meanwhile, the lensing power spectrum, C_l, traces all matter power *below* a certain redshift, weighted by a lensing kernel; i.e., the power spectrum is always enclosed within an integral over the line of sight. Hence, there is no reason that taking an η-cut should preferentially remove sensitivity to small scales in the matter power spectrum. However, confusion can arise by labeling both the lensing spectrum and the power spectrum wave number with k, and equating the two. We further discuss directly removing sensitivity to small scales in Appendix H.

APPENDIX B: FISHER MATRIX FORMALISM

Before conducting an experiment, the Fisher matrix can be used to estimate constraints and predict degeneracies for a set of parameters, $\{\theta_i\}$ [46]. We use it to estimate constraints on different regions of the matter power spectrum and lensing kernel and to predict correlations between them.

Provided the likelihood is Gaussian, the Fisher matrix for cosmic shear is

$$F_{ij} = \sum_\ell \frac{2\ell + 1}{2} \text{Tr}[C_\ell^{-1} C_{\ell,i} C_\ell^{-1} C_{\ell,j}], \qquad (B1)$$

where $C_{\ell,i}$ is the derivative of the lensing spectrum with respect to parameter θ_i.

Normally the sensitivity to the original parameters is given by the covariance matrix, C, found by inverting the Fisher matrix. However, in our analysis we produce many large, ill-conditioned, and nearly singular Fisher matrices, so inversion introduces too many numerical artifacts. Instead we use the Cramer-Rao bound. Defining σ_i as the conditional error on θ_i, the Cramer-Rao bound is

$$\frac{1}{\sigma_i} \leq \sqrt{F_{ii}}, \qquad (B2)$$

assuming all other parameters are known [46]. This measure of uncertainty does not account for the correlations between parameters. An alternative measure, which does, is defined in Appendix C.

APPENDIX C: PRINCIPAL COMPONENT ANALYSIS

For any experiment we can choose a set of parameters and estimate their covariance. A PCA finds a smaller set of independent parameters that capture the majority of the information. Informally these can be thought of as the parameters that are actually being measured, or that the data are in fact sensitive to.

For a set of N parameters, $\{\theta_i\}$, the covariance matrix, $C \equiv F^{-1}$, encodes parameter degeneracies. Since it is symmetric, it can be rotated into an eigenbasis where there are no degeneracies:

$$C = P^T D P, \qquad (C1)$$

where D is a diagonal matrix of N nonzero eigenvalues, P is a matrix formed of real eigenvectors of C, and P^T is the transpose of P.

These new parameters, $\{\zeta_i\}$, are related to the old parameters by

$$\zeta_i = \sum_j (v_i)_j \theta_j, \qquad (C2)$$

where v_i is the ith eigenvector of C and the ith row of the matrix P. When we apply this formalism to the power spectrum and lensing kernel, the set $\{\zeta_i\}$ will correspond to the amplitudes of a set of step functions, $\{f_i\}$, which can be formed from $\{v_i\}$. We refer to these functions as

"components" and the value of the jth element of v_i denotes the height of cell j in component i. For power spectrum and lensing kernel PCs, the cells will define regions in $k - z$ space and z space, respectively (see the next Appendix for more details).

The components with the smallest eigenvalues are the most tightly constrained and hence they contain the most information. Arranging the components according to ascending order in the corresponding eigenvalues, λ_i, and also the diagonal components of D, we define the fractional information content of the first m eigenvalues as

$$I(m) = \frac{\sum_{i<m}\lambda_i}{I_{\text{tot}}}, \tag{C3}$$

where $I_{\text{tot}} \equiv \sum_i \lambda_i$. The first few components, which contain the majority of the information, are called the principal components (PCs). Meanwhile, the total information content, I_{tot}, is also occasionally referred to as the "figure of merit" (FoM). This measure will be useful to compare the total constraining power of 3D and tomographic cosmic shear.

To avoid inverting the Fisher matrix which may be ill conditioned, we will calculate the PCs directly. Inverting C in Eq. (C1), we find

$$F = P^T D^{-1} P, \tag{C4}$$

and the PCs correspond to the largest eigenvalues along the diagonal of D^{-1}.

The total sensitivity to different regions can also be found without inverting F by taking a weighted sum of the components, f_i, in terms of the information content of each. This is given by the variance weighted sum defined as

$$S_{vw} = \sum_i \lambda_i |f_i|. \tag{C5}$$

The absolute value is taken because individual components have positive and negative values. Since S_{vw} is computed in PCA space, it naturally takes into account correlations between components, unlike the Cramer-Rao bound.

APPENDIX D: POWER SPECTRUM AND LENSING KERNEL PRINCIPAL COMPONENTS

To determine how 3D and tomographic cosmic shear are sensitive to the matter power spectrum and the lensing kernel, we perform a PCA, closely following the procedure in [35]. Our analysis assesses the sensitivity to the growth of structure and background evolution independently from any assumption of the underlying cosmological model.

To find the power spectrum PCs, we divide the power spectrum, $P(k, z)$, into logarithmically and linearly spaced grid cells in k and z, respectively. Inside each grid cell, i, we compute the fractional amplitude change in the power:

$$P_i(k,z,\mathcal{A}) \equiv \begin{cases} (1+\mathcal{A})P(k,z) & \text{if } (k,z) \text{ in cell } i \\ P(k,z) & \text{otherwise,} \end{cases} \tag{D1}$$

where $\mathcal{A}$ is a fixed small amplitude change. Defining each of these transformations as a parameter, θ_i, we compute a two-sided derivative:

$$\frac{\partial P}{\partial \theta_i} = \frac{P_i(k, z, \mathcal{A}) - P_i(k, z, -\mathcal{A})}{2\mathcal{A}}. \tag{D2}$$

From these we compute the Fisher matrix, and hence the PCs. In [35], the authors computed the PCs on a low-resolution matter power spectrum grid, before smoothly interpolating to higher resolution. Therefore it is unclear how much of the structure seen in their PCs is due to interpolation errors. Our method avoids this issue since the matter power spectrum is perturbed only after interpolation. Interpolation errors can thus be seen as a small change to the fiducial power spectrum.

We find the lensing kernel PCs by dividing the comoving distance into equally spaced redshift slices and making the perturbation:

$$r_i(z, \mathcal{A}) \equiv \begin{cases} (1 + \mathcal{A})r(z) & \text{if } z \text{ in slice } i \\ r(z) & \text{otherwise.} \end{cases} \tag{D3}$$

Hence the perturbed lensing kernel is

$$(F_k)_i(\mathcal{A}) = \frac{r_i(z, \mathcal{A}) - r_i'(z, \mathcal{A})}{r_i(z, \mathcal{A})r_i'(z, \mathcal{A})}. \tag{D4}$$

Again treating each perturbation as a separate parameter, θ_i, we define the two-sided derivative as

$$\frac{\partial F_k}{\partial \theta_i} = \frac{(F_k)_i(\mathcal{A}) - (F_k)_i(-\mathcal{A})}{2\mathcal{A}}, \tag{D5}$$

and compute the Fisher matrix as before.

In theory there are correlations between power spectrum PCs and lensing kernel PCs inside a much larger Fisher matrix. However, perturbations to the power spectrum have a very different effect on the lensing signal to perturbations to the lensing kernel so we assume that the two types of PCs are uncorrelated.

APPENDIX E: MODELING CHOICES

We assume a Gaussian distribution for the photometric redshift error given by

$$p(z|z_p) \equiv \frac{1}{2\pi\sigma_z(z_p)} e^{-\frac{(z - c_{\text{cal}} z_p + z_{\text{bias}})^2}{2\sigma_{z_p}}}, \tag{E1}$$

with $c_{\text{cal}} = 1$, $z_{\text{bias}} = 0$, and $\sigma_{z_p} = A(1 + z_p)$ with $A = 0.05$ [47] and

$$n(z_p) \propto \frac{a_1}{c_1} e^{-\frac{(z-0.7)^2}{b_1^2}} + e^{-\frac{(z-1.2)^2}{d_1^2}}, \qquad \text{(E2)}$$

with $(a_1/c_1, b_1, d_1) = (1.5/0.2, 0.32, 0.46)$ [48]. We assume a 15 000 degree survey with 30 galaxies per arcmin^2.

We use a fiducial lambda cold dark matter cosmology with $(\Omega_m, \Omega_k, \Omega_b, h_0, n_s, A_s, \tau) = (0.315, 0.0, 0.04, 0.67, 0.96, 2.1 \times 10^9, 0.08)$ throughout. The power spectrum is generated using CAMB [49] and the nonlinear part is generated using HALOFIT [42], produced as part of the Cosmosis [14] pipeline, each run with the default setting given in the *demo1* tutorial in Cosmosis.

APPENDIX F: CONVERGENCE CHECKS

To reduce computation time, we compute the Fisher matrix sampling sparsely in ℓ, taking

$$F_{ij} = \sum_{\ell \in L} \Delta \ell \frac{2\ell + 1}{2} \text{Tr}[C_\ell^{-1} C_{\ell,i} C_\ell^{-1} C_{l,j}], \qquad \text{(F1)}$$

where ℓ is sampled at $[2, 12, 25, 50, 75]$ below 100, then at intervals of 50 to $\ell = 2000$, and finally intervals of 100 to $\ell = 3000$. This cuts the computation time by nearly an order of magnitude. For 10-bin tomography this leads to a $<7\%$ error inside the variance weighted sum in the largest power spectrum PCA bin and $<0.1\%$ average error across all bins, compared to the Fisher where every ℓ-mode is sampled.

The lensing power spectrum is computed on an $N \times N$ grid logarithmically spaced in k and linearly space in z. For
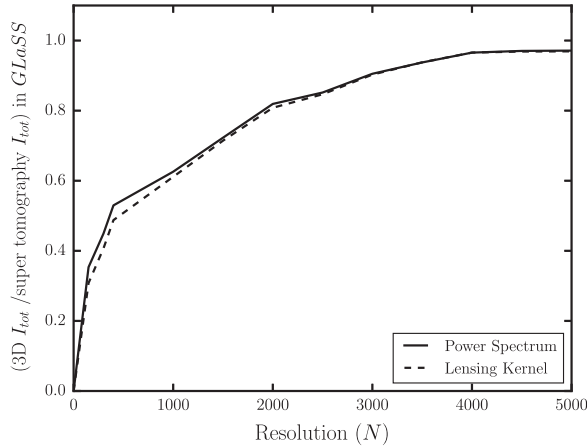

FIG. 13. Convergence of 3D cosmic shear total information content, I_{tot} [see Eq. (C3)], relative to supertomography as a function of the resolution of the computation grid. Both the lensing kernel and power spectrum I_{tot} converge to within 3% at a resolution of $N = 5000$ and 20% at $N = 2000$. I_{tot}'s were computed on a coarse PCA grid sampled very sparsely in ℓ. Due to memory constraints, $N = 2000$ throughout the rest of the paper.

our analysis, tomography and supertomography are fully converged for $N = 400$, which we have used throughout. As 3D cosmic shear converges slowly, we test the convergence of 3D cosmic shear on a coarse PCA grid: two by two and two by one for the power spectrum and lensing kernel, respectively. The convergence of the information content as a function of N, relative to supertomography, is shown in Fig. 13.

APPENDIX G: COMPARISON WITH OTHER WORK

Tomographic and 3D cosmic shear were recently compared in [16] (hereafter SM18) which reports a decrease in the error on some modified gravity parameters of 20–30% for 3D cosmic shear compared to 6 bin tomography with an equal number of galaxies per bin. Meanwhile, we only find a 15% and 2% increase in the total information, I_{tot}, for the lensing kernel and power spectrum, respectively, when going from this regime to supertomography.

The slightly smaller gains in our analysis are expected due to two differences in modeling assumptions. SM18 used $\ell_{max} = 1000$, while we used $\ell_{max} = 3000$. The higher ℓ-cut used in our analysis means we are relatively more sensitive to lower redshifts below $z = 0.5$ (see Fig. 8). However, tomography with an equal number of galaxies per bin primarily loses information at higher redshifts, beyond $z = 1$ (see Fig. 5). SM18 used a linear power spectrum, while we used a nonlinear which relatively boosts our sensitivity to high-k modes in the power spectrum. Again these modes are primarily probed at low z (see Fig. 9), where tomography with an equal number of galaxies per bin loses information.

APPENDIX H: DIRECTLY REMOVING SENSITIVITY TO SMALL SCALE POWER

As small scales, modeling errors introduce bias. Ideally it would be possible to remove sensitivity to these modes above some k_{cut}. We split the matter power spectrum into two parts: $P^{k>k_{cut}}$ and $P^{k<k_{cut}}$, where the former contains only power above the cut and the later power below. The resulting lensing spectra, $C_\ell^{k>k_{cut}}$ and $C_\ell^{k<k_{cut}}$, were calculated. The spectra have power at nearly identical modes, making it difficult to reduce the sensitivity to small scales without also losing sensitivity to the signal.

APPENDIX I: REQUISIM

RequiSim is available for download from https://github.com/astro-informatics/RequiSim. Using precomputed PCs and a user-provided knowledge matrix, RequiSim computes the P_{VWO} for a Euclid-like survey. PCs for other surveys can be computed on request.

[1] A. Refregier, R. Massey, J. Rhodes, R. Ellis, J. Albert, D. Bacon, G. Bernstein, T. McKay, and S. Perlmutter, Astron. J. **127,** 3102 (2004).

[2] C. Heymans, E. Grocutt, A. Heavens, M. Kilbinger, T. D. Kitching, F. Simpson, J. Benjamin, T. Erben, H. Hildebrandt, H. Hoekstra et al., Mon. Not. R. Astron. Soc. **432,** 2433 (2013).

[3] H. Hildebrandt, M. Viola, C. Heymans, S. Joudaki, K. Kuijken, C. Blake, T. Erben, B. Joachimi, D. Klaes, L. Miller et al., Mon. Not. R. Astron. Soc. **465,** 1454 (2017).

[4] M. Troxel, N. MacCrann, J. Zuntz, T. Eifler, E. Krause, S. Dodelson, D. Gruen, J. Blazek, O. Friedrich, S. Samuroff et al., arXiv:1708.01538.

[5] A. Refregier, A. Amara, T. Kitching, A. Rassat, R. Scaramella, J. Weller et al., arXiv:1001.0061.

[6] A. Albrecht, G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox et al., arXiv:astro-ph/0609591.

[7] R. J. Laureijs, L. Duvet, I. E. Sanz, P. Gondoin, D. H. Lumb, T. Oosterbroek, and G. S. Criado, Proc. SPIE Int. Soc. Opt. Eng. **7731,** 77311H (2010).

[8] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. Gaudi, T. Greene, O. Guyon et al., arXiv:1503.03757.

[9] Hu. Zhan and J. Anthony Tyson, Rep. Prog. Phys. **81,** 066901 (2018).

[10] W. Hu, Astrophys. J. Lett. **522,** L21 (1999).

[11] A. Heavens, Mon. Not. R. Astron. Soc. **343,** 1327 (2003).

[12] S. Bridle and L. King, New J. Phys. **9,** 444 (2007).

[13] P. L. Taylor, T. D. Kitching, J. D. McEwen, and T. Tram, Phys. Rev. D **98,** 023522 (2018).

[14] J. Zuntz, M. Paterno, E. Jennings, D. Rudd, A. Manzotti, S. Dodelson, S. Bridle, S. Sehrish, and J. Kowalkowski, Astronomy and Computing **12,** 45 (2015).

[15] T. Kitching, A. Heavens, J. Alsing, T. Erben, C. Heymans, H. Hildebrandt, H. Hoekstra, A. Jaffe, A. Kiessling, Y. Mellier et al., Mon. Not. R. Astron. Soc. **442,** 1326 (2014).

[16] A. S. Mancini, R. Reischke, V. Pettorino, B. M. Scháefer, and M. Zumalacárregui, arXiv:1801.04251.

[17] G. Efstathiou and P. Lemos, Mon. Not. R. Astron. Soc. **476,** 151 (2018).

[18] M. Troxel and M. Ishak, Phys. Rep. **558,** 1 (2015).

[19] R. Massey, H. Hoekstra, T. Kitching, J. Rhodes, M. Cropper, J. Amiaux, D. Harvey, Y. Mellier, M. Meneghetti, L. Miller et al., Mon. Not. R. Astron. Soc. **429,** 661 (2013).

[20] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, Astrophys. J. **761,** 152 (2012).

[21] D. H. Rudd, A. R. Zentner, and A. V. Kravtsov, Astrophys. J. **672,** 19 (2008).

[22] A. Heavens, T. D. Kitching, and A. Taylor, Mon. Not. R. Astron. Soc. **373,** 105 (2006).

[23] Y. Ayaita, B. M. Schäfer, and M. Weber, Mon. Not. R. Astron. Soc. **422,** 3056 (2012).

[24] B. M. Schäfer and L. Heisenberg, Mon. Not. R. Astron. Soc. **423,** 3445 (2012).

[25] T. Kitching and A. Heavens, Phys. Rev. D **95,** 063522 (2017).

[26] M. L. Brown, A. N. Taylor, D. J. Bacon, M. E. Gray, S. Dye, K. Meisenheimer, and C. Wolf, Mon. Not. R. Astron. Soc. **341,** 100 (2003).

[27] T. D. Kitching, J. Alsing, A. F. Heavens, R. Jimenez, J. D. McEwen, and L. Verde, Mon. Not. R. Astron. Soc. **469,** 2737 (2017).

[28] M. LoVerde and N. Afshordi, Phys. Rev. D **78,** 123506 (2008).

[29] B. Joachimi and S. Bridle, Astron. Astrophys. **523,** A1 (2010).

[30] H. Hildebrandt, M. Viola, C. Heymans, S. Joudaki, K. Kuijken, C. Blake, T. Erben, B. Joachimi, D. Klaes, L. Miller et al., Mon. Not. R. Astron. Soc. **465,** 1454 (2017).

[31] R. Massey, J. Rhodes, A. Leauthaud, P. Capak, R. Ellis, A. Koekemoer, A. Réfrégier, N. Scoville, J. E. Taylor, J. Albert et al., Astrophys. J. Suppl. Ser. **172,** 239 (2007).

[32] H. S. Xavier, F. B. Abdalla, and B. Joachimi, Mon. Not. R. Astron. Soc. **459,** 3693 (2016).

[33] A. S. Mancini, P. Taylor, R. Reischke, T. Kitching, V. Pettorino, B. Schäfer, B. Zieser, and P. M. Merkel, arXiv:1807.11461.

[34] B. Joachimi, P. Schneider, and T. Eifler, Astron. Astrophys. **477,** 43 (2008).

[35] S. Foreman, M. R. Becker, and R. H. Wechsler, Mon. Not. R. Astron. Soc. **463,** 3326 (2016).

[36] M. Chevallier and D. Polarski, Int. J. Mod. Phys. D **10,** 213 (2001).

[37] E. V. Linder, Phys. Rev. Lett. **90,** 091301 (2003).

[38] J. Friedman, T. Hastie, and R. Tibshirani, The Elements of Statistical Learning, Springer Series in Statistics, Vol. 1 (Springer, New York, 2001).

[39] P. M. Merkel and B. M. Schäfer, Mon. Not. R. Astron. Soc. **469,** 2760 (2017).

[40] A. Mead, J. Peacock, C. Heymans, S. Joudaki, and A. Heavens, Mon. Not. R. Astron. Soc. **454,** 1958 (2015).

[41] M. P. van Daalen, J. Schaye, C. Booth, and C. Dalla Vecchia, Mon. Not. R. Astron. Soc. **415,** 3649 (2011).

[42] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, Astrophys. J. **761,** 152 (2012).

[43] K. Heitmann, E. Lawrence, J. Kwan, S. Habib, and D. Higdon, Astrophys. J. **780,** 111 (2013).

[44] P. Castro, A. Heavens, and T. Kitching, Phys. Rev. D **72,** 023516 (2005).

[45] F. Lanusse, A. Rassat, and J.-L. Starck, Astron. Astrophys. **578,** A10 (2015).

[46] M. Tegmark, A. N. Taylor, and A. F. Heavens, Astrophys. J. **480,** 22 (1997).

[47] O. Ilbert, S. Arnouts, H. McCracken, M. Bolzonella, E. Bertin, O. Le Fevre, Y. Mellier, G. Zamorani, R. Pello, A. Iovino et al., Astron. Astrophys. **457,** 841 (2006).

[48] L. Van Waerbeke, J. Benjamin, T. Erben, C. Heymans, H. Hildebrandt, H. Hoekstra, T. D. Kitching, Y. Mellier, L. Miller, J. Coupon et al., Mon. Not. R. Astron. Soc. **433,** 3373 (2013).

[49] A. Lewis, A. Challinor, and A. Lasenby, The Astrophysical Journal **538,** 473 (2011).