

Cosmic Shear: Inference from Forward Models

Peter L. Taylor,^{1,*} Thomas D. Kitching,¹ Justing Alsing,^{2,3,4} Benjamin D. Wandelt,^{2,5,6,7} Stephen M. Feeney,² and Jason D. McEwen¹

¹*Mullard Space Science Laboratory, University College London, Holmbury St. Mary, Dorking, Surrey RH5 6NT, UK*

²*Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave, New York City, NY 10010, USA*

³*Oskar Klein Centre for Cosmoparticle Physics, Stockholm University, AlbaNova, Stockholm SE-106 91, Sweden*

⁴*Imperial Centre for Inference and Cosmology, Department of Physics, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK*

⁵*Sorbonne Université, CNRS, UMR 7095, Institut d'Astrophysique de Paris, 98 bis boulevard Arago, 75014 Paris, France*

⁶*Sorbonne Université, Institut Lagrange de Paris (ILP), 98 bis boulevard Arago, 75014 Paris, France*

⁷*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08540, USA*

Density-estimation likelihood-free inference (DELFI) has recently been proposed as an efficient method for simulation-based cosmological parameter inference. Compared to the standard likelihood-based Markov Chain Monte Carlo (MCMC) approach, DELFI has several advantages: it is highly parallelizable, there is no need to assume a possibly incorrect functional form for the likelihood and complicated effects (e.g the mask and detector systematics) are easier to handle with forward models. In light of this, we present two DELFI pipelines to perform weak lensing parameter inference with lognormal realizations of the tomographic shear field – using the C_ℓ summary statistic. The first pipeline accounts for the non-Gaussianities of the shear field, intrinsic alignments and photometric-redshift error. We validate that it is accurate enough for Stage III experiments and estimate that $\mathcal{O}(1000)$ simulations are needed to perform inference on Stage IV data. By comparing the second DELFI pipeline, which makes no assumption about the functional form of the likelihood, with the standard MCMC approach, which assumes a Gaussian likelihood, we test the impact of the Gaussian likelihood approximation in the MCMC analysis. We find it has a negligible impact on Stage IV parameter constraints. Our pipeline is a step towards seamlessly propagating all data-processing, instrumental, theoretical and astrophysical systematics through to the final parameter constraints.

I. INTRODUCTION

Weak lensing by large scale structure offers some of the tightest constraints on cosmological parameters. Over the next decade data from Stage IV experiments including Euclid¹ [1], WFIRST² [2] and LSST³ [3] will begin taking data. Extracting as much information from these ground-breaking data sets, in an unbiased way, presents a formidable challenge.

The majority of cosmic shear studies to date focus on extracting information from two-point statistics and in particular the correlation function, $\xi(\theta)$, in configuration space and the lensing power spectrum, C_ℓ , in spherical harmonic space [4–8]. While the non-Gaussian information in the shear field is accessed with higher-order statistics [9, 10], peak counts [11, 12] or machine learning [13], the impact of systematics on the two-point functions have been extensively studied [14]. For this reason we will focus on these statistics and leave the higher-order information to a future study. In particular we focus on the C_ℓ statistic because computing correlation functions from catalogues with billions galaxies

– even using an efficient code such as TREECORR [15] – is extremely computationally demanding.

Apart from [16, 17], existing studies of the shear two-point statistics [4–8] use a Gaussian likelihood analysis to infer the cosmological parameters. This approach has drawbacks. For example, with the improved statistical precision of next generation data, we will need to propagate complicated ‘theoretical systematics’ (e.g. reduced shear [18]) and detector effects [14] into the final cosmological constraints. It is difficult to derive the expected impact of these effects as is required for a likelihood analysis. It is much easier to produce forward model realizations.

It has also recently been claimed that because the true lensing likelihood is left-skewed, not Gaussian, parameter constraints from correlation functions are biased low in the $\sigma_8 - \Omega_m$ plane [19, 20]. The same argument given in these papers applies to the C_ℓ statistic. More will be said about this in Section III.

To overcome these issues, a new method called density-estimation likelihood-free inference (DELFI) [21–27] offers a way forward. DELFI is a ‘likelihood-free’ method similar to approximate Bayesian computation (ABC) [28], but much more computationally efficient. Using summary statistics (the C_ℓ statistic, in this case) generated from full forward models of the data at different points in cosmological parameter space, DELFI is used to estimate the posterior distribution.

* peterllewlyntaylor@gmail.com

¹ <http://euclid-ec.org>

² <https://www.nasa.gov/wfirst>

³ <https://www.lsst.org>

Performing inference on realizations of the data may seem computationally challenging, but using efficient data compression [25, 29] most applications require only $\mathcal{O}(10^3)$ simulations [27]. This is less than the number of simulations already required to produce a valid estimate of the inverse covariance matrix in a Stage IV likelihood analysis. DELFI is also highly parallelizable.

In this paper we develop two cosmic shear forward model pipelines for DELFI, using the publicly available `pydelfi`⁴ implementation, summarized in Figure 1. Pipeline I takes full advantage of the benefits of forward modelling and is intended for application to real datasets, while Pipeline II is intended only for comparison with the standard likelihood analysis. We also consider 3 different analyses summarized in Table I. DA1 is a DELFI analysis using shear Pipeline I. Meanwhile we compare the DELFI analysis, DA2, to the likelihood analysis, LA, to test the impact of the Gaussian likelihood approximation. It is useful for the reader to refer back to Table I and Figure 1 throughout the text.

The structure of this paper is as follows. The formalism of cosmic shear and cosmological parameter inference is reviewed in Sections II-III. While DELFI has already been applied to cosmic shear in a simple Gaussian field setting [27], in Section IV we go beyond this and present a more realistic forward model (Pipeline I) which includes the impact of intrinsic alignments and non-Gaussianities of the shear field. We also estimate the number of simulations required for a Stage IV experiment and check to confirm that we recover the input cosmology from a DA1 analysis on mock data. Next we discuss the feasibility of the DA1 analysis for Stage III data in Section V. In Section VI we test the impact of the Gaussian likelihood approximation by comparing the DA2 analysis to the LA analysis. In Section VII we discuss future prospects for DELFI in cosmic shear studies, before concluding in Section IX.

II. COSMIC SHEAR FORMALISM AND THE LOGNORMAL FIELD APPROXIMATION

A. The Lensing Spectrum

Assuming the Limber [30, 31], spatially-flat Universe [32], flat sky [31] and equal-time correlator approximations [33], the lensing spectrum, $C_{\ell, \text{GG}}^{ij}$, is given by [4]:

$$C_{\ell, \text{GG}}^{ij} = \int_0^{r_{\text{H}}} dr \frac{q_i(r)q_j(r)}{r^2} P\left(\frac{\ell}{r}, r\right), \quad (1)$$

where $P(k, r)$ is the matter power spectrum and the lensing efficiency kernel, q_i is defined as:

$$q_i(r) = \frac{3H_0^2\Omega_{\text{m}}}{2c^2} \frac{r}{a(r)} \int_r^{r_{\text{H}}} dr' n_i(r') \frac{r' - r}{r'}, \quad (2)$$

Pipeline I Pipeline II




FIG. 1. A schematic of the two forward model pipelines used in this work given model parameters p . In Pipeline I we develop a forward model of cosmic shear data for inference with DELFI which takes advantage of the forward model approach. There is no need to deconvolve the mask or pixel window function, for example. In Pipeline II we use a Gaussian field, do not use a mask, subtract off the shot-noise or deconvolve the pixel window function. These choices allow us to make a direct comparison between DELFI and a Gaussian likelihood analysis to test the Gaussian likelihood assumption.

⁴ <https://github.com/justinalsing/pydelfi/commits/master>

	DA1	DA2	LA
Inference	DELFI	DELFI	Gaussian likelihood
Pipeline	Pipeline I	Pipeline II	NA
Number of galaxies	1.56×10^9	1.56×10^9	1.56×10^9
Number of tomographic bins	6	2	2
Number of ℓ -bins	15 with $\ell \in [10, 1000]$	15 with $\ell \in [10, 1000]$	15 with $\ell \in [10, 1000]$
Field type	Lognormal	Gaussian	NA
Deconvolve Pixel Window	No	Yes	NA
Mask	Yes	No	NA
Subtract shot-noise	No	Yes	NA

TABLE I. The three analyses in this paper. In DA1 we use DELFI to infer the cosmological parameters. Since we perform inference with forward models, there is no need to deconvolve the pixel window function, deconvolve the mask or subtract off the shot noise. This analysis is applied to mock Stage IV data in Section IV. Meanwhile by comparing the DELFI analysis, DA2, with the likelihood analysis, LA, we test the impact of the Gaussian likelihood approximation. In DA2 our modeling choices are governed by the constraint that we must match the Gaussian likelihood analysis as closely as possible. Some of the map-level choices are not applicable to the likelihood analysis LA.

and we generate the $N_{\text{tomographic}}$ tomographic bins, $n_i(r')$, by dividing the radial distribution function:

$$n(z_p) \propto \frac{a_1}{c_1} e^{-\frac{(z-0.7)^2}{b_1^2}} + e^{-\frac{(z-1.2)^2}{d_1^2}}, \quad (3)$$

with $(a_1/c_1, b_1, d_1) = (1.5/0.2, 0.32, 0.46)$ [34] into bins with an equal number of galaxies per bin. To account for photometric redshift error, each bin is smoothed by the Gaussian kernel:

$$p(z|z_p) \equiv \frac{1}{2\pi\sigma_z(z_p)} e^{-\frac{(z-c_{\text{cal}}z_p+z_{\text{bias}})^2}{2\sigma_{z_p}^2}}, \quad (4)$$

with $c_{\text{cal}} = 1$, $z_{\text{bias}} = 0$ and $\sigma_{z_p} = A(1+z_p)$ with $A = 0.05$ [35].

B. Intrinsic Alignments

The tidal alignment of galaxies around massive halos adds two additional terms to the lensing spectrum. An ‘II term’ accounts for the intrinsic tidal alignment of galaxies around massive dark matter halos, while a ‘GI term’ accounts for the anti-correlation between tidally aligned galaxies at low redshifts and weakly lensed galaxies at high redshift.

We model this effect using the non-linear alignment (NLA) model [4, 36]. We also allow the intrinsic amplitude, $A(z)$, to vary as a function of redshift so that $A(z) = [(1+z_0)/(1+z)]^\eta$ [37], where z_0 is the mean redshift of the survey. This is $z_0 = 0.76$ for the $n(z)$ given in equation (3). This model was used in the joint KiDS-450+2dFLenS [38] analysis and was one of the models considered in the Dark Energy Survey Year 1 cosmic shear analysis (hereafter DESY1) [5].

In this case the II spectrum, $C_{\ell, \text{II}}^{ij}$, is given by:

$$C_{\ell, \text{II}}^{ij} = \int_0^{r_{\text{H}}} dr \frac{n_i(r)n_j(r)}{r^2} P_{\text{II}}\left(\frac{\ell}{r}, r\right), \quad (5)$$

where the II matter power spectrum is:

$$P_{\text{II}}(k, z) = F^2(z)P(k, z) \quad (6)$$

and

$$F(z) = -A_I C_1 \rho_{\text{crit}} \frac{\Omega_{\text{m}}}{D(z)} \left(\frac{1+z(r)}{1+z_0} \right)^\eta, \quad (7)$$

where ρ_{crit} is the critical density of the Universe, $D(z)$ is the growth factor and $C_1 = 5 \times 10^{-14} h^{-2} M_\odot^{-1} \text{Mpc}^3$.

The GI power spectrum is:

$$C_{\ell, \text{GI}}^{ij} = \int_0^{r_{\text{H}}} dr \frac{q_i(r)n_j(r) + n_i(r)q_j(r)}{r^2} P_{\text{GI}}\left(\frac{\ell}{r}, r\right), \quad (8)$$

and the GI matter power spectrum is:

$$P_{\text{GI}}(k, z) = F(z)P(k, z). \quad (9)$$

Altogether the theoretical lensing spectrum, $C_\ell^{T, ij}$, is given by the sum of the three contributions:

$$C_\ell^T = C_{\ell, \text{GG}}^{ij} + C_{\ell, \text{GI}}^{ij} + C_{\ell, \text{II}}^{ij} \quad (10)$$

Henceforth we will routinely drop the tomographic bin labels for convenience, as we have done here, on the left hand side.

All lensing spectra are generated inside the `Cosmosis` framework [39]. The linear power spectrum and expansion history are computed with `CAMB` [40] and the non-linear corrections are computed with `HALOFIT` [41].

C. The Lognormal Field Approximation

Generating lognormal convergence fields [42] is computationally inexpensive and captures the impact of nonlinear structure growth more accurately than Gaussian realizations. This approximation was recently used in DESY1 [5] to compute the covariance matrix from

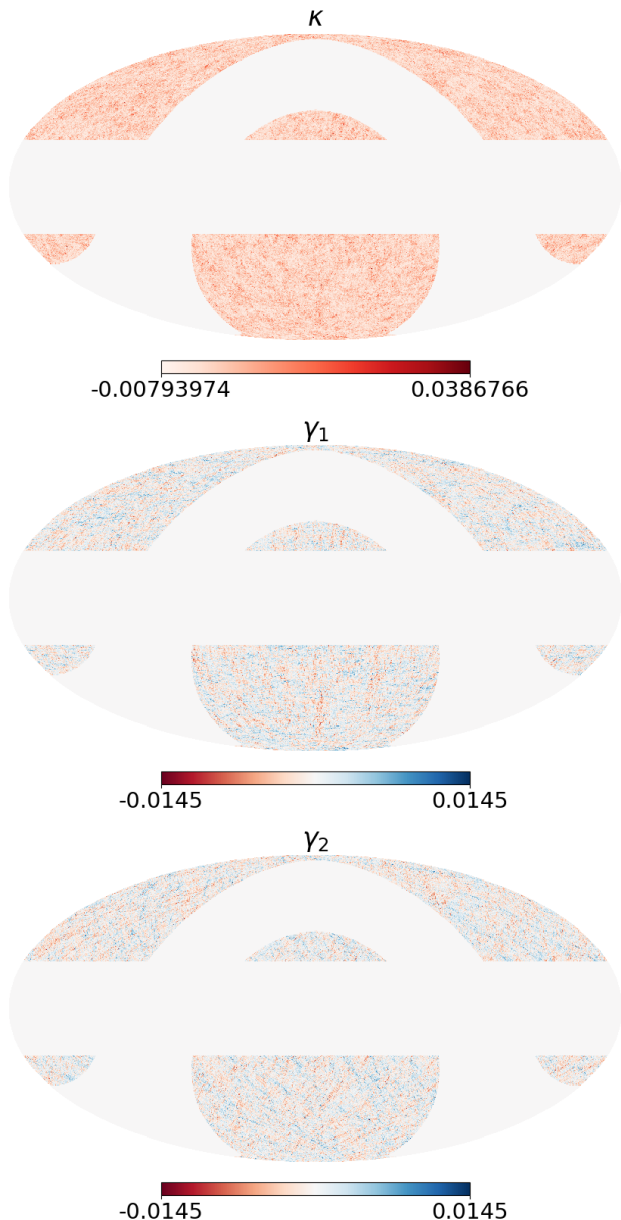


FIG. 2. A single masked data realization of the convergence field, κ , and the two observable shear components: γ_1 and γ_2 (including shape noise) for a typical Stage IV experiment. This is the lowest redshift bin of six, where the effect of non-Gaussianity is largest. The non-Gaussianity is clearly visible in the κ -map (the color scale runs between the minimum and maximum value of the κ -field to make the non-Gaussianity more visible), where the majority of pixels are very slightly negative with a small number of pixels taking very large (positive) κ -values. The mask cuts all pixels lying within 22.5 deg of the galactic and ecliptic planes.

noisy realisations of the data. No differences in parameter constraints were found when the covariance was computed using lognormal fields compared to the halo model approach [43].

In the lognormal field approximation the convergence, $\kappa^i(\theta)$, inside each tomographic bin, i , is generated by ex-

ponentiating and shifting a Gaussian realization, $g^i(\theta)$, according to:

$$\kappa^i(\theta) = \exp[g^i(\theta)] - \kappa_0^i, \quad (11)$$

where κ_0^i is a constant shift parameter.

We use `Flask` [44] to generate consistent lognormal realizations [42] of the convergence and shear fields – correlated between redshift slices. The procedure is discussed in detail in Section 5.2 of [44] (see also [45]).

`Flask` takes just two inputs:

- `Flask` takes the theoretical lensing spectrum, C_ℓ^T , defined in Sections II A-II B. Formally `Flask` uses the convergence spectrum to generate a convergence field, κ , from which it computes a consistent shear field, γ . In the flat sky approximation – which we assume throughout – the shear and convergence spectrum are the same, but care would be needed to correctly re-scale the input convergence spectrum by the appropriate ℓ -factor if the flat sky approximation was dropped [31, 46].
- `Flask` requires the shift parameter, κ_0^i for each tomographic bin i . We compute this by taking a weighted average of the shift parameter at each redshift:

$$\kappa_0^i = \int dz n_i(z) \kappa_0^i(z), \quad (12)$$

using the fitting formula:

$$\kappa_0^i(z) = 0.008z + 0.029z^2 - 0.0079z^3 + 0.0065z^4 \quad (13)$$

derived from simulations [42].

While the fitting formula will have some cosmological dependence, the shift parameter does not affect the power spectrum of the field – only impacting cosmological constraints through the covariance. Non-Gaussian corrections to the covariance already have a sub-dominant impact [47, 48], hence the dependence of these corrections on the cosmology is further sub-dominant. For this reason we ignore the cosmological dependence of the shift parameter.

A valid covariance matrix between data must be positive-definite, but this is not guaranteed for correlations between tomographic lognormal fields [44]. `Flask` overcomes this issue by perturbing the lognormal fields following the regularization procedure outlined in Section 3.1 of [44]. Provided that the regularization is applied to a small number of tomographic bins, it is found in [44] that $C_\ell^{\text{reg}}/C_\ell^{\text{ln}} \ll 1 \times 10^{-5}$, where C_ℓ^{reg} is the recovered regularized spectrum and C_ℓ^{ln} is the spectrum recovered from the unregularized map [44]. In Section V we verify that this will not impact Stage III parameter constraints.

In Figure 2 we plot a single lognormal realization generated with Pipeline I. We show the masked convergence and components of the shear field in the lowest redshift

bin. This is where the non-Gaussianities are most pronounced and clearly visible. In the convergence map the majority of the pixels take small negative values. However there are rare incidences of large positive convergence. This physically corresponds to collapsed high-density structures along the line-of-sight.

D. Band-limit Bias from the Lognormal Field

Unlike Gaussian fields, lognormal realizations are not band-limited in ℓ [44] (see Section 5.2.2 therein). In particular, Taylor expanding the lognormal convergence field, $\kappa^i(\theta)$, in terms of the Gaussian field, $g^i(\theta)$, yields quadratic and higher order terms in g^i . In harmonic space this mixes different ℓ -modes. When a band-limit is imposed, this biases the lensing spectrum recovered from the map.

III. COSMOLOGICAL PARAMETER INFERENCE

A. Gaussian likelihood Analysis

In the standard two-point cosmic shear likelihood analysis, we assume a Gaussian likelihood:

$$\ln \mathcal{L}(p) = -\frac{1}{2} \sum_{a,b} [D_a - T_a(p)] C_{ab}^{-1} [D_b - T_b(p)], \quad (14)$$

where D_a and $T_a(p)$ are the data and theory vectors respectively composed of the C_ℓ estimated from data and the theoretical expectation of C_ℓ given cosmological parameters p .

Meanwhile C_{ab}^{-1} is the inverse of the covariance matrix. Since we generate the covariance matrix from noisy simulations of the data we make the Anderson-Hartlap [49, 50] correction, to avoid bias from inverting the covariance matrix, for the remainder of the paper.

B. The Potential Insufficiency of the Gaussian likelihood Approximation in Cosmic Shear

To see why the Gaussian likelihood assumptions can lead to bias we summarize the argument given in [20]. Inside a single bin the unmasked lensing spectrum is:

$$C_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |\gamma_{\ell m}|^2. \quad (15)$$

Since the harmonic coefficients, $\gamma_{\ell m}$, are computed as a summation over a large number of pixels, they are Gaussian distributed by the central limit theorem. Squaring a Gaussian random variable gives a gamma distribution – which is left-skewed.

Taking a Gaussian rather than a gamma distribution for the likelihood could bias parameter constraints.

Since $S_8^2 = \sigma_8^2(\Omega_m/0.3)$ and Ω_m enter into the shear spectrum amplitude, we would expect these parameters to be ones which are most affected – and biased low. Only in the limit of large ℓ – as the C_ℓ itself becomes the sum over a large number of m -modes – does the central limit theorem kick in and the likelihood become Gaussian.

C. Density-estimation Likelihood-free Compression

Since density-estimation likelihood-free inference methods are most effective in low dimensions [25], we compress the C_ℓ summary statistic. As suggested in [29], the lensing spectra are compressed into a new vector, t , according to:

$$C_\ell \rightarrow t = \nabla_p \mathcal{L}_*, \quad (16)$$

where p is the set of cosmological parameters that we are inferring, \mathcal{L}_* is a proposal Gaussian likelihood centred at a fiducial set of parameters which we take to be $(\Omega_m, h_0, \Omega_b, n_s, S_8, A, \eta) = (0.3, 0.72, 0.96, 0.79, 1., 2.8)$ throughout, where $S_8 = \sigma_8(\Omega_m/0.3)^{0.5}$, A and η are the intrinsic alignment parameters defined in Section IIB and the other parameters take their standard cosmological definitions. As the assumption of a Gaussian likelihood here is only for compression purposes, it does not bias the final parameter constraints and the Fisher information is preserved provided the true likelihood is Gaussian [29]. If the true likelihood is not exactly Gaussian, some information will be lost. This is investigated in Section VI. For more advanced compression techniques using neural networks, see [51].

D. Density-estimation Likelihood-free Inference

We use `pydelfi` [27] to learn the conditional density $\mathbb{P}(t|p)$ (this software comes with many different run-mode options, but we restrict our attention to the methods used in this work). The likelihood is then given by $\mathbb{P}(t = t_{\text{data}}|p)$, where t_{data} is the mock data generated from either Pipeline I or II. Multiplying by the prior, which we take to be flat in all parameters, yields the posterior.

Using the default setting in `pydelfi`, we train five neural density estimators (NDE) (four mixture density networks (MDN) and one masked autoregressive flow (MAF), see [27] for more details) with the default network architectures described in Section 4 of [27], parameterized in terms of a set of neural network weights, w . Training multiple networks allows DELFI to avoid over-fitting and increases robustness.

We use sequential learning to learn the weights, w , updating our knowledge of the conditional density distribution $\mathbb{P}(t = t_{\text{data}}|p)$. Specifically we divide the inference task into 20 training steps with 100 simulations

per step. Given a large enough computer all the simulations in each training step could be run in parallel, so that the total time of the simulations would not exceed the time it took to perform 20 simulations.

As an initial guess for the conditional distribution, we take the multivariate Gaussian:

$$\mathbb{P}(t|p) = \mathcal{N}(t|p, F^{-1}), \quad (17)$$

where F^{-1} is the inverse of the Fisher matrix of the cosmological parameters, $F = -\langle \nabla_p \nabla_p^{Tr} \mathcal{L}_* \rangle$. At each step thereafter, we train each neural density estimator on a set of parameter realization pairs $\{p_i, t_i\}$ drawing samples from the conditional density of the previous step to ensure that the highest density regions are the most finely sampled. Meanwhile ten percent of the samples are retained as a validation set to avoid over-fitting.

At each step each NDE learns the weights, w , by minimizing the negative loss function:

$$-\ln U(w) = - \sum_{i=1}^{N_{\text{samples}}} \ln \mathbb{P}(t_i|p, w), \quad (18)$$

which is an estimate of the Kullback-Leibler divergence between the density estimator $\mathbb{P}(t = t_{\text{data}}|p)$ and the true distribution [27] (the minus sign is pulled out of $U(w)$ on the left-hand side following the convention in [27]). The final estimate for the conditional distribution is given as a weighted average over the estimates from the five networks:

$$\mathbb{P}(t|p; w) = \sum_{i \in \text{networks}} \beta_i p_i(t|p; w), \quad (19)$$

where the weights are determined by the relative likelihood of each NDE [27].

IV. THE FULL FORWARD MODEL

In this section we use Pipeline I to generate mock Stage IV data and then run analysis DA1 to recover the input cosmology. This allows us to test our pipeline and estimate the number of simulations needed for a Stage IV experiment. We describe the model choices and results below.

A. The Mask

We use a typical Stage IV survey mask shown in Figure 2. All pixels lying within 22.5 deg of either the galactic or ecliptic planes are masked. This leaves 14,490 deg² of unmasked pixels which, as a fraction of the full sky, is $f_{\text{sky}} = 0.35$.

B. Shot-Noise Model

The noise, γ_p , for each pixel, p , is drawn from a Gaussian distribution [16]:

$$\gamma_p \sim \mathcal{N}\left(0, \frac{\sigma_\epsilon}{\sqrt{\widehat{N}_P}}\right) \quad (20)$$

where \widehat{N}_P is the number of galaxies in each pixel, the orientation is angle is drawn from a uniform distribution, we take the intrinsic shape dispersion as $\sigma_\epsilon = 0.3$ [52] and use 30 galaxies per arcmin² throughout. This is a good approximation since in all our simulations there are a large number of galaxies in each pixel, so the central limit theorem applies.

C. Forward Modelling the Mask

One advantage of performing inference with full forward models of the data is that we do not need to deconvolve the mask. This is both computationally simpler and avoids the risk of bias from inaccurate deconvolution which is present in the standard likelihood analysis.

Given two masked shear fields $a(\theta)$ and $b(\theta)$, a naïve estimate of the lensing spectrum is the pixel pseudo- C_ℓ spectrum:

$$\widetilde{C}_\ell^{\text{pix}, EE} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \langle a_{\ell m}^E b_{\ell m}^E \rangle, \quad (21)$$

where the tilde is used to denote the fact that we have not corrected for the mask and the ‘pix’ superscript reminds us that we have not accounted for the pixel window function. Analogous expressions are easily found for the EB and BB spectra.

In an unmasked field, lensing by large scale structure will only induce power in the EE spectra, but to retain information leaked into the EB and BB spectra due to the presence of a mask, in Pipeline I, we use:

$$\widetilde{C}_\ell^{\text{pix}} = \widetilde{C}_\ell^{\text{pix}, EE} + \widetilde{C}_\ell^{\text{pix}, EB} + \widetilde{C}_\ell^{\text{pix}, BE} + \widetilde{C}_\ell^{\text{pix}, BB}, \quad (22)$$

as the estimator. This is computed using `HEALPy` [53, 54].

In a future pipeline it may still be desirable to use the pseudo- C_ℓ formalism to avoid mixing between E and B -modes, allowing us to immediately remove B -modes induced by unknown systematics. As long as the data and theory are treated in the same way the pseudo- C_ℓ formalism will not introduce bias, as it could in the standard likelihood analysis.

D. Mimicking a Stage IV Experiment

To estimate the number of simulations needed for a Stage IV experiment and ensure that pipeline recovers the input cosmology we produce mock data with




FIG. 3. 68% and 95% credible region parameter constraints found with DELFI analysis DA1 after 1000 simulations, for a mock Stage IV experiment. We confirm that we recover the input cosmology within statistical errors. We plot the convergence in Figure 4. In a realistic situation there may be a larger number of nuisance parameters. This would not dramatically slow convergence because we could ‘nuisance harden’ the data compression step, to only learn the posterior for the parameters of interest [55].

Pipeline I. We use 6 tomographic bins sampling 15 logarithmically spaced ℓ -bins in the range $\ell \in [10, 1000]$. We then run `pydelfi` to estimate the posterior distribution of the cosmological parameters for this data. The final parameter constraints for a lambda cold dark matter (LCDM) cosmology with two nuisance intrinsic alignment parameters are shown in Figure 3. This confirms that we recover the input parameters within errors.

In Figure 4 we plot the negative loss function defined in equation (18) for the training and validation sets. Both have converged within $\mathcal{O}(1000)$ simulations. This is similar to the number found in the simple Gaussian field pipeline presented in [27], suggesting that the inclusion of higher order effects including intrinsic alignments and non-Gaussian field corrections does not significantly increase the required number of simulations.




FIG. 4. The negative loss function defined in equation (18) for the training and validation sets as a function of the number of simulations. This suggests that $\mathcal{O}(1000)$ simulations are needed for a Stage IV experiment. This is similar to the number found in [27], which only considered a simple Gaussian field forward model with no intrinsic alignments, implying that the convergence rate is fairly insensitive to the precise details of the model.

When working with real data, we may require a large number of nuisance parameters. Nevertheless, we do not expect this to dramatically increase the number of simulations needed, since we can always tune the data compression to maximize the information retention of the parameters of interest, following the procedure in [55].

Each simulation takes approximately 33 minutes on a single thread of a 1.8 GHz Intel Xeon (E5-2650Lv3) Processor. Thus if run on 100 threads in parallel, the total simulation time of the DELFI inference step takes only 10 hours. Many of the individual modules in the pipeline are multithreaded (e.g `Flask`), so running on even more threads would further reduce the total run-time.

V. PROSPECTS FOR STAGE III DATA

In this section we discuss the viability of applying analysis DA1 to existing Stage III data. For the remainder of this section we assume a circular mask of 4951 deg^2 , similar to the final coverage of the Dark Energy Survey [56] with 10 galaxies per arcmin^2 and use Pipeline I throughout this section – except where modifications are explicitly stated.

A. Validating the Lognormal Simulations

Lognormal fields were used to generate the covariance matrix in the recent Dark Energy Survey Year 1 analysis [5]. The authors found no difference in parameter




FIG. 5. The colored lines show the absolute value of the difference between the average recovered cross and intra-bin spectra from 100 lognormal and 100 Gaussian realizations (4 tomographic bins, $N_{\text{side}} = 512$ and $\ell \in [10, 1535]$). The difference is due to the band-limit bias in the lognormal field discussed in Section II D. With these model choices the band-limit bias is safely below 1% for nearly all data points.

constraints between this analysis and one which used a halo model to generate the covariance matrix – but to verify that our pipeline is ready for Stage III data, we must also ensure that we recover an unbiased C_ℓ from the maps.

Given an accurate input C_ℓ , the only bias in Pipeline I comes from regularizing the map (see Section II C). We would not expect the band-limit bias of the lognormal field to be problematic since imposing a band-limit would affect the data in the same way. However this assumes that the true field is exactly lognormal. Nevertheless we check to ensure that the combined effect of regularization and imposing a band-limit is small. We quantify this statement by finding the difference between the average recovered pixelated C_ℓ 's from 100 Gaussian simulations (where no band-limit bias or regularization bias is present) and 100 lognormal simulations. Each 4-tomographic bin simulation takes approximately 15 minutes on a single thread and the difference in the recovered spectra is shown in Figure 5. The bias is safely below 1% in all but three data points. This confirms that once minor updates have been made (see next subsection), the pipeline will be ready for use on today's data.

B. Model Improvements

Only a small number of adjustments must be made to DA1 to apply this analysis to existing data. These are:

- We must accurately account for baryonic physics.

This can be handled using a halo model code [57], potentially in combination with the k-cut cosmic shear approach [58, 59], optimally cutting scales which can not be accurately modeled.

- We must introduce several nuisance parameters. As well as allowing for free multiplicative and additive shear biases, photo-z bias parameters will need to be allowed to vary, as in the Dark Energy Year 1 analysis [5]. This will increase the number of nuisance parameters. To avoid excessive computational costs we must ‘nuisance harden’ [55] the data compression step.

VI. TESTING THE GAUSSIAN LIKELIHOOD APPROXIMATION

In this section we compare DELFI and the standard Gaussian likelihood analysis by running the DA2 analysis and the LA analysis, on the same mock Stage IV data. We use Pipeline II to generate the mock data, produce the covariance matrix and generate the forward models in DA2. Since DELFI does not assume any particular likelihood, differences in the resulting parameter constraints are only due to the Gaussian likelihood assumption in LA. Because we can not just forward model everything in LA, care must be taken to ensure that the band-limit bias, deconvolving the mask, deconvolving the pixel window function and subtracting the shot-noise does not lead to additional bias between the two analyses. Controlling for these effects is described in the first subsection.

A. Modeling Choices In Pipeline II

To avoid the band-limit bias we use a Gaussian field, rather than the lognormal field.

We do not apply a mask in DA2 as we have found that using the pseudo- C_ℓ method [60–62] (with the public code `NaMaster` [62]) can bias parameter constraints, with our choice of `HEALpix`⁵ N_{side} and ℓ_{max} by up to 1σ . Instead we adjust the galaxy number density so that total number of galaxies and hence the signal-to-noise remains unchanged.

In LA we decide to take the C_ℓ , with no shot-noise term in the intra-bin case, as the data vector. Thus we must subtract off the expected value of the noise in DA2. This is computed by running 500 noise-only simulations, as in the analysis of [8].

We must also account for the fact that the shear spectra are computed on pixelized maps – that is, we must deconvolve the pixel window function, w_ℓ , which is defined in [63] and computed using `HEALpix`. This assumes that the scale of the signal is large relative to the




FIG. 6. The absolute value of the bias, $|b|$, due to imperfect pixel-window deconvolution and noise subtraction relative to the statistical error, σ , from 500 Pipeline II simulations. This confirms that the comparison between DELFI and the likelihood analysis presented in Section VIB will be unaffected.

pixel scale and that all pixels are the same shape. The window-corrected spectrum, C_ℓ , is given in terms of the spectrum computed from a pixelized map, C_ℓ^{pix} , by:

$$C_\ell = w_\ell^{-2} C_\ell^{\text{pix}}. \quad (23)$$

By running 500 Gaussian field simulations we have confirmed that the combined bias from deconvolving the pixel window function and subtracting the shot-noise is small, so that we can fair comparison between the DA2 analysis and the LA analysis. This is shown in Figure 6. The absolute value of the bias, $|b|$, is small relative to the statistical error, σ , with $|b|/\sigma < 0.1$ for all data points.

B. Impact of the Gaussian Likelihood Approximation

To test the impact of the Gaussian likelihood approximation we first generate 1000 mock data realizations using Pipeline II. We take 15 logarithmically spaced ℓ -bins in the range [10, 1000] and restrict our attention to the $S_8 - \Omega_m$ plane. To cut computation cost, we use only two tomographic bins. The parameters S_8 and Ω_m primarily impact the amplitude of the shear spectrum, so we do not expect to lose too much information with this choice [64, 65].

For three random data realizations, we run a DELFI and a Gaussian likelihood analysis. The resulting posteriors are shown in Figure 7. Each subplot corresponds to one of the three realizations.

In all three cases the DELFI and Gaussian likelihood contours are very similar. This suggests the Gaus-

⁵ <https://sourceforge.net/projects/healpix/>




FIG. 7. The 68% and 95% credible region parameter constraints for three random data realizations found using a MCMC Gaussian likelihood analysis and DELFI, which makes no assumption about the functional form of the likelihood. The mock data input cosmology is labeled by black dotted lines. Only in the first realization, does the input cosmology lie outside the 68% credible region – but statistically this is to be expected for a small number of realizations. The contours found using the two different analyses are very similar for all three data realizations suggesting that the Gaussian likelihood approximation has negligible impact, and the compression in equation (16) is lossless. This former statement is confirmed in Figure 8.




FIG. 8. The 68% and 95% credible region of the MLE distribution, assuming a Gaussian likelihood. The value of the input cosmology is indicated by the black dotted lines, and lies at the center of the contours. This implies that the Gaussian likelihood approximation does not lead to any measurable bias in our setup.

sian likelihood assumption does not bias parameter constraints in the $S_8 - \Omega_m$ plane and the compression defined in equation (16) is lossless.

To confirm and quantify this statement, we sample the maximum likelihood estimator (MLE) distribution assuming a Gaussian likelihood, using the 1000 data realizations generated earlier. For each realization, the MLE is found using the Nelder-Mead algorithm built into `scipy` and wrapped into `Cosmosis` using the default settings. The resulting MLE distribution is shown in Figure 8. The input cosmology lies almost exactly at the center of the 68% credible region which implies that there is no measurable bias from the Gaussian likelihood approximation.

We stress that these conclusions only hold for the C_ℓ analysis presented in this work. In particular, the ℓ -binning strategy matters. By binning ℓ -modes we are taking a sum over random variables, so by the central limit theorem broader bins correspond to more Gaussian data. The Gaussian likelihood approximation could be important for much narrower bins. Similarly the Gaussian likelihood approximation might be important in a correlation function analysis. Since correlation functions are integrals over C_ℓ 's weighted by a Bessel filter function [66], the non-Gaussianities at low- ℓ will be present in every data point.

VII. FUTURE PROSPECTS

We review the main known cosmic shear systematics which must eventually be included in the full forward model. To account for many of these effects we must first take the base model presented in this work to ‘catalog level’. This can be done by first generating a consistent density field – either with `Flask` or by taking the difference between two neighbouring tomographic bins – and then populating the density field with a realistic population of galaxies [67] assuming a biased tracer model (e.g [68]). Cosmic shear systematics break down into four broad categories: data-processing, theoretical, astrophysical and instrumental systematics.

On the data-processing side, accurately measuring the shape and photometric redshift of galaxies is the primary challenge. Both measurements are dependent on the galaxy-type [69], and this is in turn correlated with the density through the morphology-density relation [70]. Rather than using the best fit parameters for each galaxy, we can sample the posterior on each galaxy as in a Bayesian hierarchical model [16] to propagate the measurement uncertainty into the final parameter constraints, as suggested in [27]. We can also account for image blending [69, 71] more easily with forward models.

Two important theoretical systematics are the reduced shear correction [18, 66] and magnification bias [72, 73]. The former correction accounts for the fact that we measure the reduced shear $\gamma/(1-\kappa)$ with a weak lensing experiment. In a likelihood analysis, this can be computed using a perturbative expansion as in [18, 74]. This is slow and requires us to rely on potentially inaccurate fitting functions for the lensing bispectrum. Meanwhile the magnification bias accounts for the fact that galaxies of the same luminosity can fall above (below) the detectability limit in regions of high (low) lensing magnification. In both cases, these systematics can be easily handled with full forward models of consistent shear and convergence fields.

The two dominant instrumental systematics are the telescope’s point spread function (PSF) [14] and the effect of charge transfer inefficiency (CTI) in the charge-coupled devices (CCDs) [14, 75]. Efforts are underway to build pipelines which characterize these effects in upcoming experiments (e.g [76] and Paykari et al. (in prep)). Integrating these pipelines into ours would enable the propagation of instrumental errors through to the the final parameter constraints.

On the astrophysical side, the two dominant systematics are the impact of baryons on the density field [9] and the intrinsic alignment of galaxies [36, 77]. For Stage IV data, forward models will likely have to be based on high-resolution N-body lensing simulations [78, 79] to include the effects of baryons. Even with today’s highest resolution simulations the impact baryons is still uncertain [80], so it will likely be necessary to optimally cut [58] or marginalize out uncertain scales [80]. Meanwhile more sophisticated intrinsic

alignment models which account for different alignment behaviour by galaxy type [81] will need to be included.

Eventually higher-order statistics such as peak counts and the shear bispectrum can be added. Since DELFI automatically handles multiple summary statistics in a unified way, the constraints will be tighter than doing the two-point and higher-order statistic analyses separately. With a greater ability to handle systematics, DELFI may also open up the possibility of performing inference with weak lensing flux and size magnification [27, 82–86].

VIII. CONCLUSION

By comparing a Gaussian likelihood analysis to a fully likelihood-free DELFI analysis, we have found that the Gaussian likelihood approximation will have a negligible impact on Stage IV parameters constraints. Nevertheless we recommend the development of DELFI pipelines, because they offer the possibility of performing rapid parallel inference on full forward realizations of the shear data. In the future, this will allow us to seamlessly handle astrophysical and detector systematics – at a minimal computational cost.

We have taken the first steps towards this goal by developing a pipeline to rapidly generate realistic non-Gaussian shear data, including the impact of intrinsic alignments. Integrating this pipeline into `pydelfi`, we confirm that `pydelfi` is feasible, requiring only $\mathcal{O}(1000)$ simulations for Stage IV data, and we verify that the current pipeline is accurate enough for today’s data. As the pipeline is computationally inexpensive, in the future it will be useful for quickly determining which systematics are important.

We conclude that DELFI has a promising future in cosmic shear studies. Developing fast simulations that fully integrate all relevant astrophysical, detector and modelling effects is the primary hurdle. With so many clear advantages over the traditional likelihood analysis, developing these simulations should be a priority.

IX. ACKNOWLEDGEMENTS

The authors would like to thank Luke Pratley and David Alonso. We are indebted to the developers of all public code used in this work. PLT acknowledges the hospitality of the Flatiron Institute. This work was supported by a collaborative visit funded by the Cosmology and Astroparticle Student and Postdoc Exchange Network (CASPEN). PLT is supported by the UK Science and Technology Facilities Council. TDK is supported by a Royal Society University Research Fellowship. JA was partially supported by the research project grant ‘‘Fundamental Physics from Cosmological Surveys’’ funded by the Swedish Research Council (VR) under Dnr 2017-04212. BDW is supported by the Simons foundation.

The authors acknowledge the support of the Leverhulme Trust.

-
- [1] R. J. Laureijs, L. Duvet, I. E. Sanz, P. Gondoin, D. H. Lumb, T. Oosterbroek, and G. S. Criado, in *Proc. SPIE*, Vol. 7731 (2010) p. 77311H.
- [2] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. Gaudi, T. Greene, O. Guyon, *et al.*, arXiv preprint arXiv:1503.03757 (2015).
- [3] J. Anthony and L. Collaboration, in *Proc. of SPIE Vol.*, Vol. 4836, p. 11.
- [4] C. Heymans, E. Grocutt, A. Heavens, M. Kilbinger, T. D. Kitching, F. Simpson, J. Benjamin, T. Erben, H. Hildebrandt, H. Hoekstra, *et al.*, Monthly Notices of the Royal Astronomical Society **432**, 2433 (2013).
- [5] M. Troxel, N. MacCrann, J. Zuntz, T. Eifler, E. Krause, S. Dodelson, D. Gruen, J. Blazek, O. Friedrich, S. Samuroff, *et al.*, arXiv preprint arXiv:1708.01538 (2017).
- [6] T. Kitching, A. Heavens, J. Alsing, T. Erben, C. Heymans, H. Hildebrandt, H. Hoekstra, A. Jaffe, A. Kiessling, Y. Mellier, *et al.*, Monthly Notices of the Royal Astronomical Society **442**, 1326 (2014).
- [7] H. Hildebrandt, M. Viola, C. Heymans, S. Joudaki, K. Kuijken, C. Blake, T. Erben, B. Joachimi, D. Klaes, L. Miller, *et al.*, Monthly Notices of the Royal Astronomical Society (2017).
- [8] C. Hikage, M. Oguri, T. Hamana, S. More, R. Mandelbaum, M. Takada, F. Köhlinger, H. Miyatake, A. J. Nishizawa, H. Aihara, *et al.*, arXiv preprint arXiv:1809.09148 (2018).
- [9] E. Semboloni, T. Schrabback, L. van Waerbeke, S. Vafaei, J. Hartlap, and S. Hilbert, Monthly Notices of the Royal Astronomical Society **410**, 143 (2010).
- [10] L. Fu, M. Kilbinger, T. Erben, C. Heymans, H. Hildebrandt, H. Hoekstra, T. D. Kitching, Y. Mellier, L. Miller, E. Semboloni, *et al.*, Monthly Notices of the Royal Astronomical Society **441**, 2725 (2014).
- [11] A. Peel, C.-A. Lin, F. Lanusse, A. Leonard, J.-L. Starck, and M. Kilbinger, *Astronomy & Astrophysics* **599**, A79 (2017).
- [12] B. Jain and L. Van Waerbeke, *The Astrophysical Journal Letters* **530**, L1 (2000).
- [13] A. Gupta, J. M. Z. Matilla, D. Hsu, and Z. Haiman, *Physical Review D* **97**, 103515 (2018).
- [14] R. Massey, H. Hoekstra, T. Kitching, J. Rhodes, M. Cropper, J. Amiaux, D. Harvey, Y. Mellier, M. Meneghetti, L. Miller, *et al.*, Monthly Notices of the Royal Astronomical Society **429**, 661 (2012).
- [15] M. Jarvis, *Mon. Not. R. Astron. Soc.* **352**, 338 (2004).
- [16] J. Alsing, A. Heavens, A. H. Jaffe, A. Kiessling, B. Wandelt, and T. Hoffmann, Monthly Notices of the Royal Astronomical Society **455**, 4452 (2015).
- [17] J. Alsing, A. Heavens, and A. H. Jaffe, Monthly Notices of the Royal Astronomical Society **466**, 3272 (2016).
- [18] S. Dodelson, C. Shapiro, and M. White, *Physical Review D* **73**, 023009 (2006).
- [19] E. Sellentin and A. F. Heavens, Monthly Notices of the Royal Astronomical Society **473**, 2355 (2017).
- [20] E. Sellentin, C. Heymans, and J. Harnois-Déraps, Monthly Notices of the Royal Astronomical Society **477**, 4879 (2018).
- [21] F. V. Bonassi, L. You, and M. West, *Statistical applications in genetics and molecular biology* **10** (2011).
- [22] G. Papamakarios, D. C. Sterratt, and I. Murray, arXiv preprint arXiv:1805.07226 (2018).
- [23] Y. Fan, D. J. Nott, and S. A. Sisson, *Stat* **2**, 34 (2013).
- [24] G. Papamakarios, I. Murray, and T. Pavlakou, “Advances in neural information processing systems,” (2016).
- [25] J. Alsing, B. Wandelt, and S. Feeney, *Monthly Notices of the Royal Astronomical Society* **477**, 2874 (2018).
- [26] J.-M. Lueckmann, G. Bassetto, T. Karaletsos, and J. H. Macke, arXiv preprint arXiv:1805.09294 (2018).
- [27] S. F. B. W. Justin Alsing, Tom Charnock, arXiv preprint arXiv:1903.00007 (2019).
- [28] E. Ishida, S. Vitenti, M. Penna-Lima, J. Cisewski, R. de Souza, A. Trindade, E. Cameron, V. Busti, C. collaboration, *et al.*, *Astronomy and Computing* **13**, 1 (2015).
- [29] J. Alsing and B. Wandelt, *Monthly Notices of the Royal Astronomical Society: Letters* **476**, L60 (2018).
- [30] M. LoVerde and N. Afshordi, *Physical Review D* **78**, 123506 (2008).
- [31] T. D. Kitching, J. Alsing, A. F. Heavens, R. Jimenez, J. D. McEwen, and L. Verde, arXiv preprint arXiv:1611.04954 (2016).
- [32] P. L. Taylor, T. D. Kitching, J. D. McEwen, and T. Tram, *Phys. Rev. D* **98**, 023522 (2018).
- [33] T. D. Kitching and A. Heavens, *Physical Review D* **95**, 063522 (2017).
- [34] L. Van Waerbeke, J. Benjamin, T. Erben, C. Heymans, H. Hildebrandt, H. Hoekstra, T. D. Kitching, Y. Mellier, L. Miller, J. Coupon, *et al.*, Monthly Notices of the Royal Astronomical Society **433**, 3373 (2013).
- [35] O. Ilbert, S. Arnouts, H. McCracken, M. Bolzonella, E. Bertin, O. Le Fevre, Y. Mellier, G. Zamorani, R. Pello, A. Iovino, *et al.*, *Astronomy & Astrophysics* **457**, 841 (2006).
- [36] C. M. Hirata and U. Seljak, *Physical Review D* **70**, 063526 (2004).
- [37] N. MacCrann, J. Zuntz, S. Bridle, B. Jain, and M. R. Becker, Monthly Notices of the Royal Astronomical Society **451**, 2877 (2015).
- [38] S. Joudaki, C. Blake, A. Johnson, A. Amon, M. Asgari, A. Choi, T. Erben, K. Glazebrook, J. Harnois-Déraps, C. Heymans, *et al.*, Monthly Notices of the Royal Astronomical Society **474**, 4894 (2017).
- [39] 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).
- [40] A. Lewis and A. Challinor, *Astrophysics Source Code Library* (2011).
- [41] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, *The Astrophysical Journal* **761**, 152 (2012).
- [42] S. Hilbert, J. Hartlap, and P. Schneider, *Astronomy & Astrophysics* **536**, A85 (2011).
- [43] E. Krause and T. Eifler, Monthly Notices of the Royal Astronomical Society **470**, 2100 (2017).

- [44] H. S. Xavier, F. B. Abdalla, and B. Joachimi, *Monthly Notices of the Royal Astronomical Society* **459**, 3693 (2016).
- [45] A. S. Mancini, P. Taylor, R. Reischke, T. Kitching, V. Pettorino, B. Schäfer, B. Zieser, and P. M. Merkel, *Physical Review D* **98**, 103507 (2018).
- [46] P. Castro, A. Heavens, and T. Kitching, *Physical Review D* **72**, 023516 (2005).
- [47] M. Sato and T. Nishimichi, *Physical Review D* **87**, 123538 (2013).
- [48] T. Eifler, E. Krause, P. Schneider, and K. Honscheid, *Monthly Notices of the Royal Astronomical Society* **440**, 1379 (2014).
- [49] J. Hartlap, P. Simon, and P. Schneider, *Astronomy & Astrophysics* **464**, 399 (2007).
- [50] T. W. Anderson, T. W. Anderson, T. W. Anderson, T. W. Anderson, and E.-U. Mathématicien, *An introduction to multivariate statistical analysis*, Vol. 2 (Wiley New York, 1958).
- [51] T. Charnock, G. Lavaux, and B. D. Wandelt, *Physical Review D* **97**, 083004 (2018).
- [52] M. Brown, A. Taylor, N. Hambly, and S. Dye, *Monthly Notices of the Royal Astronomical Society* **333**, 501 (2002).
- [53] K. M. Gorski, B. D. Wandelt, F. K. Hansen, E. Hivon, and A. J. Banday, arXiv preprint astro-ph/9905275 (1999).
- [54] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *Astrophys. J.* **622**, 759 (2005), arXiv:astro-ph/0409513.
- [55] J. Alsing and B. Wandelt, arXiv preprint arXiv:1903.01473 (2019).
- [56] M. Troxel and M. Ishak, *Physics Reports* **558**, 1 (2015).
- [57] A. Mead, J. Peacock, C. Heymans, S. Joudaki, and A. Heavens, *Monthly Notices of the Royal Astronomical Society* **454**, 1958 (2015).
- [58] P. L. Taylor, F. Bernardeau, and T. D. Kitching, *Physical Review D* **98**, 083514 (2018).
- [59] F. Bernardeau, T. Nishimichi, and A. Taruya, *Monthly Notices of the Royal Astronomical Society* **445**, 1526 (2014).
- [60] B. D. Wandelt, E. Hivon, and K. M. Gorski, *Physical Review D* **64**, 083003 (2001).
- [61] M. Brown, P. Castro, and A. Taylor, arXiv preprint astro-ph/0410394 (2008).
- [62] D. Alonso, J. Sanchez, and A. Slosar, *Monthly Notices of the Royal Astronomical Society* (2018).
- [63] D. Jeong, J. Chluba, L. Dai, M. Kamionkowski, and X. Wang, *Physical Review D* **89**, 023003 (2014).
- [64] P. L. Taylor, T. D. Kitching, and J. D. McEwen, *Phys. Rev. D* **99**, 043532 (2019).
- [65] A. Spurio Mancini, R. Reischke, V. Pettorino, B. Schäfer, and M. Zumalacárregui, *Monthly Notices of the Royal Astronomical Society* **480**, 3725 (2018).
- [66] M. Bartelmann and P. Schneider, *Physics Reports* **340**, 291 (2001).
- [67] L. Miller, T. Kitching, C. Heymans, A. Heavens, and L. Van Waerbeke, *Monthly Notices of the Royal Astronomical Society* **382**, 315 (2007).
- [68] J. Elvin-Poole, M. Crocce, A. Ross, T. Giannantonio, E. Rozo, E. Rykoff, S. Avila, N. Banik, J. Blazek, S. Bridle, *et al.*, *Physical Review D* **98**, 042006 (2018).
- [69] A. Kannawadi, H. Hoekstra, L. Miller, M. Viola, I. F. Conti, R. Herbonnet, T. Erben, C. Heymans, H. Hildebrandt, K. Kuijken, *et al.*, arXiv preprint arXiv:1812.03983 (2018).
- [70] R. Houghton, *Monthly Notices of the Royal Astronomical Society* **451**, 3427 (2015).
- [71] S. Samuroff, S. Bridle, J. Zuntz, M. Troxel, D. Gruen, R. Rollins, G. Bernstein, T. Eifler, E. Huff, T. Kacprzak, *et al.*, *Monthly Notices of the Royal Astronomical Society* **475**, 4524 (2017).
- [72] T. Hamana, *Monthly Notices of the Royal Astronomical Society* **326**, 326 (2001).
- [73] J. Liu, Z. Haiman, L. Hui, J. M. Kratochvil, and M. May, *Physical Review D* **89**, 023515 (2014).
- [74] C. Shapiro, *The Astrophysical Journal* **696**, 775 (2009).
- [75] J. Rhodes, A. Leauthaud, C. Stoughton, R. Massey, K. Dawson, W. Kolbe, and N. Roe, *Publications of the Astronomical Society of the Pacific* **122**, 439 (2010).
- [76] R. D. Vavrek, R. J. Laureijs, J. L. Alvarez, J. Amiaux, Y. Mellier, R. Azzollini, G. Buenadicha, G. S. Criado, M. Cropper, C. Dabin, *et al.*, in *Modeling, Systems Engineering, and Project Management for Astronomy VI*, Vol. 9911 (International Society for Optics and Photonics, 2016) p. 991105.
- [77] A. Kiessling, M. Cacciato, B. Joachimi, D. Kirk, T. D. Kitching, A. Leonard, R. Mandelbaum, B. M. Schäfer, C. Sifón, M. L. Brown, *et al.*, *Space Science Reviews* **193**, 67 (2015).
- [78] A. Izard, P. Fosalba, and M. Crocce, *Monthly Notices of the Royal Astronomical Society* **473**, 3051 (2017).
- [79] A. Kiessling, A. Heavens, A. Taylor, and B. Joachimi, *Monthly Notices of the Royal Astronomical Society* **414**, 2235 (2011).
- [80] H.-J. Huang, T. Eifler, R. Mandelbaum, and S. Dodelson, arXiv preprint arXiv:1809.01146 (2018).
- [81] S. Samuroff, J. Blazek, M. Troxel, N. MacCrann, E. Krause, C. Leonard, J. Prat, D. Gruen, S. Dodelson, T. Eifler, *et al.*, arXiv preprint arXiv:1811.06989 (2018).
- [82] J. Alsing, D. Kirk, A. Heavens, and A. H. Jaffe, *Monthly Notices of the Royal Astronomical Society* **452**, 1202 (2015).
- [83] C. A. Duncan, B. Joachimi, A. F. Heavens, C. Heymans, and H. Hildebrandt, *Monthly Notices of the Royal Astronomical Society* **437**, 2471 (2013).
- [84] E. M. Huff and G. J. Graves, *The Astrophysical Journal Letters* **780**, L16 (2013).
- [85] H. Hildebrandt, L. van Waerbeke, D. Scott, M. Béthermin, J. Bock, D. Clements, A. Conley, A. Cooray, J. Dunlop, S. Eales, *et al.*, *Monthly Notices of the Royal Astronomical Society* **429**, 3230 (2013).
- [86] A. Heavens, J. Alsing, and A. H. Jaffe, *Monthly Notices of the Royal Astronomical Society: Letters* **433**, L6 (2013).