

Taking the weight off: mitigating parameter bias from catastrophic outliers in 3×2 pt analysis

Carolyn McDonald Mill ^a, Leonard C. Danielle ^a, Markus Michael Rau ^a,
Cora Uhlemann ^b and Shahab Joudaki ^c

^a*School of Mathematics, Statistics and Physics, Newcastle University,
Herschel Building, NE1 7RU Newcastle-upon-Tyne, U.K.*

^b*Fakultät für Physik, Universität Bielefeld,
Postfach 100131, 33501 Bielefeld, Germany*

^c*Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT),
Av. Complutense 40, E-28040 Madrid, Spain*

*E-mail: c.mill2@newcastle.ac.uk, danielle.leonard@newcastle.ac.uk,
markus.rau@newcastle.ac.uk, cuhlemann@physik.uni-bielefeld.de,
shahab.joudaki@ciemat.es*

ABSTRACT: Stage IV cosmological surveys will map the universe with unprecedented precision, reducing statistical uncertainties to levels where unmodelled systematics can significantly bias inference. In particular, photometric redshift (photo- z) errors and intrinsic alignments (IA) must be robustly accounted for to ensure accurate inference of cosmological parameters. The increasing depth of Stage IV surveys exacerbates these challenges by producing low signal-to-noise galaxy populations prone to inaccurate photo- z measurements. Catastrophically misidentified redshifts are especially problematic for 3×2 pt inferences that combine weak lensing and galaxy clustering information. We demonstrate that even modest outlier fractions (e.g. 5%) can lead to substantial biases in cosmological parameter estimates: up to 1.8σ in Ω_M and σ_8 , and over 8σ in the IA redshift evolution parameter η . To address this, we introduce a flexible weighting scheme at the likelihood level that down-weights the most contamination-sensitive elements of the data vector during inference. This method mitigates biases without inflating the parameter space, reducing cosmological parameter biases to below 1σ without substantially degrading constraining power. Our approach offers a practical solution for future analyses, enabling robust cosmological inference in the presence of catastrophic redshift errors.

KEYWORDS: Bayesian reasoning, cosmological parameters from LSS, galaxy clustering, weak gravitational lensing

ARXIV EPRINT: [2509.08052](https://arxiv.org/abs/2509.08052)

Contents

1	Introduction	1
2	Set-up and modelling	3
2.1	Photometric redshifts	3
2.2	3×2 pt	5
2.3	Parameter inference	7
3	Results	9
3.1	Inferred cosmology response to outlier misspecification	11
3.2	IA and photo- z parameters	14
4	Discussion and conclusions	14
A	Full parameter constraints	17

1 Introduction

Upcoming and ongoing Stage IV surveys, such as the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) [2], *Euclid* [3, 4], and the Nancy Grace Roman Space Telescope [5], will offer an unprecedented scale and precision of late-time cosmological data. These surveys promise to help us answer pressing cosmological questions such as the nature of dark energy [4, 6–8].

Weak lensing [9, 10] is a powerful probe of cosmological parameters, particularly those which characterise dark energy [11–13]. Typical analyses often detect this effect via two-point correlations of the shapes of galaxies lensed by the matter distribution between them and the observer. This allows us to extract information about the intervening matter along our line of sight to them - a signal known as cosmic shear. Cosmic shear can be complemented with two-point correlations capturing galaxy clustering and galaxy-galaxy lensing - the cross-correlation between shapes and positions - to produce a combined set of two-point correlations, often referred to as 3×2 pt statistics.

However, as the precision of our data improves, the dominant source of error shifts from statistical uncertainties to systematic errors, i.e., those arising from our modelling [14–18]. It is therefore essential to adequately include systematic effects in our models and understand how they may impact our analysis, ensuring that this influx of high-quality data can be fully exploited.

In order to probe the nature of dark energy, we require information on how the large-scale structure has evolved over time. Galaxy redshifts are essential for tracking this evolution. Given the billions of galaxies targeted by surveys like LSST and the corresponding hypothetical cost of obtaining spectra for all galaxies, redshifts for such surveys are estimated primarily through photometric methods - observing galaxies in a set of broad bands and applying redshift estimation techniques such as template fitting and machine learning [19]. While photometric

redshifts (photo- z) require less exposure time and are less computationally intensive to obtain than spectroscopic redshifts (spec- z), they are also significantly less accurate [20]. To account for photo- z uncertainties in tomographic analyses, nuisance parameters that characterise uncertainties in galaxy redshift distributions are introduced in the data model. Often they are parametrised as small shifts in the mean and width of redshift distributions, and their value is inferred alongside cosmological parameters [16, 21]. However, it remains unclear whether this approach sufficiently captures the whole range of ways in which photo- z errors qualitatively impact source galaxy redshift distributions.

One type of photo- z error which is generally un-modelled in state-of-the-art cosmological analyses is the presence of catastrophic outliers: redshift estimates whose residuals are several times larger than the typical photo- z scatter, placing the galaxy in the wrong tomographic bin or at a vastly different cosmic epoch. Multiple photometric redshifts are often plausible for a given photometry, as demonstrated in [22]. Such outliers commonly originate from degeneracies in spectral energy distribution (SED) fitting, where distinct spectral breaks produce similar broadband colours — for instance, confusion between the Lyman break at high redshift and the Balmer break at low redshift [23]. Failing to account for catastrophic outliers has been shown to bias our inference of cosmological parameters [24, 25]. These outlier populations are difficult to model because of the scarcity of spectroscopic redshifts at the depths probed by Stage IV surveys. This makes it challenging to construct representative training samples to reconstruct these populations [26].

In addition to photo- z systematics, weak lensing measurements are also impacted by intrinsic galaxy alignments (IA) [27–29]. The observed cosmic shear signal consists of correlations in both the distortion of galaxy images due to the cosmic structure and the intrinsic galaxy shapes. Correlated intrinsic shapes, caused by physical effects such as tidal forces, can contaminate the weak lensing signal. Understanding and correctly modelling intrinsic alignments is thus crucial for extracting accurate cosmological information, as incorrect parametrisation of this contribution has been shown to be susceptible to biasing cosmological parameters [30, 31].

Previous studies have found that significant bias in cosmological parameters can be induced by misspecifying the modelling of either IA [30–32] or photo- z uncertainty [33–35], where misspecification refers to adopting an analysis model that omits or inaccurately represents relevant physical observational effects. Moreover, a coupling between these systematics has been identified [36, 37], and it has been demonstrated that the degeneracies between their parametrisations can further exacerbate these biases [38, 39].

In this work, we investigate the sensitivity of cosmological parameter inference to the misspecification of photo- z and IA modelling. We perform a Markov Chain Monte Carlo (MCMC) analysis on a synthetic 3×2 pt LSST-like data set, implementing parametrisations for IA and photo- z , along with a model for an outlier population arising from the confusion between Lyman and Balmer breaks. We examine the impact of misspecifying this outlier population on our cosmological and IA parameters and investigate the potential of modifying the likelihood as a method to reduce this effect.

We describe the models used for inference and for quantifying model misspecification in section 2. In section 3 we describe the numerical choices made for this analysis and present




Figure 1. The cosmological inference process from a galaxy sample over redshift range z , to summary statistic C_ℓ , to parameters of interest θ_i .

our results, focusing on constraints on cosmological parameters Ω_M, σ_8 , dark energy equation of state parameters w_0, w_a , and IA parameters for the non-linear alignment model a, η . We discuss our findings and outline our conclusions in section 4.

2 Set-up and modelling

In figure 1 we sketch the inference process for determining cosmological parameter values; we will provide details of how each step was carried out for this work in the indicated subsections. From a population of galaxies described by a photo- z distribution (2.1) we obtain summary statistics (in our case the angular correlation functions, C_ℓ) that we predict from a model with a chosen parameter set (2.2). We can then use these statistics to infer information about cosmological parameters from the data (2.3). However, it is possible that our model does not perfectly represent the data, which can lead to biases in the inferred values of parameters θ_i . In this section, we also discuss our method to test how well a given model defined by a set of parameters represents the data. We then introduce composite likelihoods as a method of improving model performance under misspecification.

2.1 Photometric redshifts

Photo- z methods allow redshift information to be obtained with less exposure time and more cost-effectively than with spec- z methods across the vast number of galaxies observed in large-scale sky surveys. In cosmological analysis, galaxies with similar redshifts are typically grouped into tomographic redshift bins, enabling us to probe the statistical behaviour of structures at different epochs rather than on a per-galaxy basis. In this work, we use mock redshift distributions for the lens and source samples (used for clustering and cosmic shear analyses respectively, with both in use for galaxy-galaxy lensing), each convolved with Gaussian distributions to represent a Gaussian model for photo- z mean shift and variance uncertainty.




Figure 2. Redshift distributions for the five bins of the lens (left) and source (right) samples.

Throughout this work, tomographic bins are indexed from 0 to 4 (rather than the conventional 1 to 5) for consistency with our analysis pipeline, with 0 being the lowest-redshift bin.

Mean shifts. Small errors in estimated galaxy redshifts can arise from noise in the photometric measurements, limitations of the template or training set used for estimation, and mild degeneracies between spectral features. To account for these uncertainties in inference, cosmological analyses often introduce nuisance parameters that allow small shifts in the mean of each redshift bin to be marginalised over [16, 21, 40]. In this analysis, we model photo- z uncertainty using two sets of five nuisance parameters, Δz^i , which allow for independent shifts in the mean redshift of each tomographic lens and source bin. These parameters modify the redshift bins such that

$$z_{new}^i = z^i + \Delta z^i. \quad (2.1)$$

Catastrophic Outliers. Catastrophic outliers concern any case where the photometric redshift distribution diverges significantly from the true redshift distribution of the sample. This can be caused by the blending of samples or the redshifting of distinctive spectroscopic features outside of the photometric broadbands [22].

A dominant case of outliers is caused by the degeneracy between high- z Lyman-break galaxies and low- z Balmer break galaxies, producing a contaminating population of galaxies that are determined to have redshift values $z \sim 0 - 0.5$ when their true redshift will lie around $z \sim 2.5 - 3.5$ [41].

In the case of LSST, the source sample is expected to have a more prominent presence of outliers due to the less stringent requirements on the photo- z uncertainty in these objects relative to the lenses.

We examine the effects of outliers on parameter inference by including the effect of a Lyman-Balmer confusion in the mock data. We do this by allowing for a fraction f of the total photo- z distribution of $\frac{dN_s^0}{dz}$ (the redshift distribution of source bin 0) to be an outlier population with a Gaussian redshift distribution centred at $z = 3.2$ with a width of 0.5, such that the total distribution of the photo- z in source bin 0 is given by

$$\frac{dN_s^0(f)}{dz} = (1 - f) \frac{dN_s^{0, fid}}{dz} + f \mathcal{N}(3.2, 0.5) \quad (2.2)$$

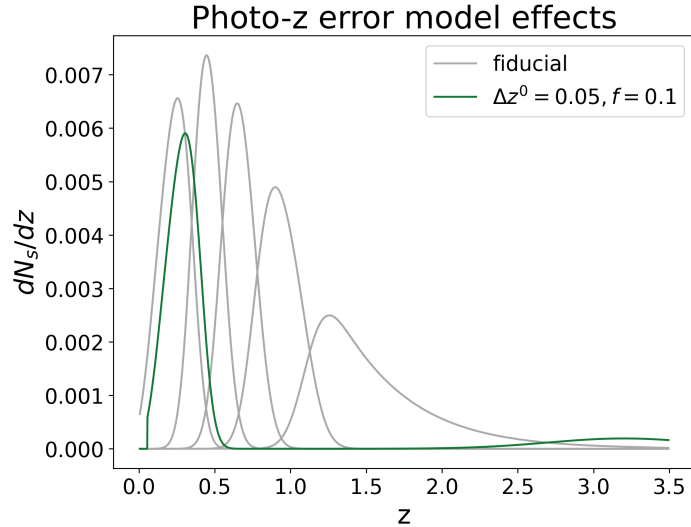


Figure 3. The photo- z distribution of source bin $\frac{dN_s^0}{dz}$ with mean shift $\Delta z^0 = 0.05$ and an outlier fraction of $f = 0.1$. We use unrealistically large values here for clear visualisation. The fiducial source bins are plotted in grey for comparison.

This distribution is chosen to approximate the prominent case of outliers due to the confusion between Lyman break and Balmer break galaxies [41]. We demonstrate the joint effect that a mean shift $\Delta z^0 = 0.05$ and an outlier fraction of $f = 0.1$ has on the redshift distribution as the green line in figure 3.

2.2 3×2 pt

3×2 pt is the expected baseline approach for analysing weak lensing and photometric galaxy clustering in Stage IV surveys. It combines the constraining power from two-point statistics of galaxy clustering, galaxy-galaxy lensing, and weak lensing using their two-point angular power spectra. The total datavector in our model is constructed using the 5 autocorrelations between the galaxy clustering redshift bins, 7 galaxy-galaxy lensing redshift bin combinations¹ and the 15 unique cosmic shear auto and cross redshift bin combinations.

The two-point angular auto-power spectra for galaxy clustering can be formulated as

$$C_{g^i g^i}^\ell = b_i^2 \int_0^{\chi_\infty} d\chi \left(\frac{dN_l^i}{d\chi} \right)^2 P_{mm} \left(\frac{\ell}{\chi}, \chi \right), \quad (2.3)$$

where b_i denotes the linear galaxy bias, $\frac{dN_l^i}{d\chi}$ is the photometric redshift distribution in photo- z bin i with respect to comoving distance χ , χ_∞ is the comoving distance at the horizon, and P_{mm} is the non-linear matter power spectrum. We assume a flat universe and adopt the Limber approximation throughout.

The galaxy-galaxy lensing angular power spectra are given by

$$C_{\kappa^i g^j}^\ell = b_j \int_0^{\chi_\infty} d\chi \frac{W_i(\chi)}{\chi} \frac{dN_l^j}{d\chi} P_{mm} \left(\frac{\ell}{\chi}, \chi \right), \quad (2.4)$$

¹[0,2], [0,3], [0,4], [1,3], [1,4], [2,4], [3,4], taken from [42]

where the lensing efficiency kernel $W_i(\chi)$ defined as

$$W_i(\chi) = \frac{3H_0^2\Omega_M}{2c^2} \frac{\chi}{a(\chi)} \int_{\chi}^{\chi_\infty} \frac{\chi' - \chi}{\chi'} \frac{dN_s^i}{d\chi'} d\chi'. \quad (2.5)$$

Finally, the weak lensing angular power spectra are given by

$$C_{\kappa^i \kappa^j}^\ell = \int_0^{\chi_\infty} d\chi \frac{W_i(\chi)W_j(\chi)}{\chi^2} P_{mm} \left(\frac{\ell}{\chi}, \chi \right). \quad (2.6)$$

We computed the angular power spectra for the MCMC using the `pyccl`² Python package [43], where the matter power spectra were calculated using 1-loop perturbation theory, with calculations carried out using FASTPT [44, 45]. We assumed a w_0w_a CDM cosmology, parametrised with Ω_B , Ω_M , n_s , σ_8 , h and dark energy equation of state parameters w_0 , w_a .

Intrinsic Alignments. The observed projected galaxy shape (or ellipticity) field receives contributions both from weak lensing and from IA. This total effective shear field can be expressed as

$$\gamma^\epsilon = \gamma^\kappa + \gamma^{IA}, \quad (2.7)$$

where we denote the observable shear with ϵ , the shear induced by lensing with κ and the intrinsic component due to galaxy shapes with IA. Intrinsic alignments arise from the tendency of galaxies to align their shapes with the tidal gravitational field of their surrounding matter environment, together with other environmental and galaxy-evolution related effects. In the widely used Non-Linear Alignment (NLA) model, the IA shear field is assumed to respond linearly to the nonlinear local tidal field tensor [46]. The IA shear is then given by

$$\gamma_{ij}^{IA} = A_1 s_{ij}, \quad (2.8)$$

where $s_{ij} = \partial_i \partial_j \phi$ is the tidal field tensor derived from the gravitational potential ϕ . The amplitude of the intrinsic alignment signal is often taken to be a redshift dependent function $A_1(z)$, which captures the strength and redshift evolution of the alignment. Following the formulation in [17], this can be expressed as

$$A_1(z) = -a \frac{\bar{C}_1 \rho_c \Omega_M}{D(z)} \left(\frac{1+z}{1+z_{piv}} \right)^\eta, \quad (2.9)$$

where ρ_c is the critical density, $D(z)$ is the linear growth factor, the pivot redshift $z_{piv} = 0.62$, and the normalisation factor $\bar{C}_1 = 5 \times 10^{-14} M_\odot h^{-2} \text{Mpc}^2$.

For cosmic shear, the total angular power spectra for the sum of the cosmic shear κ and the intrinsic component I can be written as

$$C_{\epsilon^i \epsilon^j}^\ell = C_{\kappa^i \kappa^j}^\ell + C_{\kappa^i I^j}^\ell + C_{I^i \kappa^j}^\ell + C_{I^i I^j}^\ell \quad (2.10)$$

and the galaxy-galaxy lensing case as

$$C_{\epsilon^i g^j}^\ell = C_{\kappa^i g^j}^\ell + C_{I^i g^j}^\ell. \quad (2.11)$$

²<https://github.com/LSSTDESC/CCL>

These IA contributions are calculated through line-of-sight projections of their respective 3D power spectra

$$C_{\kappa^i I^j}^\ell = \int_0^{\chi_\infty} d\chi \frac{W_i(\chi)}{\chi} \frac{dN_s^j}{d\chi} P_{mI} \left(\frac{\ell}{\chi}, \chi \right), \quad (2.12)$$

$$C_{I^i I^j} = \int_0^{\chi_\infty} \frac{1}{\chi^2} \frac{dN_s^i}{d\chi} \frac{dN_s^j}{d\chi} P_{II} \left(\frac{\ell}{\chi}, \chi \right), \quad (2.13)$$

$$C_{I^i g^j} = b_j \int_0^{\chi_\infty} \frac{dN_s^i}{d\chi} \frac{dN_l^j}{d\chi} P_{Im} \left(\frac{\ell}{\chi}, \chi \right). \quad (2.14)$$

Each of the 3D power spectra can be expressed in terms of the redshift-dependent amplitude $A_1(z)$, enabling IA modelling to be folded into the same parameter inference framework as lensing and clustering — see [47] for a detailed theoretical description. In this work we choose to use the NLA model for IA over other more complex modelling frameworks (e.g. Tidal Alignment - Tidal Torquing [47], the halo model for IA [48], or Effective Field Theory of IA [49]). This is because we are primarily focused on the impact of the photo- z model and particularly the inclusion of an outlier population. When introducing this additional complexity in the photo- z space, we benefit from employing a minimal realistic IA model to avoid over-expanding complexity in multiple domains simultaneously and further inflating the parameter space dimensionality. Future work considering outlier rates in combination with more complex IA models would be of value.

2.3 Parameter inference

The cosmological inference process involves sampling from a posterior distribution, which combines prior knowledge (expressed as probability distributions) with a likelihood function that quantifies how well a given parametric model fits the observed data. Under a typical Gaussian likelihood approximation the natural logarithm of the likelihood \mathcal{L} is given by

$$\log \mathcal{L} \propto -0.5(\mathcal{D} - \mathcal{M}_\theta)^T \Sigma^{-1} (\mathcal{D} - \mathcal{M}_\theta) \quad (2.15)$$

where in a typical real analysis, \mathcal{D} would be a measured datavector. For us, in the simulated analyses of this work, it is a synthetic datavector produced with the fiducial parameters given in table 1, acting in place of our observed data. \mathcal{M}_θ is the datavector produced by our model for a parameter set θ , and Σ is the data covariance matrix, assumed to be parameter-independent. In this section we describe the construction of each of these components as used in this analysis.

To sample efficiently over the parameter space, we employ the Affine Invariant Markov Chain Monte Carlo (MCMC) method implemented using the `emcee` package [50]. This method is well-suited to exploring high-dimensional parameter spaces, as it adapts to the geometry of the posterior distribution and mitigates issues related to correlated parameters.

The full parameter space θ consists of 7 $w_0 w_a$ CDM parameters together with 5 linear galaxy bias parameters $\{b_i\}$, 2 intrinsic alignment parameters for the NLA model $\{a, \eta\}$, and 10 photo- z uncertainty parameters $\{\Delta z_l^i, \Delta z_s^i\}$.

Posterior Predictive Checks: To assess the suitability of a given model for our data, we employ posterior predictive checks (PPCs), a standard technique in Bayesian model evaluation [51]. In the PPC framework, we first draw N parameter vectors θ^i from the posterior distribution $p(\theta|\mathcal{D})$. For each draw, we generate a replicated datavector \mathcal{M}_θ^i from the likelihood $p(\mathcal{M}_\theta|\theta^i)$. $p(\theta|\mathcal{D})$ and $p(\mathcal{M}_\theta|\theta)$ together define the posterior predictive distribution:

$$p(\mathcal{M}_\theta|\mathcal{D}) = \int p(\mathcal{M}_\theta|\theta)p(\theta|\mathcal{D})d\theta. \quad (2.16)$$

This produces a collection of N replicated datavectors, \mathcal{M}_θ^i , that can be directly compared to the observed data \mathcal{D} . In this work, we perform a visual comparison by overplotting the model predictions and mock data to identify potential model misspecification, an example of which is shown in figure 5.

In addition, we implement the goodness-of-fit test proposed in [51] (Equation 3). For this purpose, we use the Mahalanobis distance [52], which we define as

$$\chi^2 = (y - \overline{\mathcal{M}}_\theta)^T \Sigma^{-1} (y - \overline{\mathcal{M}}_\theta), \quad (2.17)$$

where y is the datavector which we are testing the goodness-of-fit of and $\overline{\mathcal{M}}_\theta$ is the mean of the posterior predictive datavectors \mathcal{M}_θ . This metric provides a quantitative measure of how well the model predictions match the observed data. If the model is a good representation of the data, we would expect the χ^2 value of the observed (or mock) data to lie within the distribution of χ^2 values computed from the replicated datavectors \mathcal{M}_θ^i .

Composite Likelihoods: In this work, we investigate whether leveraging composite likelihoods, following [53], can help reduce the effects of model misspecification. A comprehensive review of these methods is provided by [54]. During inference we usually resolve the likelihood given in equation (2.15) to determine the probability of the proposed parameter set describing the data. After identifying components of the likelihood that are poorly modelled due to misspecification using PPCs, we can isolate these components and down-weight them, arriving at the composite likelihood

$$\log \mathcal{L} \propto \gamma \log \mathcal{L}_P + \log \mathcal{L}_R, \quad (2.18)$$

where subscripts P and R refer to the punitive and representative components, respectively and $\gamma \in [0, 1]$ is a weighting factor that limits the importance of the punitive part. To obtain the respective parts $\mathcal{L}_{P/R}$, we split the data vector \mathcal{D} and the associated covariance matrix Σ into blocks

$$\mathcal{D} = \begin{pmatrix} \mathcal{D}_R \\ \mathcal{D}_P \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \Sigma_{RR} & \Sigma_{RP} \\ \Sigma_{PR} & \Sigma_{PP} \end{pmatrix}. \quad (2.19)$$

The likelihood depends on the inverse of the data covariance, which we obtain by considering the covariance as a 2×2 block matrix to identify the separate parts. It is useful to define the two objects

$$\overline{\mu} = \mathcal{M}_P + \Sigma_{PR} \Sigma_{RR}^{-1} (\mathcal{D}_R - \mathcal{M}_R), \quad \overline{\Sigma} = \Sigma_{PP} - \Sigma_{PR} \Sigma_{RR}^{-1} \Sigma_{RP}, \quad (2.20)$$

where \mathcal{D} and \mathcal{M} indicate the mock data and model datavectors respectively. Here $\bar{\mu}$ and $\bar{\Sigma}$ are the mean and covariance of the conditional distribution

$$P(\mathcal{D}_P | \mathcal{M}_P, \Sigma_{PP}, \mathcal{M}_R, \Sigma_{RP}, \mathcal{D}_R) = \mathcal{N}(\bar{\mu}, \bar{\Sigma}). \quad (2.21)$$

We can then extract the punitive component given by the conditional Gaussian distribution of \mathcal{D}_P given \mathcal{D}_R

$$\log \mathcal{L}_P \propto -0.5 (\mathcal{D}_P - \bar{\mu})^T \bar{\Sigma}^{-1} (\mathcal{D}_P - \bar{\mu}), \quad (2.22)$$

and the representative component

$$\log \mathcal{L}_R \propto -0.5 (\mathcal{D}_R - \mathcal{M}_R)^T \Sigma_{RR}^{-1} (\mathcal{D}_R - \mathcal{M}_R). \quad (2.23)$$

We can then repeat the inference process, instead resolving this composite likelihood, with identifiably misspecified components given less importance. Our choice of γ can be informed by the PPC from the original inference. As MCMC methods only require the relative probability between samples to evaluate the posteriors, it is not required that $\log \mathcal{L} = \gamma \log \mathcal{L}_P + \log \mathcal{L}_R$, but only that the change in probability between samples is conserved. Therefore we do not require the down-weighting of $\log \mathcal{L}_P$ to be counter-weighted in $\log \mathcal{L}_R$.

3 Results

We performed this analysis in an LSST Year 1 context, which has an expected survey area of $12.3 \times 10^3 \text{ deg}^2$ [38]. The fiducial values for the linear galaxy biases b_i , the galaxy-galaxy lensing bin combinations, and the scale cuts in ℓ , as well as the source and lens redshift distributions (figure 2) and the non-Gaussian data covariance matrix, were taken from the LSST DESC SRD [42]. This data covariance matrix was also used to produce the noise for scenarios where noise was included in the mock data.

We used the fiducial parameters shown in table 1 to generate the mock data used throughout our analysis. In this table we also present the priors placed on these parameters during inference for most scenarios considered in this work and provide the sources that informed these choices. However, when ‘‘centred priors’’ are used during inference, all Gaussian priors are defined such that their mean lies at the fiducial value for that parameter. This only affects w_a and Δz^i priors when implemented, as all other Gaussian priors are always centred on the fiducial parameters. The uniform priors applied to b_g^i remain unchanged throughout.

In this work, we consider six scenarios summarised in table 2:

1. We produce mock data using the fiducial parameters from table 1, and include no outlier fraction and no noise. We use centred priors, modifying the priors applied such that all Gaussian priors are centred on the fiducial parameters used to produce the data. We then carry out the inference process using the model described in section 2. This could be considered a perfectly modelled case, and we use this as a reference for the other scenarios we consider.
2. We repeat the inference from Scenario 1, but implement the priors as described in table 1. This simulates a standard inference process where all systematics are modelled, and there is no outlier fraction present in the data.

θ	Ω_M	σ_8	h	Ω_B	n_s	w_0	w_a	a	η
θ_{fid}	0.3156	0.830	0.6727	0.0492	0.9645	-1.00	0.05	0.44	-0.7
μ_p	0.3156,	0.831	0.6727	0.0492	0.9645	-1.00	0.00	0.44	-0.7
σ_p	$\Sigma_{\mathbf{p}}$ [42]							0.38 [55]	2.2 [55]

θ	b_g^i	$\Delta z_l^i [10^{-3}]$	$\Delta z_s^i [10^{-3}]$
θ_{fid}	{1.56, 1.73, 1.91, 2.10, 2.29}	{7.1, 1.6, 3.9, 9.0, 7.5}	{-3.9, 3.8, -0.6, 0.4, 1.6}
prior	$\mathcal{U}(0.8, 3)$ [17, 55]	$\mathcal{N}(0, 0.004)$ [56]	

Table 1. Fiducial parameters used for the mock data in the MCMC analysis, along with the priors applied during inference and their source. Priors in the upper table are Gaussians with a mean μ_p and standard deviation σ_p . The first seven cosmological parameters are drawn from a joint multivariate Gaussian prior $\mathcal{N}(\mu_p, \Sigma_{\mathbf{p}})$, where $\Sigma_{\mathbf{p}}$ corresponds to the Stage III prior parameter covariance matrix including w_0, w_a from [42].

Scenario	f_D	Noise	Centred Priors	Composite \mathcal{L}
1	0.00	No	Yes	No
2	0.00	No	No	No
3	0.02	No	No	No
4	0.02	Yes	No	No
5	0.05	No	No	No
6	0.05	No	No	Yes

Table 2. Summary of scenarios considered in this work. f_D : the outlier fraction included in the datavector. *Noise*: Whether Gaussian noise was included in the mock data (Yes) or not (No). *Centred Priors*: Whether all priors were centred on fiducial parameter values (Yes) or were as described in table 1 (No). *Composite \mathcal{L}* : Whether the composite likelihood method described in section 2.3 was applied during inference (Yes) or not (No).

3. We add an outlier fraction of $f_D = 0.02$ to the mock data. As a parametrisation of the outlier fraction is not included in the model, this allows us to determine the robustness of the inference process in the case of a small misspecification error in outlier fraction.
4. We include an outlier fraction of $f_D = 0.02$ in the mock data and add Gaussian noise to determine the effect of a small outlier fraction on parameter inference in a more realistic case.
5. We include an outlier fraction of $f_D = 0.05$ in the mock data to determine the effect of an increased outlier fraction on parameter inference.
6. We include an outlier fraction of $f_D = 0.05$ in the mock data and down-weight the likelihood components impacted by the outlier fraction by a factor of $\gamma = 0.05$ using the method described in section 2.3. We use this to test whether this method can improve parameter constraints in the presence of misspecification.

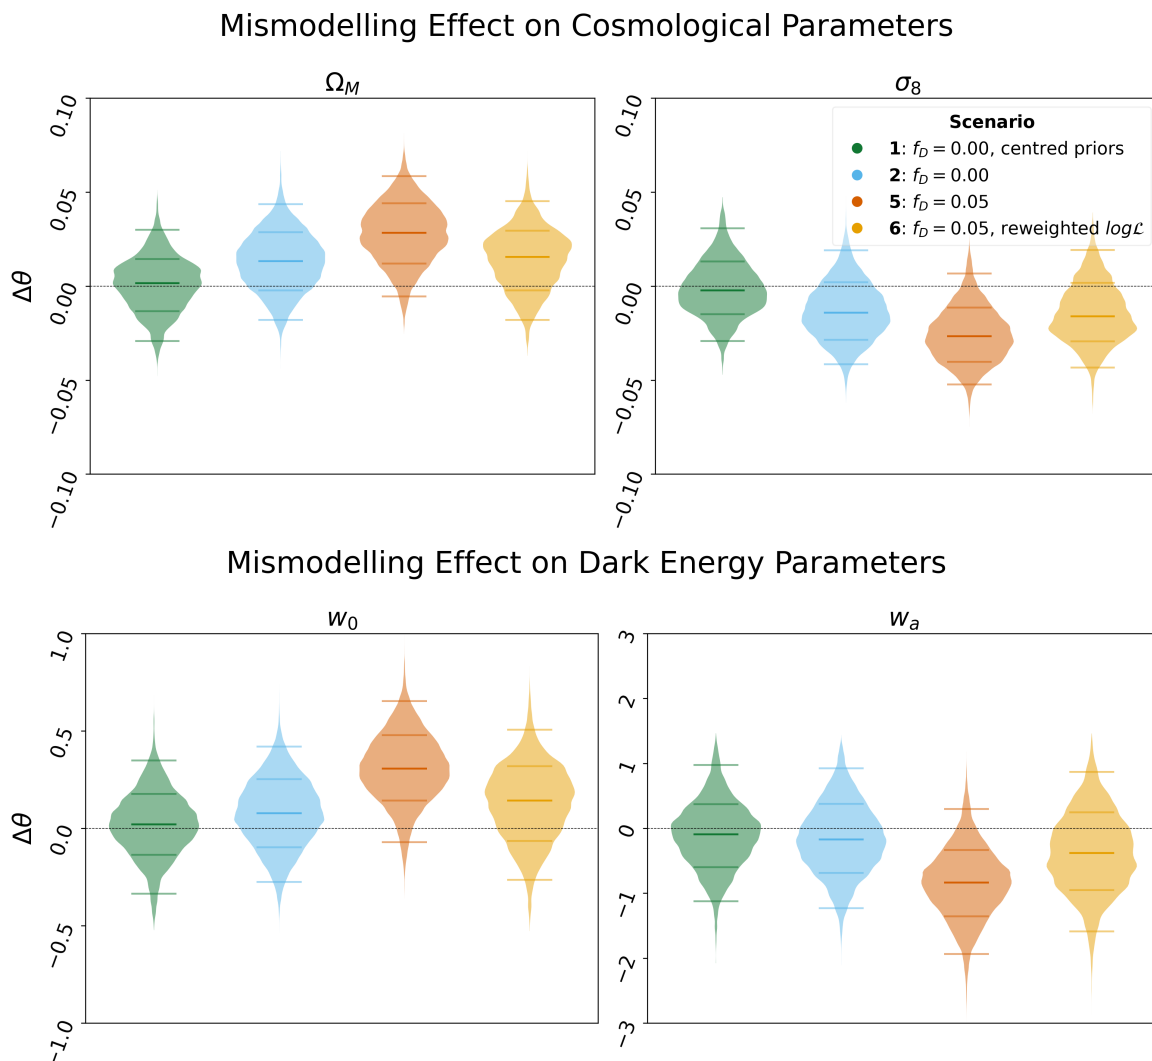


Figure 4. The posterior probability distributions for cosmological parameters Ω_M, σ_8 (upper) and dark energy parameters (lower) for a subset of the scenarios described in table 2. $\Delta\theta$ indicates the difference between the maximum a posteriori and the fiducial value (table 1), where $\Delta\theta = 0$ is indicated with the dashed black line. Mean and 1σ and 2σ are indicated by the dark and light lines.

We confirmed that MCMC chains were converged by calculating the Gelman-Rubin statistic [57] between the first half and second half of the chain along iterations, after removing burn-in.

3.1 Inferred cosmology response to outlier misspecification

We show the results of the parameter inference process for the cosmological and dark energy parameters in figure 4, focussing here on the most illustrative of the Scenarios from table 2 (with results for all Scenarios shown in appendix A). We found that we were able to obtain cosmological parameters (Ω_M and σ_8) near the fiducial values when we included no outliers in the data and centred all priors on the fiducial values used to build the datavector (Scenario 1), which we can consider a perfectly modelled scenario. A bias of 1σ in the cosmological

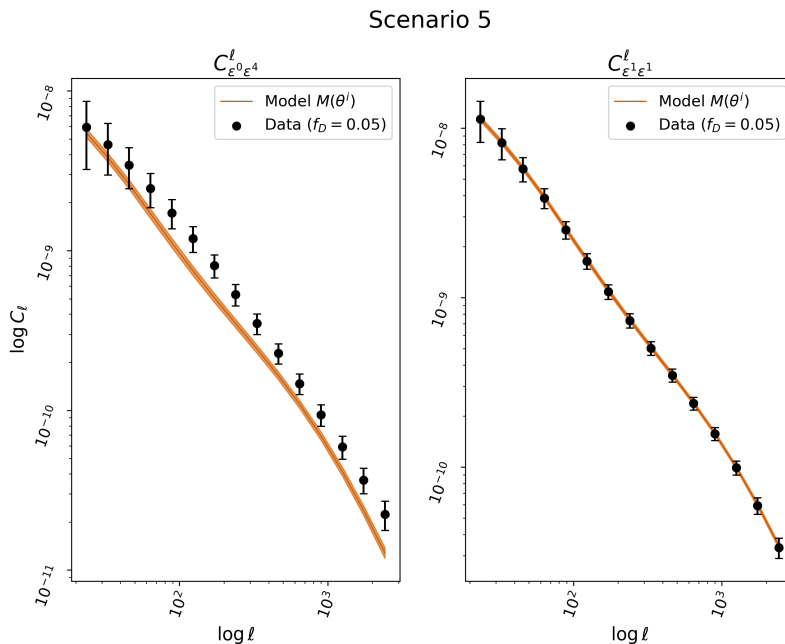


Figure 5. The posterior predictive diagnostic plots for Scenario 5 for two weak lensing angular power spectra in the mock data. Catastrophic outliers included in the lowest bin affect cross-correlations significantly (left) while leaving the autocorrelation in higher bins unaffected (right). We plot the model as the mean of the 500 datavectors produced from sampling the posterior, with the shaded region indicating the 2σ error, and plot the mock data as data points with 2σ error bars drawn from the data covariance matrix.

parameters was produced by applying priors which are not centred on the true parameter values as described in table 1 (Scenario 2), with a milder effect on the dark energy parameters. Once an unmodelled outlier fraction of $f_D = 0.05$ was included in the data (Scenario 5), a bias of 1.8σ was produced in Ω_M , σ_8 , and w_0 with a slightly smaller bias produced in w_a .

We can detect the poor fit of the model to the data visually by implementing PPCs, as shown in figure 5. Considering Scenario 5, we plot the mock data for the cosmic shear cross-correlation $C_{\epsilon^0 \epsilon^4}^\ell$ along with the mean and standard deviation of the reconstructed data from 500 samples from the posterior (from here on referred to as replicated data). We also plot this for $C_{\epsilon^1 \epsilon^1}^\ell$; a component of the datavector that is not directly affected by misspecification in the outlier fraction. The outlier population induces a bias in the data that the model is unable to fit to, producing a poor reproduction of the angular power spectra. This misspecification only produces a poor fit in angular power spectra that include the tomographic redshift bin $\frac{dN_s^0}{dz}$ containing the outlier population (i.e. $C_{\epsilon^0 \epsilon^0-4}^\ell$), while spectra for other bin combinations (e.g. $C_{\epsilon^1 \epsilon^1}^\ell$) remain unaffected. By visual comparison then, we can identify cases of misspecification and inform where in the model the misspecification may be implemented.

We draw attention to the difficulty of identifying this misspecification when noise is included on the mock data in figure 6, where we display the mock data and replicated data for Scenarios 3 and 4. These have identical set-ups apart from the inclusion of Gaussian noise in the mock data in Scenario 4. In the noiseless case of Scenario 3, a consistent bias




Figure 6. The posterior predictive diagnostic plots for Scenario 3 and 4 for two weak lensing angular power spectra in the mock data. With an outlier fraction of $f = 0.02$ a bias is produced (left) that is not present in model components that don't contain the outlier population (centre left). However, this bias can be obscured when noise is added to the mock data (centre right). Nevertheless, one can still see disagreement between the replicated data and the mock data when comparing to components of the model not affected by the outlier fraction (right). We plot the model as the mean of the 500 datavectors produced from sampling the posterior, with the shaded region indicating the 2σ error, and plot the mock data as data points with 2σ error bars drawn from the data covariance matrix.

is still clear for an outlier fraction of $f = 0.02$, even though the replicated data are within 2σ of the mock data. In the noisy case, there is a clear divergence between the mock and replicated data, with the mock data biased by over 2σ at some values of ℓ . However, there is not a consistent positive or negative bias between the mock data and the replicated data in the presence of noise. Therefore, it may be difficult to identify cases of misspecification when analysing noisy data.

We implement Equation 2.17 to determine $\chi_{M^i}^2$ for each of the 500 replicated datavectors, as well as χ_D^2 for the mock data, to determine whether the model is representative of the data. We plot the results of this test in figure 7. We find that in Scenarios 3 and 5, it is clear that $C_{\epsilon^0 \epsilon^4}^{\ell}$ is misspecified while for $C_{\epsilon^1 \epsilon^1}^{\ell}$ the model is representative of the data. This is less clear when we consider Scenario 4, where the noise causes both C^{ℓ} s to be indicative of a poor fit. However, the $C_{\epsilon^0 \epsilon^4}^{\ell}$ component is a significantly worse fit when compared to $C_{\epsilon^1 \epsilon^1}^{\ell}$. Therefore, when considering the relative goodness-of-fit of components of the data it is possible to identify poorly modelled components of the data even in the presence of noise.

Now that the components of the datavector impacted by the misspecification have been identified using PPCs, we can choose to down-weight the likelihood of these as discussed in section 2.3. Due to the large discrepancy between the model and data for this section of the datavector, we choose to heavily down-weight the punitive component with $\gamma = 0.05$. Once this down-weighting is implemented, the bias in the cosmological and dark energy parameters

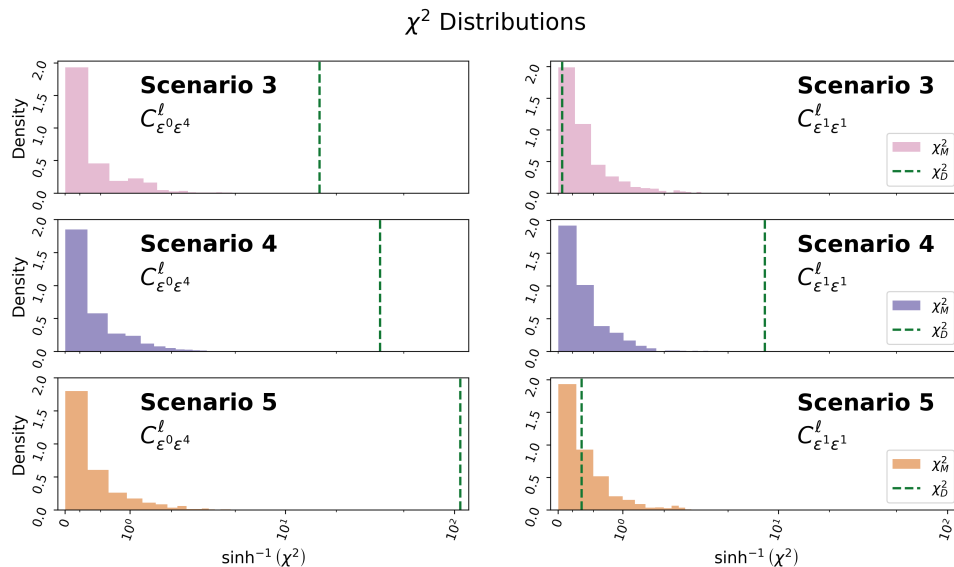


Figure 7. The distributions of χ_M^2 for the 500 replicated datavectors for Scenarios 3, 4, and 5 for a component of the datavector where an outlier fraction is present in one of the cross-correlated redshift bins (left) vs. not present (right), with the χ_D^2 value in the mock data shown as a green dashed line. The x-axis is rescaled as $\sinh^{-1}(\chi^2)$ to allow for clear visual comparison over the large ranges considered.

is reduced to less than 1σ . This shows that the composite likelihood approach is promising in mitigating the impact of unmodelled photo- z outliers.

3.2 IA and photo- z parameters

We also visualise the posteriors on the IA parameter and on a subset of photo- z parameters for the scenarios described in table 2, in addition to the case of the misspecification of a smaller outlier population, in figure 8. We see that we are able to achieve reasonably unbiased constraints when no outlier rate is included. An outlier fraction of $f_D = 0.02$ produces a 1σ bias in the overall IA amplitude a and a 2σ bias in the IA redshift-dependence η . This bias increases to 4σ in a and 8σ in η when this outlier fraction was increased to $f_D = 0.05$. When a reweighting of the likelihood was applied, this bias was reduced to a level comparable to the $f_D = 0.02$ case. The presence of the outlier population produces a bias in the Δz_s^i parameter, both in the bin containing the population and in the neighbouring bin. We expect that this is due to the high degeneracy between the photo- z uncertainty parameters. We find that the bins further from the one containing the outlier population (not shown) are not significantly biased. We find that the likelihood reweighting is also able to reduce the bias in the Δz_s^i parameters.

4 Discussion and conclusions

We first note that in considering the results shown in figures 4 and 8, the only case where the peak of the posterior distribution very closely aligns with the fiducial (true) parameters is that where the prior distributions were artificially centred on those fiducial parameters. We see that in all other cases considered, even with a completely accurate model, constraints

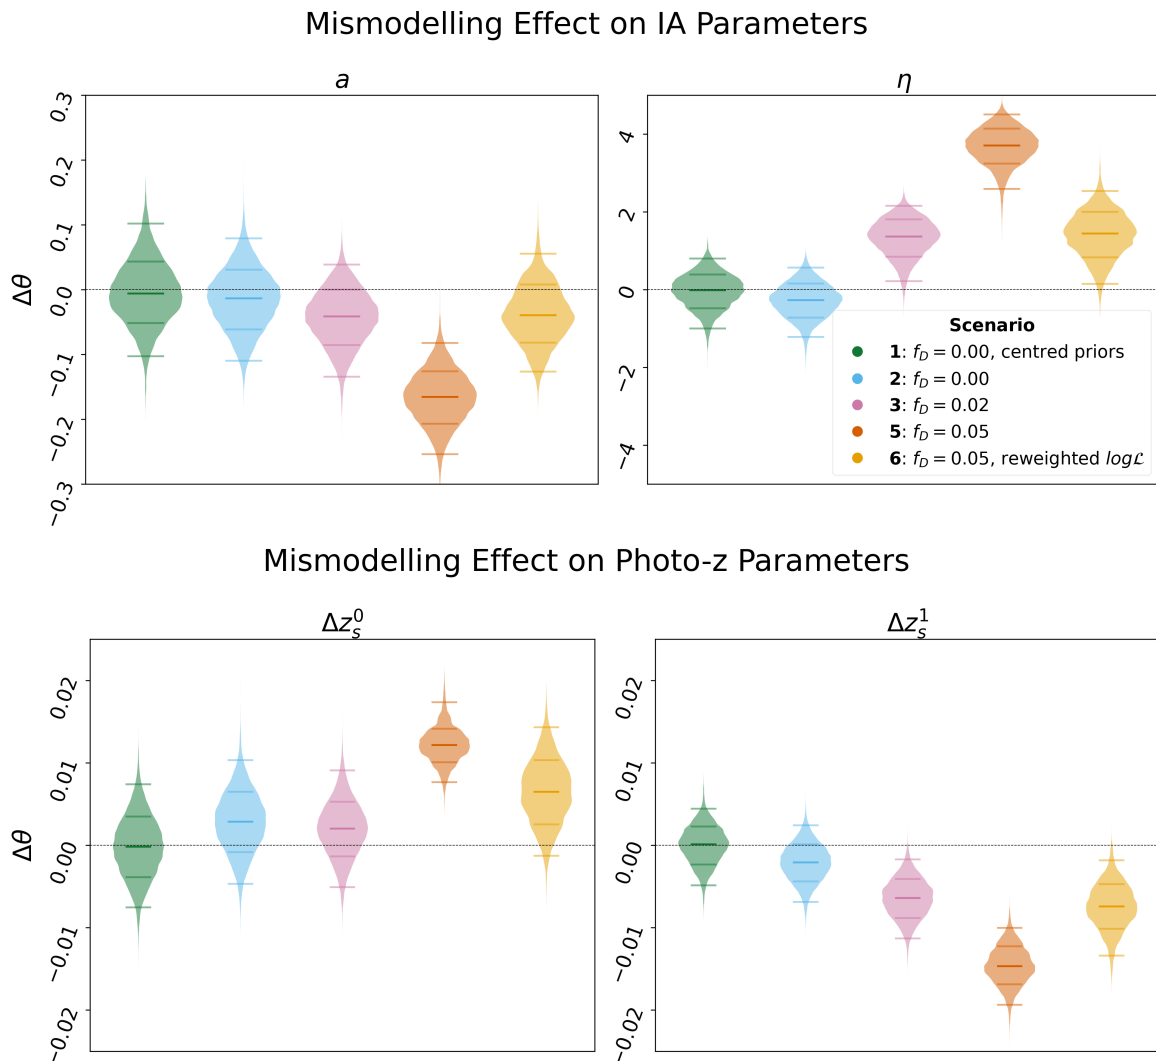


Figure 8. The posterior probability distributions for the NLA IA parameters; the overall amplitude a and the redshift-dependence η (upper) and for the Δz_s^i parameters for the first two source redshift bins (lower) for the scenarios described in table 2. $\Delta\theta$ indicates the difference between the maximum *a posteriori* and fiducial value (table 1), where $\Delta\theta = 0$ is indicated with the dashed black line. Mean and 1σ and 2σ are indicated by the dark and light lines.

are biased, in some parameters to nearly 1σ . This reinforces the fact we must be aware of possible effects on our constraints from informative prior distributions, which in reality cannot be generically centred on the true underlying parameter values describing our data.

In this work, we have considered a somewhat optimistic case where the model can perfectly reproduce the data, there is only one tomographic redshift bin with outliers, and those outliers are localised in one redshift region. This is instead of a more realistic case like that considered in [15] where all bins contain outlier populations, some with multiple such populations. However, even in the optimistic case considered here, a single outlier population can bias cosmological and dark energy parameters up to almost 2σ when it is not included in the model. We also see that misspecification of the outlier fraction can significantly bias

the IA parameters. This is not a surprising effect, since here we have considered an outlier population in the 0th redshift bin, and it is generally understood that lower-redshift cosmic shear holds more information about IA than the higher-redshift cases.

In order to avoid these biases, it may be tempting to propose that we simply model these populations during the inference stage. However, this would require an understanding of the distribution of these outlier galaxies that is not currently possible to achieve with the spectroscopic samples available at the depths considered in this work [20]. Furthermore, in this work we have considered only a single outlier population in a singular tomographic redshift bin, which only requires one additional model parameter. It is more likely that real data will contain multiple outlier populations across multiple bins, with shapes that are not well understood. To accurately model this would massively expand the parameter space of the analysis, and the degeneracy between these and other nuisance parameters (as shown in figure 8) risks making the likelihood intractably complex. However, as long as this misspecification can be identified through methods like posterior predictive checks, we have shown that the parameter bias can be reduced by down-weighting the poorly modelled components of the likelihood, with minimal effect on the constraining power.

In this case, with a single outlier population, we are aided somewhat in the application of the composite likelihood method by the simplicity of the scenario. A question which might arise in the case with multiple outlier populations (as in [15]) is whether it would be possible to effectively apply the composite likelihood method to down-weight relevant components of the data vector in the case where there are outlier populations in multiple redshift bins. Fortunately, the PPC method can help us to identify the most severely affected components of the datavector in this scenario, and an iterative approach can be used to achieve a balance between retaining information and achieving unbiased results. The composite likelihood method is also flexible: the likelihood can in fact be split into more than two components, with different γ factors being chosen to differentially down-weight different segments of the data vector that are more or less impacted by misspecification.

One might also be justifiably concerned that down-weighting the portions of the likelihood containing the outlier rate would negatively impact the constraints on the IA and photo- z parameters, due to the importance of lower redshift information for constraining IA and the degeneracy between these parameters. We have found that, promisingly, the constraints on these parameters do not worsen with this method (figure 8). In fact, we find that applying this method reduces the bias in IA parameters of the NLA model to be comparable with a reduction in the outlier fraction. Future applications to more complex or realistic cases will require careful tuning of the weightings of the likelihood, potentially informed by both posterior predictive checks and simulations.

While we have shown that the poorly fitting components of the model are easily visually identified in the case of noiseless simulated data, we have also demonstrated that this is non-trivial when considering more realistic noisy data (figure 6). We have suggested methods to nevertheless use PPCs and related tools to support the composite likelihood method in this case, however, further work to build up an infrastructure for using composite likelihoods in the case of realistic noisy data would be a valuable step towards applying this method in a real analysis.

We have shown that the composite likelihood approach has promise for mitigating the impact of an outlier population. We further suggest that it has the potential to be a more generally useful tool to mitigate model misspecification. There is scope for further exploration of the applicability of this method and its utility to mitigate issues related to other systematic effects, including e.g. baryonic feedback or other non-linear structure formation processes.

We provide now a clear, point-by-point summary of the main take-aways of this work:

- Model misspecification for the outlier fraction has the potential to significantly bias cosmological and dark energy parameter constraints, with shifts approaching 2σ .
- Even when the data is generated from a perfectly known model, miscentred priors can introduce parameter biases up to 1σ .
- The outlier fraction induces a particularly large bias in the IA parameters, which can reach up to 8σ for reasonable outlier fractions.
- We demonstrated that posterior predictive checks can effectively identify poorly modelled components of the data vector, specifically those stemming from an unmodelled outlier population.
- Using a composite likelihood approach to down-weight affected components — guided by posterior predictive diagnostics — reduced parameter biases to below 1σ without degrading constraints on intrinsic alignment and photometric redshift shifts, despite the down-weighted observables being important for constraining them.

These findings highlight the importance of identifying and mitigating model misspecification in cosmological analyses from realistically complex redshift distributions. While our analysis focused on an optimistic and simplified scenario, the techniques developed here can inform future studies dealing with more realistic data. In such cases, tuning the composite likelihood weights—potentially using a combination of posterior predictive checks and simulation-based diagnostics—will be essential. Additionally, further work is needed to assess the robustness of these methods under the presence of noise, which will be a key step toward applying them to real data. It would also be interesting to extend this work to assess the impact of IA and photo- z systematics on beyond 2-point and map-based statistics. Future work could investigate the potential of such techniques to identify systematic effects and break degeneracies, building on e.g. [58–62].

A Full parameter constraints

We present the cosmological, dark energy, and IA parameter constraints for each scenario considered in this work in table 3. We also show these in the form of violin plots in figure 10.

Acknowledgments

CMM acknowledges the U.K. Science and Technology Facilities Council (STFC) for support from grant No. ST/W006790/1. CMM was supported by an STFC-funded Long-Term Attachment ST/W006790/1. MMR acknowledges support by the IES\R1\241442 Royal

Scenario	1	2	3	4	5	6
Ω_M	$0.317^{+0.013}_{-0.015}$	$0.329^{+0.015}_{-0.016}$	$0.330^{+0.014}_{-0.015}$	$0.347^{+0.018}_{-0.016}$	$0.344^{+0.016}_{-0.016}$	$0.331^{+0.014}_{-0.018}$
σ_8	$0.828^{+0.015}_{-0.013}$	$0.816^{+0.016}_{-0.014}$	$0.0816^{+0.015}_{-0.014}$	$0.801^{+0.015}_{-0.015}$	$0.803^{+0.015}_{-0.014}$	$0.814^{+0.018}_{-0.013}$
w_0	$-0.979^{+0.155}_{-0.156}$	$-0.922^{+0.175}_{-0.176}$	$-0.882^{+0.116}_{-0.175}$	$-0.683^{+0.179}_{-0.188}$	$-0.693^{+0.172}_{-0.165}$	$-0.857^{+0.176}_{-0.207}$
w_a	$-0.041^{+0.461}_{-0.506}$	$-0.122^{+0.547}_{-0.514}$	$-0.249^{+0.518}_{-0.509}$	$-0.837^{+0.569}_{-0.519}$	$-0.786^{+0.500}_{-0.521}$	$-0.331^{+0.627}_{-0.572}$
a	$0.434^{+0.049}_{-0.046}$	$0.427^{+0.044}_{-0.048}$	$0.399^{+0.041}_{-0.044}$	$0.389^{+0.041}_{-0.046}$	$0.274^{+0.040}_{-0.041}$	$0.401^{+0.047}_{-0.043}$
η	$-0.717^{+0.405}_{-0.464}$	$-0.971^{+0.424}_{-0.446}$	$0.669^{+0.438}_{-0.523}$	$0.789^{+0.632}_{-0.640}$	$3.009^{+0.430}_{-0.470}$	$0.740^{+0.558}_{-0.608}$
Δz_s^0	$-0.004^{+0.004}_{-0.004}$	$-0.001^{+0.004}_{-0.004}$	$-0.002^{+0.003}_{-0.003}$	$-0.004^{+0.004}_{-0.004}$	$0.008^{+0.002}_{-0.002}$	$0.003^{+0.004}_{-0.004}$
Δz_s^1	$0.004^{+0.002}_{-0.002}$	$0.002^{+0.002}_{-0.002}$	$-0.003^{+0.002}_{-0.002}$	$-0.005^{+0.002}_{-0.002}$	$-0.011^{+0.002}_{-0.002}$	$-0.004^{+0.003}_{-0.003}$

Table 3. Cosmological parameter constraints and their 1σ values for each scenario. Columns correspond to those in table 2.




Figure 9. The posterior probability distributions for cosmological parameters Ω_M, σ_8 (upper) and dark energy parameters (lower) for a range of scenarios described in table 2. $\Delta\theta$ indicates the difference between the maximum a posteriori and the fiducial value (table 1), where $\Delta\theta = 0$ is indicated with the dashed black line. Mean and 1σ and 2σ are indicated by the dark and light lines.

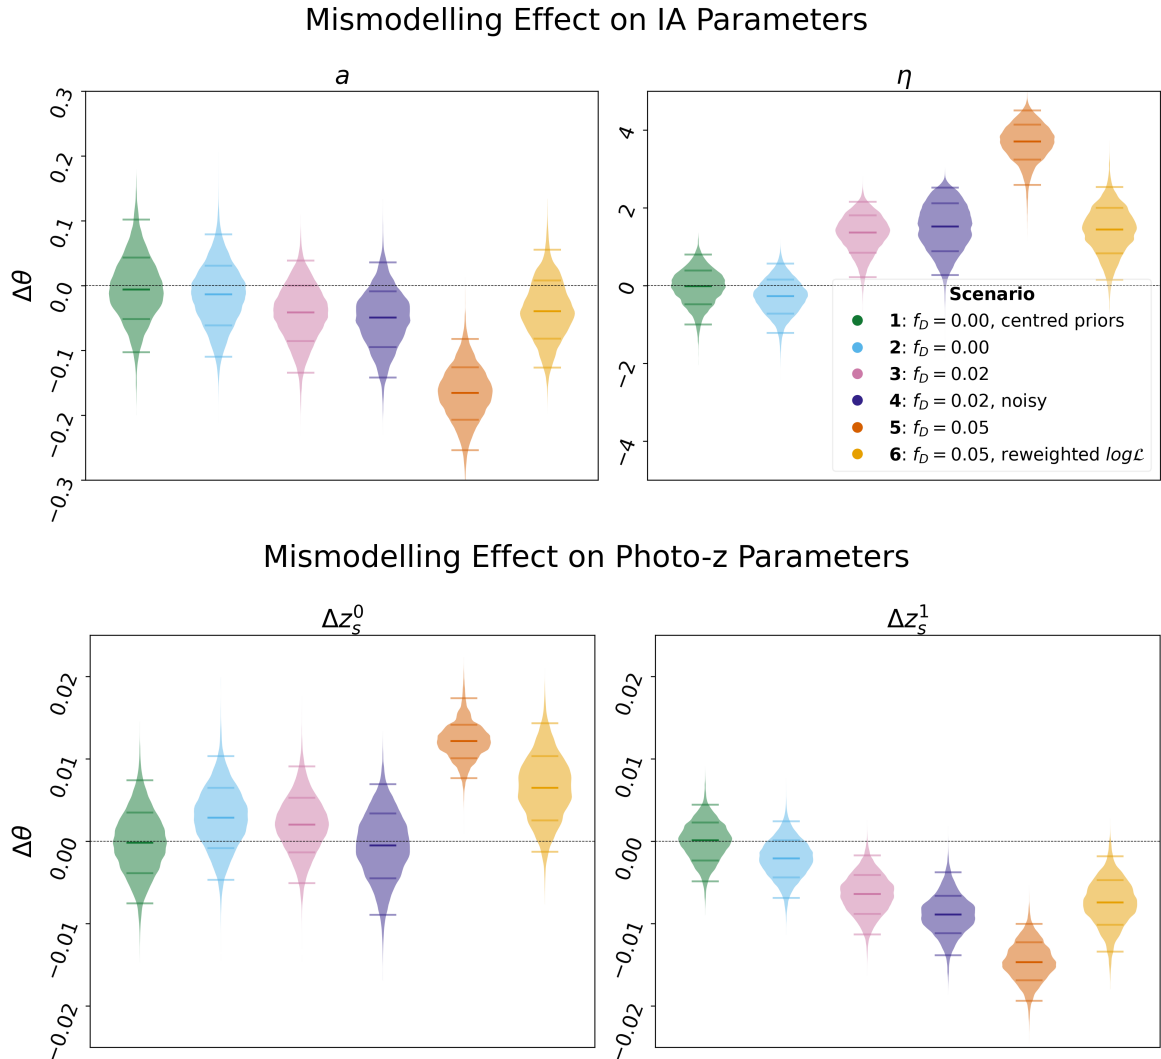


Figure 10. The posterior probability distributions for IA parameters and photo- z uncertainty parameters (lower) for a range of scenarios described in table 2. $\Delta\theta$ indicates the difference between the maximum a posteriori and the fiducial value (table 1), where $\Delta\theta = 0$ is indicated with the dashed black line. Mean and 1σ and 2σ are indicated by the dark and light lines.

Society International Exchanges 2024 grant. CU was supported by the STFC Astronomy Theory Consolidated Grant ST/W001020/1 from U.K. Research & Innovation and by the European Union (ERC StG, LSS_BeyondAverage, 101075919). CDL was supported by the Science and Technology Facilities Council (STFC) [grant No. UKRI1172]. This work used the DiRAC Data Intensive service (DIaL2 / DIaL3) at the University of Leicester, managed by the University of Leicester Research Computing Service on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC service at Leicester was funded by BEIS, UKRI and STFC capital funding and STFC operations grants. DiRAC is part of the UKRI Digital Research Infrastructure. The figures in this work were generated using Matplotlib [63] making use of NumPy [64] and SciPy [65].

Data Availability Statement. This article has associated data in a data repository. Visualisation of the full parameter posterior distributions for all scenarios considered in this work are provided at https://github.com/carolymmm/3x2pt_contours.

References

- [1] C. M. Mill, C. D. Leonard, M. M. Rau, C. Uhlemann, S. Joudaki, *Data Availability Statement Repository*, https://github.com/carolymmm/3x2pt_contours.
- [2] LSST collaboration, *LSST: from Science Drivers to Reference Design and Anticipated Data Products*, *Astrophys. J.* **873** (2019) 111 [[arXiv:0805.2366](#)] [[INSPIRE](#)].
- [3] EUCLID collaboration, *Euclid Definition Study Report*, [arXiv:1110.3193](#) [[INSPIRE](#)].
- [4] EUCLID collaboration, *Euclid. I. Overview of the Euclid mission*, *Astron. Astrophys.* **697** (2025) A1 [[arXiv:2405.13491](#)] [[INSPIRE](#)].
- [5] D. Spergel et al., *Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report*, [arXiv:1503.03757](#) [[INSPIRE](#)].
- [6] COSMOVERSE NETWORK collaboration, *The CosmoVerse White Paper: Addressing observational tensions in cosmology with systematics and fundamental physics*, *Phys. Dark Univ.* **49** (2025) 101965 [[arXiv:2504.01669](#)] [[INSPIRE](#)].
- [7] LSST DARK ENERGY SCIENCE collaboration, *Large Synoptic Survey Telescope: Dark Energy Science Collaboration*, [arXiv:1211.0310](#) [[DOI:10.2172/1156445](#)] [[INSPIRE](#)].
- [8] EROSITA collaboration, *eROSITA Science Book: Mapping the Structure of the Energetic Universe*, [arXiv:1209.3114](#) [[INSPIRE](#)].
- [9] M. Bartelmann and P. Schneider, *Weak gravitational lensing*, *Phys. Rept.* **340** (2001) 291 [[astro-ph/9912508](#)] [[INSPIRE](#)].
- [10] R. Mandelbaum, *Weak lensing for precision cosmology*, *Ann. Rev. Astron. Astrophys.* **56** (2018) 393 [[arXiv:1710.03235](#)] [[INSPIRE](#)].
- [11] D. Huterer, *Weak lensing, dark matter and dark energy*, *Gen. Rel. Grav.* **42** (2010) 2177 [[arXiv:1001.1758](#)] [[INSPIRE](#)].
- [12] DES collaboration, *Dark Energy Survey Year 3 results: Constraints on extensions to Λ CDM with weak lensing and galaxy clustering*, *Phys. Rev. D* **107** (2023) 083504 [[arXiv:2207.05766](#)] [[INSPIRE](#)].
- [13] S. Joudaki et al., *KiDS-450: Testing extensions to the standard cosmological model*, *Mon. Not. Roy. Astron. Soc.* **471** (2017) 1259 [[arXiv:1610.04606](#)] [[INSPIRE](#)].
- [14] D. Huterer, M. Takada, G. Bernstein and B. Jain, *Systematic errors in future weak lensing surveys: Requirements and prospects for self-calibration*, *Mon. Not. Roy. Astron. Soc.* **366** (2006) 101 [[astro-ph/0506030](#)] [[INSPIRE](#)].
- [15] LSST DARK ENERGY SCIENCE collaboration, *Forecasting the Impact of Source Galaxy Photometric Redshift Uncertainties on the LSST 3×2 pt Analysis*, [arXiv:2507.01374](#) [[INSPIRE](#)].
- [16] S. Sugiyama et al., *Hyper Suprime-Cam Year 3 results: Cosmology from galaxy clustering and weak lensing with HSC and SDSS using the minimal bias model*, *Phys. Rev. D* **108** (2023) 123521 [[arXiv:2304.00705](#)] [[INSPIRE](#)].

- [17] DES collaboration, *Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to modeling uncertainty*, *Phys. Rev. D* **105** (2022) 023515 [[arXiv:2105.13544](#)] [[INSPIRE](#)].
- [18] H. Hildebrandt et al., *KiDS-450: Cosmological parameter constraints from tomographic weak gravitational lensing*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 1454 [[arXiv:1606.05338](#)] [[INSPIRE](#)].
- [19] M. Salvato, O. Ilbert and B. Hoyle, *The many flavours of photometric redshifts*, [arXiv:1805.12574](#).
- [20] J.A. Newman and D. Gruen, *Photometric Redshifts for Next-Generation Surveys*, *Ann. Rev. Astron. Astrophys.* **60** (2022) 363 [[arXiv:2206.13633](#)] [[INSPIRE](#)].
- [21] H. Hildebrandt et al., *KiDS+VIKING-450: Cosmic shear tomography with optical and infrared data*, *Astron. Astrophys.* **633** (2020) A69 [[arXiv:1812.06076](#)] [[INSPIRE](#)].
- [22] DES collaboration, *Phenotypic redshifts with self-organizing maps: A novel method to characterize redshift distributions of source galaxies for weak lensing*, *Mon. Not. Roy. Astron. Soc.* **489** (2019) 820 [[arXiv:1901.05005](#)] [[INSPIRE](#)].
- [23] S. Jouvel et al., *Designing Future Dark Energy Space Missions: II. Photometric Redshift of Space Weak Lensing Optimized Survey*, *Astron. Astrophys.* **532** (2011) A25 [[arXiv:1003.4294](#)] [[INSPIRE](#)].
- [24] A.P. Hearin, A.R. Zentner, Z. Ma and D. Huterer, *A General Study of the Influence of Catastrophic Photometric Redshift Errors on Cosmology with Cosmic Shear Tomography*, *Astrophys. J.* **720** (2010) 1351 [[arXiv:1002.3383](#)] [[INSPIRE](#)].
- [25] G. Bernstein and D. Huterer, *Catastrophic photometric redshift errors: weak lensing survey requirements*, *Mon. Not. Roy. Astron. Soc.* **401** (2010) 1399 [[arXiv:0902.2782](#)] [[INSPIRE](#)].
- [26] D. Masters et al., *Mapping the Galaxy Color–redshift Relation: Optimal Photometric Redshift Calibration Strategies for Cosmology Surveys*, *Astrophys. J.* **813** (2015) 53 [[arXiv:1509.03318](#)] [[INSPIRE](#)].
- [27] M.A. Troxel and M. Ishak, *The Intrinsic Alignment of Galaxies and its Impact on Weak Gravitational Lensing in an Era of Precision Cosmology*, *Phys. Rept.* **558** (2014) 1 [[arXiv:1407.6990](#)] [[INSPIRE](#)].
- [28] B. Joachimi et al., *Galaxy alignments: An overview*, *Space Sci. Rev.* **193** (2015) 1 [[arXiv:1504.05456](#)] [[INSPIRE](#)].
- [29] C. Lamman et al., *The IA Guide: A Breakdown of Intrinsic Alignment Formalisms*, [arXiv:2309.08605](#) [[DOI:10.21105/astro.2309.08605](#)] [[INSPIRE](#)].
- [30] E. Krause, T. Eifler and J. Blazek, *The impact of intrinsic alignment on current and future cosmic shear surveys*, *Mon. Not. Roy. Astron. Soc.* **456** (2016) 207 [[arXiv:1506.08730](#)] [[INSPIRE](#)].
- [31] D. Kirk et al., *Galaxy alignments: Observations and impact on cosmology*, *Space Sci. Rev.* **193** (2015) 139 [[arXiv:1504.05465](#)] [[INSPIRE](#)].
- [32] J. Yao, M. Ishak, W. Lin and M.A. Troxel, *Effects of Self-Calibration of Intrinsic Alignment on Cosmological Parameter Constraints from Future Cosmic Shear Surveys*, *JCAP* **10** (2017) 056 [[arXiv:1707.01072](#)] [[INSPIRE](#)].

- [33] EUCLID collaboration, *Euclid preparation. XXXI. The effect of the variations in photometric passbands on photometric-redshift accuracy*, *Astron. Astrophys.* **681** (2024) A66 [[arXiv:2310.14716](#)] [[INSPIRE](#)].
- [34] LSST DARK ENERGY SCIENCE collaboration, *Impact of large-scale structure systematics on cosmological parameter estimation*, *JCAP* **03** (2025) 064 [[arXiv:2409.14265](#)] [[INSPIRE](#)].
- [35] A.H. Wright et al., *KiDS-Legacy: Cosmological constraints from cosmic shear with the complete Kilo-Degree Survey*, *Astron. Astrophys.* **703** (2025) A158 [[arXiv:2503.19441](#)] [[INSPIRE](#)].
- [36] A.H. Wright et al., *KiDS+VIKING-450: Improved cosmological parameter constraints from redshift calibration with self-organising maps*, *Astron. Astrophys.* **640** (2020) L14 [[arXiv:2005.04207](#)] [[INSPIRE](#)].
- [37] B. Stölzner et al., *Self-calibration and robust propagation of photometric redshift distribution uncertainties in weak gravitational lensing*, *Astron. Astrophys.* **650** (2021) A148 [[arXiv:2012.07707](#)] [[INSPIRE](#)].
- [38] C.D. Leonard, M.M. Rau and R. Mandelbaum, *Photometric redshifts and intrinsic alignments: Degeneracies and biases in the 3×2 pt analysis*, *Phys. Rev. D* **109** (2024) 083528 [[arXiv:2401.06060](#)] [[INSPIRE](#)].
- [39] S. Fischbacher, T. Kacprzak, J. Blazek and A. Refregier, *Redshift requirements for cosmic shear with intrinsic alignment*, *JCAP* **01** (2023) 033 [[arXiv:2207.01627](#)] [[INSPIRE](#)].
- [40] DES collaboration, *Dark Energy Survey Year 1 Results: Redshift distributions of the weak lensing source galaxies*, *Mon. Not. Roy. Astron. Soc.* **478** (2018) 592 [[arXiv:1708.01532](#)] [[INSPIRE](#)].
- [41] M.L. Graham et al., *Photometric Redshifts with the LSST: Evaluating Survey Observing Strategies*, *Astron. J.* **155** (2017) 1 [[arXiv:1706.09507](#)] [[INSPIRE](#)].
- [42] LSST DARK ENERGY SCIENCE collaboration, *The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document*, [arXiv:1809.01669](#) [[DOI:10.2172/1471560](#)] [[INSPIRE](#)].
- [43] LSST DARK ENERGY SCIENCE collaboration, *Core Cosmology Library: Precision Cosmological Predictions for LSST*, *Astrophys. J. Suppl.* **242** (2019) 2 [[arXiv:1812.05995](#)] [[INSPIRE](#)].
- [44] J.E. McEwen, X. Fang, C.M. Hirata and J.A. Blazek, *FAST-PT: a novel algorithm to calculate convolution integrals in cosmological perturbation theory*, *JCAP* **09** (2016) 015 [[arXiv:1603.04826](#)] [[INSPIRE](#)].
- [45] X. Fang, J.A. Blazek, J.E. McEwen and C.M. Hirata, *FAST-PT II: an algorithm to calculate convolution integrals of general tensor quantities in cosmological perturbation theory*, *JCAP* **02** (2017) 030 [[arXiv:1609.05978](#)] [[INSPIRE](#)].
- [46] S. Bridle and L. King, *Dark energy constraints from cosmic shear power spectra: impact of intrinsic alignments on photometric redshift requirements*, *New J. Phys.* **9** (2007) 444 [[arXiv:0705.0166](#)] [[INSPIRE](#)].
- [47] J. Blazek, N. MacCrann, M.A. Troxel and X. Fang, *Beyond linear galaxy alignments*, *Phys. Rev. D* **100** (2019) 103506 [[arXiv:1708.09247](#)] [[INSPIRE](#)].
- [48] M.C. Fortuna et al., *The halo model as a versatile tool to predict intrinsic alignments*, *Mon. Not. Roy. Astron. Soc.* **501** (2021) 2983 [[arXiv:2003.02700](#)] [[INSPIRE](#)].
- [49] Z. Vlah, N.E. Chisari and F. Schmidt, *An EFT description of galaxy intrinsic alignments*, *JCAP* **01** (2020) 025 [[arXiv:1910.08085](#)] [[INSPIRE](#)].
- [50] J. Goodman and J. Weare, *Ensemble samplers with affine invariance*, *Commun. Appl. Math. Comput. Sc.* **5** (2010) 65 [[INSPIRE](#)].

- [51] A. Gelman, X.-L. Meng and H. Stern, *Posterior predictive assessment of model fitness via realized discrepancies*, *Stat. Sin.* **6** (1996) 733.
- [52] P.C. Mahalanobis, *On the generalized distance in statistics*, *Proc. Natl. Inst. Sci. India* **2** (1936) 49.
- [53] LSST DARK ENERGY SCIENCE collaboration, *A composite likelihood approach for inference under photometric redshift uncertainty*, *Mon. Not. Roy. Astron. Soc.* **509** (2021) 4886 [[arXiv:2101.01184](#)] [[INSPIRE](#)].
- [54] C. Varin, N. Reid and D. Firth, *An overview of composite likelihood methods*, *Stat. Sin.* **21** (2011) 0.
- [55] DES collaboration, *Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing*, *Phys. Rev. D* **98** (2018) 043526 [[arXiv:1708.01530](#)] [[INSPIRE](#)].
- [56] M.M. Rau et al., *Weak lensing tomographic redshift distribution inference for the Hyper Suprime-Cam Subaru Strategic Program three-year shape catalogue*, *Mon. Not. Roy. Astron. Soc.* **524** (2023) 5109 [[arXiv:2211.16516](#)] [[INSPIRE](#)].
- [57] A. Gelman and D.B. Rubin, *Inference from Iterative Simulation Using Multiple Sequences*, *Statist. Sci.* **7** (1992) 457 [[INSPIRE](#)].
- [58] O. Friedrich, L. Castiblanco, A. Halder and C. Uhlemann, *Bye binormal: analysing the joint PDF of galaxy density and weak lensing convergence*, [arXiv:2507.16957](#) [[INSPIRE](#)].
- [59] N. Porqueres et al., *Field-level inference of cosmic shear with intrinsic alignments and baryons*, [arXiv:2304.04785](#) [[INSPIRE](#)].
- [60] D. Zürcher et al., *Towards a full w CDM map-based analysis for weak lensing surveys*, *Mon. Not. Roy. Astron. Soc.* **525** (2023) 761 [[arXiv:2206.01450](#)] [[INSPIRE](#)].
- [61] A. Halder et al., *Beyond 3×2 -point cosmology: the integrated shear and galaxy 3-point correlation functions*, *JCAP* **10** (2023) 028 [[arXiv:2305.17132](#)] [[INSPIRE](#)].
- [62] LSST DARK ENERGY SCIENCE (LSST DESC) collaboration, *Forecasting the power of higher order weak-lensing statistics with automatically differentiable simulations*, *Astron. Astrophys.* **679** (2023) A61 [[arXiv:2305.07531](#)] [[INSPIRE](#)].
- [63] J.D. Hunter, *Matplotlib: A 2D Graphics Environment*, *Comput. Sci. Eng.* **9** (2007) 90 [[INSPIRE](#)].
- [64] C.R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357 [[arXiv:2006.10256](#)] [[INSPIRE](#)].
- [65] P. Virtanen et al., *SciPy 1.0—Fundamental Algorithms for Scientific Computing in Python*, *Nature Meth.* **17** (2020) 261 [[arXiv:1907.10121](#)] [[INSPIRE](#)].