Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **522,** 5037–5048 (2023) https://doi.org/10.1093/mnras/stad1192
Advance Access publication 2023 April 21

Analytical marginalization over photometric redshift uncertainties in cosmic shear analyses

Jaime Ruiz-Zapatero [1]★ Boryana Hadzhiyska,[2,3] David Alonso,[1] Pedro G. Ferreira,[1] Carlos García-García[1] and Arrykrishna Mootoovaloo[1]

[1]*Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*
[2]*Miller Institute for Basic Research in Science, University of California, Berkeley, CA 94720, USA*
[3]*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

Accepted 2023 April 18. Received 2023 April 11; in original form 2023 January 27

ABSTRACT

As the statistical power of imaging surveys grows, it is crucial to account for all systematic uncertainties. This is normally done by constructing a model of these uncertainties and then marginalizing over the additional model parameters. The resulting high dimensionality of the total parameter spaces makes inferring the cosmological parameters significantly more costly using traditional Monte Carlo sampling methods. A particularly relevant example is the redshift distribution, $p(z)$, of the source samples, which may require tens of parameters to describe fully. However, relatively tight priors can be usually placed on these parameters through calibration of the associated systematics. In this paper, we show, quantitatively, that a linearization of the theoretical prediction with respect to these calibrated systematic parameters allows us to analytically marginalize over these extra parameters, leading to a factor of $\sim$30 reduction in the time needed for parameter inference, while accurately recovering the same posterior distributions for the cosmological parameters that would be obtained through a full numerical marginalization over 160 $p(z)$ parameters. We demonstrate that this is feasible not only with current data and current achievable calibration priors but also for future Stage-IV data sets.

Key words: gravitational lensing: weak – methods: data analysis – large-scale structure of Universe.

1 INTRODUCTION

In recent years, unprecedentedly precise observations in cosmology have uncovered a number of tensions between data sets that may constitute both tantalizing hints of new physics and a manifestation of a lack of control over theoretical and observational systematics (Heymans et al. 2021; Riess et al. 2022).

At its simplest, the current cosmological paradigm, the Λ (denoting the cosmological constant) cold dark matter (ΛCDM) model, can be described by only five parameters: $\Omega_{\rm m}$, $\Omega_{\rm b}$, $A_{\rm s}$, $n_{\rm s}$, and h [see e.g. Scott (2018) for a detailed review]. However, in order to relate the theoretical predictions of this model to actual physical observables, it is necessary to extend it. Phenomenological models that describe the astrophysical systems that form the basis of our observations, as well as observational sources of systematic uncertainty, are then appended to the core ΛCDM model. In the presence of large statistical uncertainties, these models may consist of simple relationships in terms of a handful of parameters. However, more precise data require an equally precise characterization of these relationships, which leads to an increase in the complexity of the model. Thus, the number of parameters associated with these bridging models, colloquially referred to as 'nuisance' parameters, has steadily grown over the years.

The term 'nuisance' is accurate when describing these parameters. Not only are they generally uninteresting by comparison with the fundamental cosmological parameters we aim to constrain, but the increase in parameter dimensionality of the model makes exploring their posterior distribution significantly more computationally costly as well. Standard Markov chain Monte Carlo (MCMC) and other rejection-based sampling methods (Metropolis et al. 1953; Foreman-Mackey et al. 2013; Alsing & Handley 2021, among others) suffer from the so-called curse of dimensionality, whereby the acceptance rate of new samples decreases sharply with the number of parameters (exponentially in the worst cases).

Nuisance parameters can be divided into two groups based on their prior distributions: calibrated and non-calibrated parameters. The non-calibrated parameters can only be constrained by the data and, as such, typically have largely non-constraining priors. On the other hand, we can place tighter priors on the calibrated parameters, either by accurately characterizing the instrument measurements or by using independent external observations. In the case of cosmic shear analyses, the impact of galaxy intrinsic alignments (IAs; Hirata & Seljak 2004) is often a non-calibrated systematic. On the calibrated side, the two best examples are multiplicative shape measurement systematics and the uncertainties in the redshift distribution of the target source galaxies (Hoyle et al. 2018; Sánchez & Bernstein 2019; Hildebrandt et al. 2020a; Stölzner et al. 2021; Zhang et al. 2023). It is important to stress that whether a parameter is calibrated or not is not inherent to the parameter but to whether a sufficiently tight prior

★ E-mail: jaime.ruiz-zapatero@physics.ox.ac.uk

on it can be placed (e.g. from external data). For example, IAs can be calibrated as shown in Yao et al. (2020), among others.

Of these calibrated systematics, the dominant source of uncertainty in photometric surveys is the accuracy of redshift distributions, which are known to strongly affect the accuracy of cosmological constraints. The vital quantity to determine is the redshift distribution of each tomographic sample of galaxies, $p(z)$. The fact that the uncertainties in $p(z)$ can be calibrated with external spectroscopic data, e.g. via direct calibration (DIR; Lima et al. 2008; Wright et al. 2020), clustering redshifts (Schneider et al. 2006; Newman 2008; Matthews & Newman 2010; Schmidt et al. 2013), and shear ratios (Prat et al. 2018; Sánchez et al. 2022), enables us to place relatively strong priors on the redshift distribution, which in turn makes it possible to use approximate methods to efficiently marginalize over these uncertainties.

Analytical marginalization schemes for photometric redshift uncertainties have already been proposed in the literature. In Stölzner et al. (2021), an analytical marginalization scheme for photometric redshift uncertainties was proposed based on Gaussian mixture models and applied to the analysis of KV450 data (Hildebrandt et al. 2020b). In Zhang et al. (2023), a resampling approach to marginalize over these uncertainties was proposed and applied to the analysis of Subaru Hyper Suprime-Cam (HSC) data. Similarly, Reischke (2023) proposed an approach based on functional derivatives and applied it to Euclid-like (Spergel et al. 2015) and KV450 (Hildebrandt et al. 2020b) data. Alternatively, Cordero et al. (2022) proposed an approach using hyperrank, a method based on ranking discrete samples from the space of all possible redshift distributions of discrete realizations, and applied it to DESY3 data (Sevilla-Noarbe et al. 2021). Here, we will explore the method initially proposed in Hadzhiyska et al. (2020), further exploited in García-García et al. (2023), and recently characterized in the context of the Laplace approximation in Hadzhiyska et al. (2023). The method is based on linearizing the dependence of the theoretical prediction with respect to the parameters defining the redshift distribution around their calibration priors. This then allows one to analytically marginalize over these parameters by modifying the covariance matrix of the data, effectively assigning higher variance (as allowed by the calibration prior) to the data modes most sensitive to variations in the $p(z)$.

The goal of this paper is to exhaustively validate this approximate marginalization scheme in the context of cosmic shear analyses. We will do so by proving that we are able to obtain the same constraints on cosmological parameters using this scheme, as well as employing brute-force methods that sample the full parameter space exactly. We will show this both for simple parametrizations of the $p(z)$ uncertainties, in terms of shifts to the mean of the distribution, and using completely general 'non-parametric' models that treat the amplitude of the $p(z)$ in narrowly spaced intervals of z as calibrated variables, leading to a model with more than $\sim$100 nuisance parameters. In order to numerically marginalize over such large parameter spaces, we develop an autodifferentiable code to obtain theoretical predictions for the cosmic shear observables. This allows us to employ gradient-based sampling algorithms, such as Hamiltonian Monte Carlo (HMC), to beat the aforementioned curse of dimensionality. Finally, we will show that the method is valid not only for current data, but also for futuristic Stage-IV surveys, where photometric redshift uncertainties will likely make up a large fraction of the total error budget. Interestingly, our analysis will show that, in the context of cosmic shear data, relatively inexpensive parametrizations of photometric redshift uncertainties based on one free parameter per redshift bin [e.g. mean shifts, or ranked discrete realizations; Cordero et al. 2022] return effectively the same posterior distribution on cosmological parameters as the most general non-parametric models.

This paper is structured as follows. In Section 2, we describe the methods used in this work, including the theory behind weak lensing observables, the calibration of redshift distributions, and the mathematics of analytical marginalization via first-order expansion. Section 3 presents the Dark Energy Survey data used to produce realistic source redshift distributions and their associated uncertainties, as well as the models used to simulate future data sets. In Section 4, we describe the likelihood used to analyse these data, as well as the different parametrizations used to describe $p(z)$ uncertainties. Section 5 presents our results, quantifying the performance of analytical marginalization methods. Finally, we present our conclusions in Section 6.

2 METHODS

2.1 Cosmic shear power spectra

It is now commonplace to carry out the analysis of galaxy weak lensing data tomographically. The full sample is split into redshift bins and the two-point correlation functions of all pairs of bins are measured and compared with their theoretical prediction. Let $\gamma_\alpha(\hat{\boldsymbol{n}})$ be a map of the spin-2 lensing shear field inferred from the sources in the αth redshift bin. Its relation with the three-dimensional matter overdensity $\delta_m(\boldsymbol{x})$ is (Bartelmann & Schneider 2001; Krause et al. 2017)

$$\gamma_\alpha(\hat{\boldsymbol{n}}) = \int_0^{\chi_H} \mathrm{d}\chi\, q_\alpha(\chi)\, \left[-\chi^{-2}\eth\eth\nabla^{-2}\delta_m(\chi\hat{\boldsymbol{n}}, z)\right], \tag{1}$$

where $\hat{\mathbf{n}}$ is the sky direction, χ is the comoving radial distance at redshift z, χ_H is the distance to the horizon, $q^\alpha(\chi)$ is the weak lensing radial kernel, and $\eth$ is the spin-raising differential operator, acting on a spin-s quantity as (Newman & Penrose 1966):

$$\eth_s f(\theta,\varphi) = -(\sin\theta)^s \left(\frac{\partial}{\partial\theta} + \frac{i}{\sin\theta}\frac{\partial}{\partial\varphi}\right)(\sin\theta)^{-s}{}_s f \tag{2}$$

and turning it into a spin-$(s+1)$ quantity. The weak lensing kernel is[1]

$$q_\alpha(\chi) \equiv \frac{3}{2}H_0^2\Omega_\mathrm{m}\frac{\chi}{a(\chi)}\int_{z(\chi)}^{\infty}\mathrm{d}z'\, p_\alpha(z')\frac{\chi(z')-\chi}{\chi(z')}, \tag{3}$$

where $H_0 \equiv H(z=0)$ is the Hubble expansion rate today, Ω_m is the current matter density parameter, and most importantly for our discussion, $p_\alpha(z)$ is the redshift distribution in bin α.

The angular power spectrum of the E-mode components of two maps α and β, $C_\ell^{\alpha\beta}$ can be related to the three-dimensional matter power spectrum $P(k, z)$ via

$$C_\ell^{\alpha\beta} = G_\ell^2\int\frac{\mathrm{d}\chi}{\chi^2}\, q_\alpha(\chi)\, q_\beta(\chi)\, P\left(k=\frac{\ell+1/2}{\chi}, z(\chi)\right), \tag{4}$$

where we have assumed the Limber approximation (Limber 1953; Afshordi, Loh & Strauss 2004), which is valid for the broad weak lensing kernels considered in this work. The scale-dependent lensing pre-factor,

$$G_\ell \equiv \sqrt{\frac{(\ell+2)!}{(\ell-2)!}}\frac{1}{(\ell+1/2)^2}, \tag{5}$$

[1]Note that this is only strictly valid in ΛCDM (Ferreira 2019).

accounts for the difference between angular and three-dimensional derivatives in equation (1) (i.e. $\chi^2\eth^2\nabla^{-2} \not\equiv 1$). This pre-factor leads to sub-per cent differences for $\ell > 11$ and can therefore be neglected on small scales (Kilbinger et al. 2017). In this work, we will use the `Halofit` fitting function of Smith et al. (2003) and Takahashi et al. (2012) to describe the matter power spectrum.

The IA of galaxies due to local interactions (gravitational or otherwise), is an important contaminant for cosmic shear data that must be taken into account (Brown et al. 2002). For simplicity however and since the focus of this work is the impact of the marginalization over redshift distribution uncertainties, we will ignore the contribution from IAs in this analysis. In the same spirit, we will also assume no baryonic effects on the non-linear power spectrum.

2.2 Redshift distribution uncertainties

The $p(z)$ can however be calibrated through various methods, e.g. weighted DIR with a sufficiently complete spectroscopic sample (Lima et al. 2008; Wright et al. 2020), clustering redshifts (Schneider et al. 2006; Newman 2008; Matthews & Newman 2010; Schmidt et al. 2013), and shear ratios (Kuijken et al. 2015; Prat et al. 2018; Sánchez et al. 2022). This typically leads to relatively tight priors on the $p(z)$, but the residual uncertainties in this prior must be propagated into the final parameter constraints.

To characterize these uncertainties, we will make use of two different methods, which encompass the range of model complexity which we may reasonably expect from current and future data.

(i) **Method 1: z shifts.** Most cosmic shear analyses to date (Miyazaki et al. 2012; Troxel et al. 2018; Hildebrandt et al. 2020b; Heymans et al. 2021; Secco et al. 2022, among others) have summarized the uncertainty in the calibrated $p_\alpha(z)$ into a single parameter Δz^α that shifts the mean of the redshift distribution. That is, let $\hat{p}_\alpha(z)$ be the best guess redshift distribution. The true redshift distribution is then

$$p_\alpha(z) = \hat{p}_\alpha(z + \Delta z^\alpha). \tag{6}$$

A prior on Δz^α can be derived using the calibration methods listed earlier. We will refer to this method as *parametric*.

This simple model turns out to be relatively well suited to describe the impact of $p(z)$ uncertainties in the case of cosmic shear data (Ma, Hu & Huterer 2006) even from Stage-IV surveys (The LSST Dark Energy Science Collaboration 2018). Since weak lensing is a radially cumulative effect, the amplitude of the weak lensing kernel (equation 3) is mostly sensitive to the mean redshift of the sample, and thus, much of the effect on cosmic shear observables is well described by this parameter (Bonnett et al. 2016).

Other modes of $p(z)$ uncertainty, such as the distribution width, may be more relevant for galaxy clustering observables, or for the intrinsic alignment contribution to cosmic shear. Near-future cosmic shear samples may indeed require a more sophisticated description of the $p(z)$ uncertainty, and thus, we turn to a more general method.

(ii) **Method 2: $p(z)$ bin heights.** Most $p(z)$ calibration methods (e.g. DIR or clustering redshifts) will produce a binned measurement of the $p(z)$ with deterministic redshift bin ranges, and uncertain bin heights. The most general method to propagate these uncertainties is therefore to treat each bin height $p_i \equiv p(z_i)$ as a free parameter in the model, with a prior given by the calibration uncertainties. The latter may be in the form of individual 1σ errors for each bin height, if the uncertainties are approximately uncorrelated, or a full covariance matrix covering all bin heights.

The resulting parametrization thus sidesteps any attempt at summarizng the uncertainty into effective parameters, and thus, we will refer to this method as *non-parametric*. The method therefore fully propagates all calibration uncertainties into the final constraints with minimal approximations.

The key practical difference between both methods, in the context of error propagation, is the additional complexity they incur. The parametric approach (Method 1) introduces one free parameter per redshift bin. For $O(5)$ bins, this is already enough to significantly impact the performance of standard MCMC algorithms. In turn, the non-parametric approach (Method 2) introduces tens or hundreds of parameters per redshift bin, and one must resort to advanced sampling methods in order to fully explore the resulting model without assumptions.

2.3 Linearization and analytical marginalization

Let $\boldsymbol{\Omega}$ be the set of non-calibrated parameters of a model (in our case this is the set of cosmological and non-calibrated nuisance parameters) and $\boldsymbol{\nu}$ the set of calibrated parameter such that the total set of parameters is given by $\boldsymbol{\theta} = \boldsymbol{\Omega} \cup \boldsymbol{\nu}$. Now, consider the general case of a Gaussian posterior distribution of the form

$$\begin{aligned} -2\log P(\boldsymbol{\Omega}, \boldsymbol{\nu}|\boldsymbol{d}) = (\boldsymbol{d}-\boldsymbol{t})^T C^{-1}(\boldsymbol{d}-\boldsymbol{t}) + (\boldsymbol{\nu}-\bar{\boldsymbol{\nu}})^T P^{-1}(\boldsymbol{\nu}-\bar{\boldsymbol{\nu}}) \\ -2\log P(\boldsymbol{\Omega}) + \text{const.}, \end{aligned} \tag{7}$$

where $\boldsymbol{d}$ is the data. We assume a Gaussian calibration prior with mean $\bar{\boldsymbol{\nu}}$ and covariance P, while $P(\boldsymbol{\Omega})$ is the prior on $\boldsymbol{\Omega}$ (which is, as per our assumption, broad). $\mathbf{t}(\boldsymbol{\Omega}, \boldsymbol{\nu})$ is the theoretical prediction for the data $\boldsymbol{d}$ which implicitly depends on both calibrated and non-calibrated parameters. C is the covariance matrix of $\boldsymbol{d}$, which is parameter-independent.

Assuming a tight prior on $\boldsymbol{\nu}$, we start by expanding the theory prediction around $\bar{\boldsymbol{\nu}}$

$$\boldsymbol{t} \simeq \bar{\boldsymbol{t}} + T(\boldsymbol{\nu}-\bar{\boldsymbol{\nu}}), \quad \text{where} \quad \bar{\boldsymbol{t}} \equiv \boldsymbol{t}(\boldsymbol{\Omega}, \bar{\boldsymbol{\nu}}), \quad T \equiv \left.\frac{\mathrm{d}\boldsymbol{t}}{\mathrm{d}\boldsymbol{\nu}}\right|_{\boldsymbol{\nu}=\bar{\boldsymbol{\nu}}}. \tag{8}$$

Substituting this approximation in equation (7), the posterior becomes Gaussian in $\boldsymbol{\nu}$, and thus, the calibrated parameters can be marginalized analytically. As shown in Hadzhiyska et al. (2020), the resulting marginalized posterior is

$$\begin{aligned} -2\log P(\boldsymbol{\Omega}|\boldsymbol{d}) \simeq \quad & (\boldsymbol{d}-\bar{\boldsymbol{t}})^T \tilde{C}^{-1}(\boldsymbol{d}-\bar{\boldsymbol{t}}) - 2\log P(\boldsymbol{\Omega}) \\ & + \log\left[\det\left(T^T C^{-1} T + P^{-1}\right)\right] + \text{const.}, \end{aligned} \tag{9}$$

where the modified covariance is

$$\tilde{C} \equiv C + TPT^T. \tag{10}$$

Note that, strictly speaking, both the modified covariance and the term in the second line of equation (9) depend on $\boldsymbol{\Omega}$, which would in principle complicate the evaluation of the likelihood. In practice, this parameter dependence can be neglected such that the value of $\boldsymbol{\Omega}$ at which these terms are evaluated can be fixed during exploration of the posterior. However, fixing $\boldsymbol{\Omega}$ at values with a bad fit to the data will result in a mischaracterization of the response of the theory vector to the nuisance parameters leading to inaccurate marginalized posteriors. Ideally, $\boldsymbol{\Omega}$ is fixed to its maximum a posteriori (MAP) value. However, as shown in Hadzhiyska et al. (2020) and in preliminary results, no appreciable differences are found in the marginalized posteriors for $\boldsymbol{\Omega}$ within 2σ of the MAP. Note that the size of the 2σ region will depend on how constraining the data are.

This result is intuitively simple to understand if we think of T as the response of the data to variations in the nuisance parameters. After marginalizing over the calibrated parameters, the resulting distribution is a multivariate Gaussian, where the data covariance

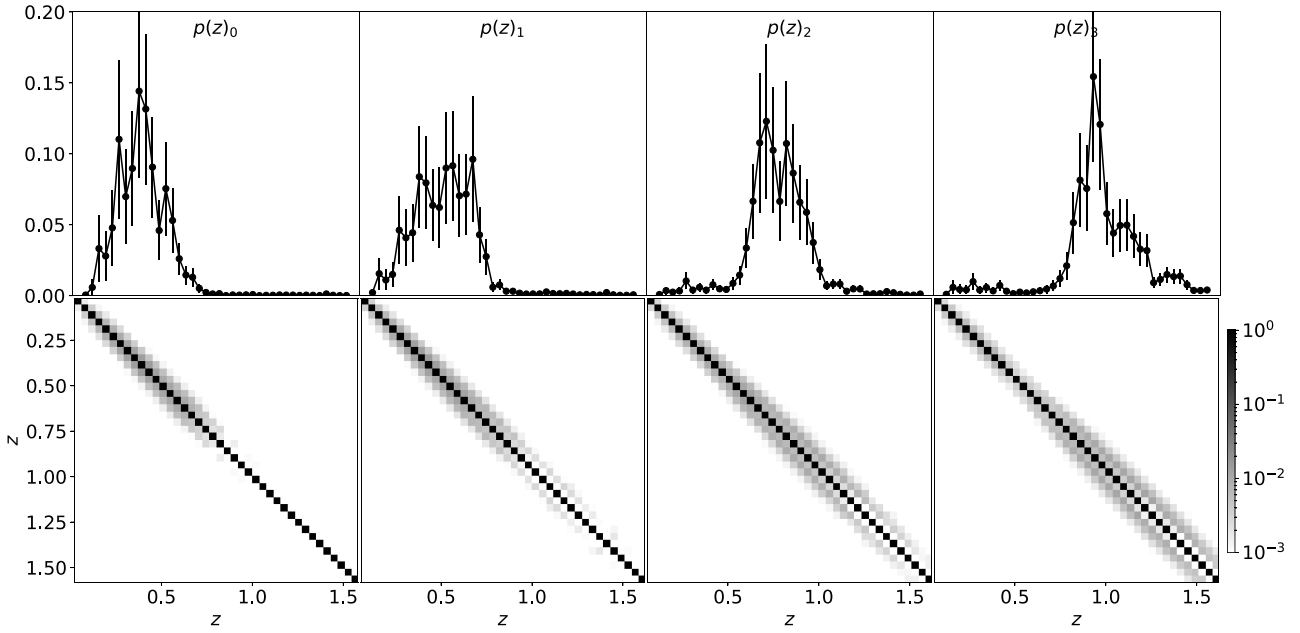

Figure 1. Top row: normalized galaxies' redshift distributions for each of the four redshift bins. Bottom row: correlation matrix obtained using the DIR algorithm for each of the four galaxies' redshift distributions. Note that for visualization purposes we display the absolute values of the each correlation matrix in logarithmic scale. In this plot, we can see that the covariance matrices obtained through the DIR algorithm are mostly diagonal.

has been updated in equation (10) by increasing the uncertainty in the data modes that most prominently respond to variations in the nuisance parameters.

In this work, $\boldsymbol{\nu}$ corresponds to the parameters describing the redshift distribution uncertainties, i.e. one shift parameter per redshift bin when using the parametric approach, or a set of $p(z)$ bin heights in the non-parametric scheme. The method described earlier however is fully general and has in the past been applied to marginalize over other types of nuisance parameters, including multiplicative shape measurement biases (Hildebrandt et al. 2020b), as well as truly linear parameters such as shot-noise (García-García et al. 2021) or systematic template amplitudes (Koukoufilippas et al. 2020). The aim of this paper is thus to determine the applicability of this method to the case of redshift distribution uncertainties.

3 DATA

In order to evaluate the performance of the analytical marginalization approach described in the previous section in the context of current and future surveys, we make use of data from the first-year cosmic shear analysis of the Dark Energy Survey (DES-Y1; Abbott et al. 2018b). The aim of this is two-fold: first, to demonstrate that the method can be successfully implemented in real data, with real-life complications (e.g. noisy $p(z)$s, numerical covariances, astrophysical and observational systematics) and, secondly, to demonstrate this validity for future Stage-IV data sets in the presence of $p(z)$ calibration uncertainties already achieved on current data. This section describes the DES-Y1 data used, and the models used to generate simulated future Stage-IV data.

3.1 DES-Y1 data and redshift distributions

The Dark Energy Survey is a photometric, 5-yr survey, that has observed 5000 deg^2 of the sky using five different filter bands (*grizY*). The observations were made with the 4 m Blanco Telescope, provided with the 570-Mpix Dark Energy Camera (DECam), from the Cerro Tololo Inter-American Observatory, in Chile. In this paper, we use cosmic shear data from the first data release (Abbott et al. 2018b), which covers 1786 deg^2 before masking. In particular, we use the public METACALIBRATION source catalogue,[2] which is divided in four redshift bins covering the range $z \lesssim 1.6$ (Hoyle et al. 2018).

We use the calibrated redshift distributions of the METACALIBRATION sample provided by García-García et al. (2023). The $p(z)$s were estimated via DIR (Lima et al. 2008), using the COSMOS 30-band catalogue (Laigle et al. 2016) as a calibrating sample. The uncertainties of the measured redshift distributions were estimated analytically, as described in García-García et al. (2023), accounting for both shot-noise and sample variance, and represent a realistic level of $p(z)$ uncertainty achieved by current existing data sets. The redshift distributions were sampled on 40 bins of width $\delta z = 0.04$ covering the range $0 \leq z \leq 1.6$. Fig. 1 shows, in the first row, the redshift distributions of the four METACALIBRATION samples and their statistical uncertainties. Note that we estimated the full covariance matrix of the $p(z)$ bin heights. The covariance is dominated by the diagonal, as can be seen in the bottom panels of Fig. 1.

We will also use the cosmic shear angular power spectra provided by Nicola et al. (2021). A full description of the methods used to estimate these power spectra, and their associated covariance matrix, from the DES-Y1 data is provided by the authors.

3.2 Future Stage-IV data

We generate a simulated data vector corresponding to a Stage-IV cosmic shear survey, such as the Legacy Survey of Space Time, at the Rubin Observatory (LSST Dark Energy Science Collaboration 2012), or the Euclid survey (Spergel et al. 2015). Our aim is to

[2] https://desdr-server.ncsa.illinois.edu/despublic/y1a1_files/

effectively test the analytical marginalization method in the low-noise regime, where the inferred posterior is likely more sensitive to residual $p(z)$ uncertainties, and the error budget may become dominated by these, rather than the statistical errors in the data themselves.

For simplicity, we simulate the Stage-IV survey as having the same redshift distributions as the DES-Y1 sample. This includes both the $p(z)$s themselves, and their calibration uncertainties. While it is possible that techniques for inferring redshifts from photometry, or the size and quality of calibrating spectroscopic samples, will improve substantially by the time Stage-IV data are available, we prefer to err on the side of caution and assume the same performance as currently achieved. For instance, it is possible that redshift estimates will suffer commensurately with the increase in survey depth. The results presented here are therefore conservative, and their validity will only be reinforced if better $p(z)$ calibration samples are used in the future.

We generate cosmic shear power spectra using CCL (Chisari et al. 2019) for the best-fitting *Planck* 2018 cosmological parameters (Planck Collaboration 2020): $\Omega_b h^2 = 0.02237$, $\Omega_c h^2 = 0.12$, $h = 0.6736$, $10^9 A_s = 2.0830$, $n_s = 0.9649$, $w_0 = -1$, and $w_a = 0$. We use the same sampling in ℓ used for the DES-Y1 power spectra, and use only scales in the range $\ell \in [30, 2000]$.

We compute the covariance matrix of these power spectra analytically, including a disconnected 'Gaussian' component, and a connected super-sample covariance contribution (SSC).

$$\mathrm{Cov}\left(C_\ell^\alpha, C_{\ell'}^{\rho\sigma}\right) = \mathrm{Cov}_G\left(C_\ell^{\alpha\beta}, C_{\ell'}^{\rho\sigma}\right) + \mathrm{Cov}_{\mathrm{SSC}}\left(C_\ell^{\alpha\beta}, C_{\ell'}^{\sigma\rho}\right). \quad (11)$$

We estimate the Gaussian covariance using a simple mode-counting approximation (Efstathiou 2004) as

$$\mathrm{Cov}_G\left(C_\ell^{\alpha\beta}, C_{\ell'}^{\rho\sigma}\right) = \delta_{\ell\ell'} \frac{C_\ell^{\alpha\rho} C_\ell^{\beta\sigma} + C_\ell^{\alpha\sigma} C_\ell^{\beta\rho}}{(2\ell+1)\,\Delta\ell\, f_{\mathrm{sky}}}, \quad (12)$$

where f_{sky} is the fraction of the sky covered by the experiment. We assume $f_{\mathrm{sky}} = 0.4$, as in the case of LSST (LSST Dark Energy Science Collaboration 2012). The angular power spectra above contain the contribution from shape noise in the autocorrelation, of the form

$$N_\ell^{\alpha\alpha} = \frac{\sigma_\gamma^2}{\bar{n}_\alpha}. \quad (13)$$

Here, $\sigma_\gamma = 0.28$ is the per-component ellipticity dispersion in each source, and $\bar{n}_\alpha$ is the angular number density of sources in the αth redshift bin. We assume $n_\alpha = 4\,\mathrm{arcmin}^{-2}$ in each redshift bin.

We compute the SSC following:

$$\mathrm{Cov}_{\mathrm{SSC}}(C_\ell^{\alpha\beta}, C_{\ell'}^{\rho\sigma}) = \int \mathrm{d}\chi\ \frac{q^\alpha(\chi) q^\beta(\chi) q^\rho(\chi) q^\sigma(\chi)}{\chi^4} \times \quad (14)$$

$$\frac{\partial P(\ell/\chi, z)}{\partial \delta_{\mathrm{LS}}} \frac{\partial P(\ell'/\chi, z)}{\partial \delta_{\mathrm{LS}}} \sigma^2_{\mathrm{LS}}(z), \quad (15)$$

as in Nicola et al. (2021). $\partial P(k, z)/\partial \delta_{\mathrm{LS}}$ is the response of the matter power spectrum to a large-scale density fluctuation δ_{LS}, and the quantity $\sigma_b^2(z)$ is the variance of the long-wavelength mode over the survey footprint. We estimate the latter as in Krause & Eifler (2017), modelling the footprint simply as a circular cap of area $4\pi f_{\mathrm{sky}}$. We estimate the response function using perturbation theory and the halo model, as described in Krause & Eifler (2017), and as implemented in CCL.

4 LIKELIHOOD

We extract cosmological parameter constraints using a Gaussian likelihood as described in Section 2.3. In order to validate the analytical marginalization approach, we will either use the full posterior distribution in equation (7), or the analytically marginalized version in equation (9).[3] In the first case, $\boldsymbol{\nu}$ includes all nuisance parameters describing the redshift distribution uncertainties, and in both cases $\boldsymbol{\Omega}$ includes all other model parameters. Specifically, $\boldsymbol{\Omega}$ contains the five ΛCDM cosmological parameters (Ω_m, Ω_b, σ_8, n_s, h).

When marginalizing over redshift distribution uncertainties, $\boldsymbol{\nu}$ will contain either one redshift shift parameter Δz_α for each redshift bin, when employing the parametric description of $p(z)$ uncertainties (Method 1), or a set of bin heights for each redshift bin determining $p_\alpha(z)$, when using the non-parametric approach (Method 2). The first case will introduce four new parameters to the model, while the latter will introduce $4 \times 40 = 160$ new amplitude parameters, as described in Section 3.1.

Table 1. Prior distributions for the parameters considered in this work. Note that the redshift calibration section contains the priors for both the Δz and $p_\alpha(z)$ models which are not sampled simultaneously.

Parameter priors			
Parameter	Prior	Parameter	Prior
Cosmology		**Redshift calibration**	
Ω_m	$U(0.1, 0.9)$	Δz_1	$N(0.0, 0.016)$
Ω_b	$U(0.03, 0.07)$	Δz_2	$N(0.0, 0.017)$
h	$U(0.55, 0.91)$	Δz_3	$N(0.0, 0.013)$
n_s	$U(0.87, 1.07)$	Δz_4	$N(0.0, 0.015)$
σ_8	$U(0.6, 0.9)$	p_i	$N(\bar{p}_i, C)$
Shear multiplicative bias			
m^i	0.012		

Table 1 shows the parameter priors used in this work. All cosmological parameters take uniform, largely uninformative priors. For simplicity, the multiplicative bias parameters were fixed at the centre of the Gaussian priors from the official analysis of DES-Y1 (Abbott et al. 2018a). When using Method 1 to numerically marginalize over the $p(z)$ uncertainties, we used Gaussian priors on each of the shift parameters Δz_α, following those used by DES-Y1 (Abbott et al. 2018a). When using Method 2 (marginalization over $p(z)$ bin amplitudes), we assume a multivariate Gaussian prior, with the $p(z)$ covariance described in Section 3.1 and shown in Fig. 1.

For both $p(z)$ uncertainty models, when using analytical marginalization, we use equation (9) and modify the covariance as in equation (10), with P given by the priors described earlier. When using numerical marginalization, we simply explore the posterior distribution of the full model, including all the $p(z)$ and p_i parameters. In the case of Method 2, this involves sampling a distribution with 165 parameters, of which the bulk (160 parameters) describe the $p(z)$ uncertainty. This is not feasible for standard Metropolis–Hastings MCMC methods (Metropolis et al. 1953; Hastings 1970) due to the curse of dimensionality, and therefore, we resort to an HMC approach.

HMC (MacKay 2002; Betancourt 2017) uses notions of Hamiltonian dynamics to draw trajectories on the parameter space along which the sampler moves. This results in a much greater acceptance rate, and allows HMC to beat the dimensionality curse. HMC can thus efficiently explore parameter spaces with large numbers of dimensions in far less time than Metropolis–Hastings or nested sampling techniques (Alsing & Handley 2021). The main difficulty of using HMC is the need to calculate gradients of the log-posterior

[3] Recall that we treat the term in the second line of equation (9) as a constant.

to calculate the Hamiltonian equations of motion. The additional computational cost of obtaining these derivatives numerically (e.g. via adaptive finite differences) may outweigh the gains caused by the higher acceptance rates of HMC. To overcome this problem, we make use of automatic differentiation (AD). To take advantage of AD, we have developed a cosmological theoretical prediction code natively written in the JULIA programming language (Ruiz-Zapatero et al., in preparation). JULIA is a just-in-time compiled language with C-like performance and seamless AD integration, which can thus be used to efficiently sample complex cosmological posteriors using HMC. To sample the posterior distribution, we use the No-U-Turns Sampler (Hoffman & Gelman 2011) implementation of HMC within the TURING.JL package (Ge, Xu & Ghahramani 2018).

5 RESULTS

5.1 Linearizing Δz

Let us begin the discussion of our results by considering the simplest of the two models of the photometric uncertainties studied in this work, the Δz model (called Method 1 earlier). As discussed in Section 4, this model introduces four new shift parameters Δz (one per redshift bin) in addition to the five ΛCDM parameters. All other nuisance parameters are kept fixed. For the DES-Y1 and LSST-like data sets, we will compare the result of analytically marginalizing over the Δz parameters against performing the full numerical marginalization on the corresponding cosmological constraints. In order to quantify the contribution of redshift uncertainties to the total error budget, we will also present results for the case when the Δz parameters are fixed (i.e. assuming perfect knowledge of the redshift distributions).

Our results for DES-Y1 data are shown in Fig. 2, with the errors on all parameters listed in Table 2. On the one hand, we find that marginalizing analytically or numerically over the Δz parameters leads to the same marginalized posterior for the cosmological parameters. On the other hand, fixing the Δz parameters returns a posterior distribution that is only mildly narrower than the marginal distribution. For the DES-Y1 data, the impact of redshift uncertainties in the final cosmological errors is relatively small (although not negligible). Thus, if we truly wish to study the effect of marginalizing analytically as opposed to numerically over the Δz parameters, we will have to consider futuristic LSST-like data, where the impact of these uncertainties will likely be higher.

We show results for futuristic LSST-like data on Fig. 3, with the parameter constraints listed in Table 2. First of all, in the case of LSST-like data we observe that not marginalizing over the Δz parameters in the model results in significantly narrower posteriors, with the final uncertainties shrinking by a factor of $\sim$2. The impact of redshift distribution uncertainties in this case is thus much more relevant, and the accuracy of the analytical marginalization scheme becomes paramount. However, comparing the contours obtained by numerical and analytical marginalization, we observe that both methods return largely equivalent posterior distributions, with the final uncertainties changing by much less than 10 per cent. This holds even in the case, where the Δz prior worsen by a factor of 4 as seen in Fig. A1, in Appendix A. Therefore, linearizing the likelihood around the Δz parameters will be a good enough approximation for LSST-like data, at least for relatively simple parametrizations of the $p(z)$ uncertainty, which will allow us to reduce the dimensionality of the model and make parameter inference more efficient.

It is worth emphasizing that the results in this section are not meant to be interpreted as forecasts on the constraining power of LSST on cosmological parameters, but only on our ability to analytically marginalize over photometric uncertainties in inferring the underlying cosmology. The recovered constraints depend strongly on assumptions such as the redshift calibration that LSST will be able to achieve for the different samples involved. As such, the results presented here are only a conservative estimate of the effect of analytical marginalization on cosmological constraints.

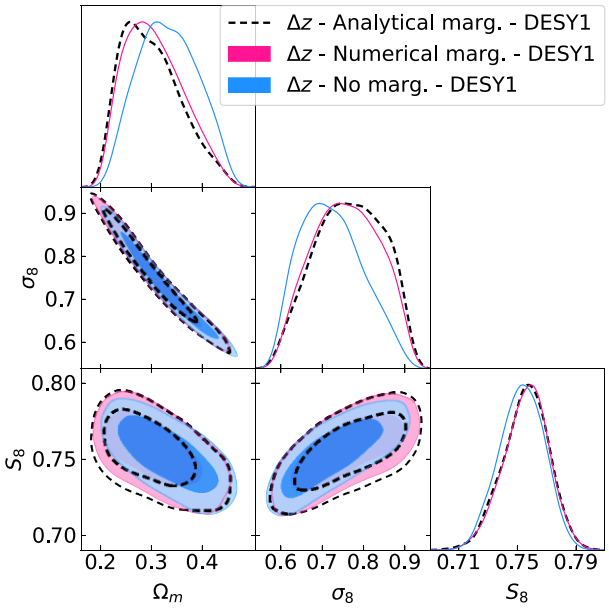

Figure 2. Marginalized posterior distributions for the combination of parameters Ω_m, σ_8, and S_8 obtained when considering the Δz model for photometric uncertainties for DES-Y1 data. The blue contours correspond to the case where the Δz parameter are fixed. The magenta contours are obtained when numerically marginalizing over the Δz parameters. Finally, the black dashed contours are obtained when analytically marginalizing over the Δz parameters. We can observe that the analytical and numerical marginalization return nearly identical posteriors.

Table 2. Numerical values for the mean and 1σ confidence intervals for the 1D marginalized posterior distributions of the cosmological parameters Ω_m, σ_8, and S_8 obtained when considering the first method (z shifts) to characterize the photometric redshift uncertainties. The first column shows the values obtained when the Δz parameters were kept fixed, the second column when they were marginalized numerically and the third column when they were marginalized analytically. In each row, we display the constraints obtained when using DES-Y1 or LSST-like data to constrain the models.

	Δz model	Fixed	Numerical	Analytical
Ω_m	DES-Y1	0.333 ± 0.055	0.3 ± 0.056	0.306 ± 0.055
–	LSST	0.311 ± 0.011	0.317 ± 0.02	0.317 ± 0.02
σ_8	DES-Y1	0.724 ± 0.072	0.765 ± 0.077	0.758 ± 0.076
–	LSST	0.82 ± 0.015	0.821 ± 0.027	0.823 ± 0.027
S_8	DES-Y1	0.753 ± 0.015	0.756 ± 0.015	0.756 ± 0.015
–	LSST	0.833 ± 0.002	0.833 ± 0.005	0.833 ± 0.006

5.2 Linearizing $p_\alpha(z)$

In the previous section, we have shown that even for futuristic LSST-like data, it is possible to marginalize over redshift uncertainties analytically, assuming a relatively simple parametrization of these

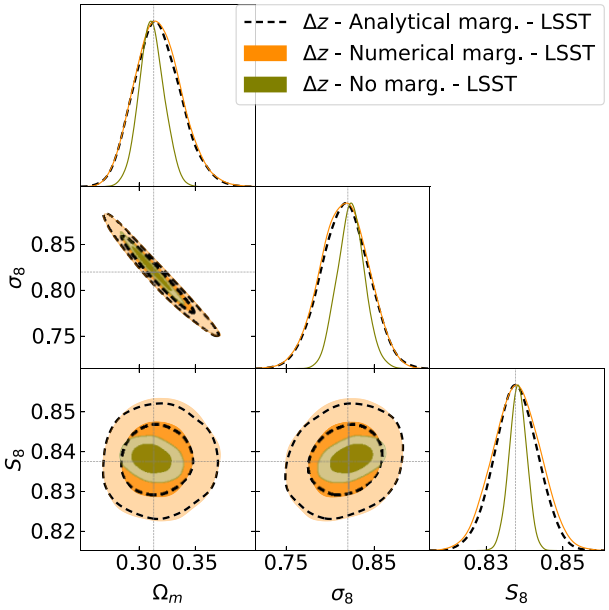

Figure 3. Marginalized posterior distributions for the combination of parameters Ω_m, σ_8, and S_8 obtained when considering the Δz model for photometric uncertainties for futuristic LSST-like data. The green contours correspond to the case where the Δz parameter are fixed. The orange contours are obtained when numerically marginalizing over the Δz parameters. The black dashed contours are obtained when analytically marginalizing over the Δz parameters. Finally, the dotted lines mark the values of the fiducial cosmology used to generate the data. We can observe that the analytical and numerical marginalizations return nearly identical posteriors.

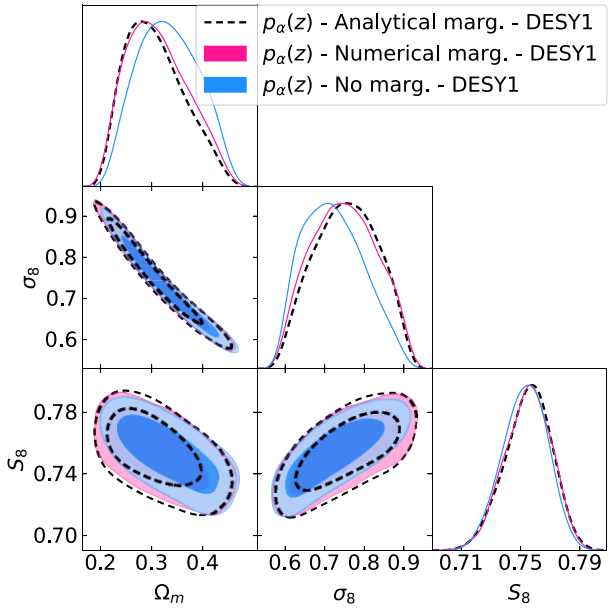

Figure 4. Marginalized posterior distributions for the combination of parameters Ω_m, σ_8, and S_8 obtained when considering the $p_\alpha(z)$ model for photometric uncertainties for DES-Y1 data. The blue contours correspond to the case where the $p_\alpha(z)$ parameter are fixed. The magenta contours are obtained when numerically marginalizing over the $p_\alpha(z)$ parameters. Finally, the black dashed contours are obtained when analytically marginalizing over the $p_\alpha(z)$ parameters. We can observe that the analytical and numerical marginalization return nearly identical posteriors.

uncertainties. We now turn to more complex models to characterize these uncertainties.

In order to do so, we consider the previously discussed $p_\alpha(z)$ model (called Method 2 above), which turns the height of each bin in the redshift distribution histograms into a free parameter. This results in 40 new free parameters per redshift bin with a total of 160 parameters for the data considered in this work.

We start by revisiting the DES-Y1 data analysis, presenting our results in Fig. 4. As we observed in the previous section, we find that even for the far more general $p_\alpha(z)$ model, there is no significant difference between numerically marginalizing over the $p_\alpha(z)$, or doing so through our approximate analytical approach. Furthermore, as before, fixing the shape of the redshift distribution leads to only mildly tighter constraints. On the one hand, this means that the result found for the Δz model is not reliant on the simplicity of the model, but instead inherent to the sensitivity of DES-Y1 data. On the other hand, this also means that we must turn once again to futuristic LSST-like data to study the impact of a more general parametrization of photometric uncertainties.

The results for futuristic LSST-like data are shown in Fig. 5. As in the case of the Δz parametrization, we find that, in the case of LSST-like data, not including the $p_\alpha(z)$ parameters in the model results in significantly narrower posteriors. By looking at the corresponding numerical values in Table 3, we see that the S_8 constraints become twice as tight when the $p_\alpha(z)$ parameters are fixed. Most importantly, we find that marginalizing over the $p_\alpha(z)$ parameters analytically or numerically yields almost indistinguishable posteriors. Thus, the results found in Section 5.1 for the simple Δz parametrization, in fact, hold for significantly more general models of the uncertainty in the galaxy redshift distributions.

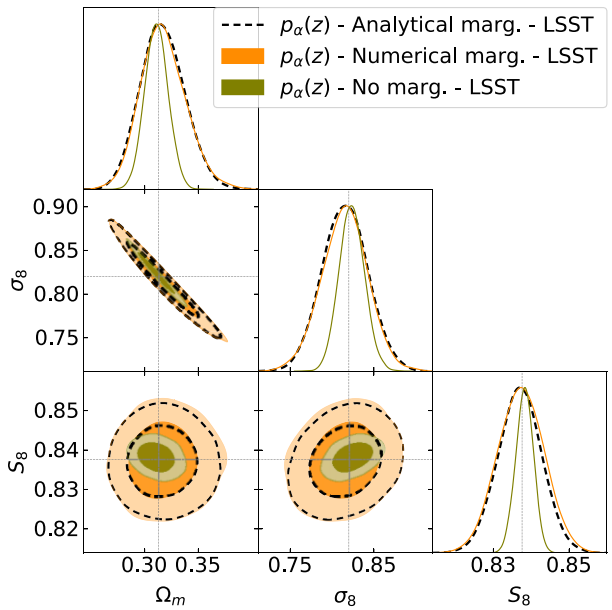

Figure 5. Marginalized posterior distributions for the combination of parameters Ω_m, σ_8, and S_8 obtained when considering the $p_\alpha(z)$ model for photometric uncertainties for LSST-like futuristic data. The green contours correspond to the case, where the $p_\alpha(z)$ parameter are fixed. The orange contours were obtained when numerically marginalizing over the $p_\alpha(z)$ parameters. The black dashed contours were obtained when analytically marginalizing over the $p_\alpha(z)$ parameters. Finally, the dotted lines mark the values of the fiducial cosmology used to generate the data. We can observe that the analytical and numerical marginalization return nearly identical posteriors.

Table 3. Numerical values for the mean and 1σ confidence intervals for the 1D marginalized posterior distributions of the cosmological parameters Ω_m, σ_8, and S_8 obtained when considering the second method [$p(z)$ bin heights] to characterize the photometric redshift uncertainties. The first column shows the values obtained when the $p_\alpha(z)$ parameters are kept fixed, the second column when they are marginalized numerically, and the third column when they are marginalized analytically. In each row, we display the constraints obtained when using DES-Y1 or LSST-like data to constrain the models.

$p_\alpha(z)$ model		Fixed	Numerical	Analytical
Ω_m	DES-Y1	0.333 ± 0.056	0.308 ± 0.055	0.312 ± 0.057
	LSST	0.311 ± 0.011	0.317 ± 0.02	0.317 ± 0.021
σ_8	DES-Y1	0.723 ± 0.073	0.755 ± 0.075	0.75 ± 0.077
	LSST	0.824 ± 0.015	0.816 ± 0.026	0.815 ± 0.027
S_8	DES-Y1	0.753 ± 0.015	0.755 ± 0.015	0.755 ± 0.015
	LSST	0.838 ± 0.002	0.837 ± 0.006	0.837 ± 0.006

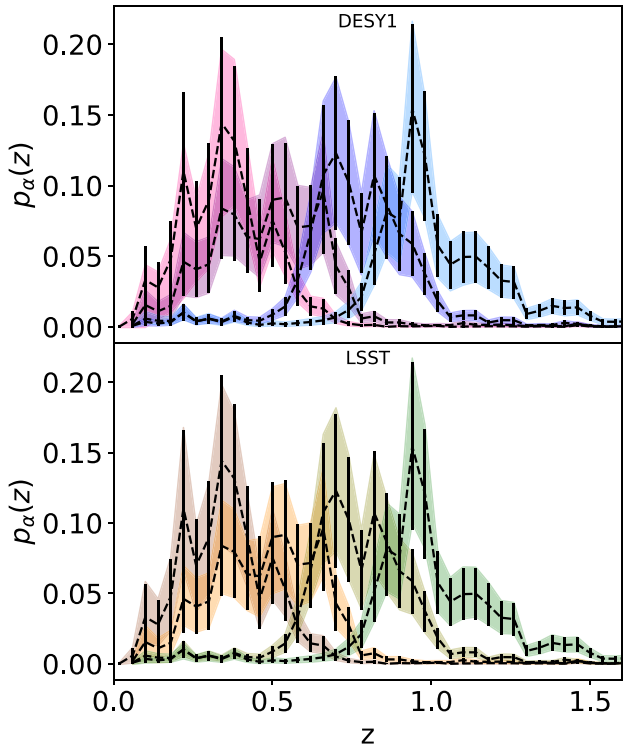

Figure 6. Posterior distributions for the $p_\alpha(z)$ parameters when considering DES-Y1 data (top row) and futuristic LSST-like data (bottom row). The black dashed line shows the mean of the Gaussian prior of the $p_\alpha(z)$ parameters. The error bars show their corresponding error.

Finally, in Fig. 6 we present the constraints obtained for the 160 $p_\alpha(z)$ parameters for both the DES-Y1 (top panel) and LSST-like data (bottom panel) in colour bands. We observe that the posterior distributions are largely dominated by the prior (shown in dashed black line with error bars), and thus, the redshift distribution is not significantly self-calibrated by the data in either case.

Before moving to the next section, it is worth stressing that constraining such a large parameter space has only been possible thanks to the autodifferentiable nature of the code used to obtain theoretical predictions, allowing us to use gradient-based samplers, much more efficient than standard samplers. The development of such autodifferentiable codes will therefore become imperative in the near future given the increasing complexity of models used in cosmological analyses.

5.3 Δz versus $p_\alpha(z)$

In the previous sections, we have focused on the impact of how we marginalize over the different parametrizations of photometric redshift uncertainties. In this section, we will focus instead on what we marginalize over, i.e. the impact of the choice of parametrization. The question is then: Can a one-parameter-per-bin model (Δz model) capture all the meaningful modifications to photometric redshift distributions?

In order to answer this question, we constrain the cosmological parameters for the Δz and $p_\alpha(z)$ models in the case with futuristic LSST-like data. In both cases, we marginalize numerically over their respective nuisance parameters. As shown in Fig. 7 and Tables 2 and 3, both methods recover the same posterior distributions with small differences. Moreover, we observed no biases with respect to the fiducial cosmology regardless of the parametrization of the photometric redshift uncertainty or the method used to marginalize over it.

Thus, it is in principle possible that even Stage-IV surveys will be able to use relatively simple models to describe the redshift distribution of cosmic shear samples.[4] Note, however, that this result is subject to the modelling of the rest of the systematic effects in the survey. More complex modelling of IA might require better control on the $p(z)$ systematics. In such case, the Δ_z parametrization might become insufficient and one would expect a difference between said parametrization and sampling the full $p_\alpha(z)$.

6 CONCLUSIONS

One of the most significant obstacles to overcome in photometric weak lensing surveys is the accurate modelling of redshift distributions, $p(z)$. Not only are our measurements prone to error, which can bias the inferred cosmological parameters, but accounting for these uncertainties is also a major inhibitor of efficient parameter inference. In this paper, we investigate the impact of analytically marginalizing over the uncertainties in the redshift distribution of galaxies in weak lensing surveys, as initially proposed in Hadzhiyska et al. (2020). In particular, we thoroughly quantify the validity of this approach for a current weak lensing survey, DES, as well as for a futuristic LSST-like survey, testing whether a fast analytical method proposed in this work is capable of reproducing the posterior distributions and constraints one arrives at when adopting the traditional method of diligently varying tens or hundreds of nuisance parameters.

Our results show that, for present surveys, marginalizing over the uncertainty in the redshift distribution of galaxies has only a mild impact on the constraints on cosmological parameters, although one that our analytical approximation is able to reproduce accurately. This is true for the two parametrizations of the uncertainties considered in this work, in terms of shifts in the mean redshift or redshift distribution histogram heights. However, the impact of redshift distribution uncertainties changes dramatically for future LSST-like surveys. In this case, redshift uncertainties commensurate with current calibration samples lead to a degradation in the final

[4]Note, however, that this is likely not the case for photometric galaxy clustering studies where other properties of the redshift distribution (e.g. its width) have a stronger impact on the theoretical prediction (Nicola et al. 2020).

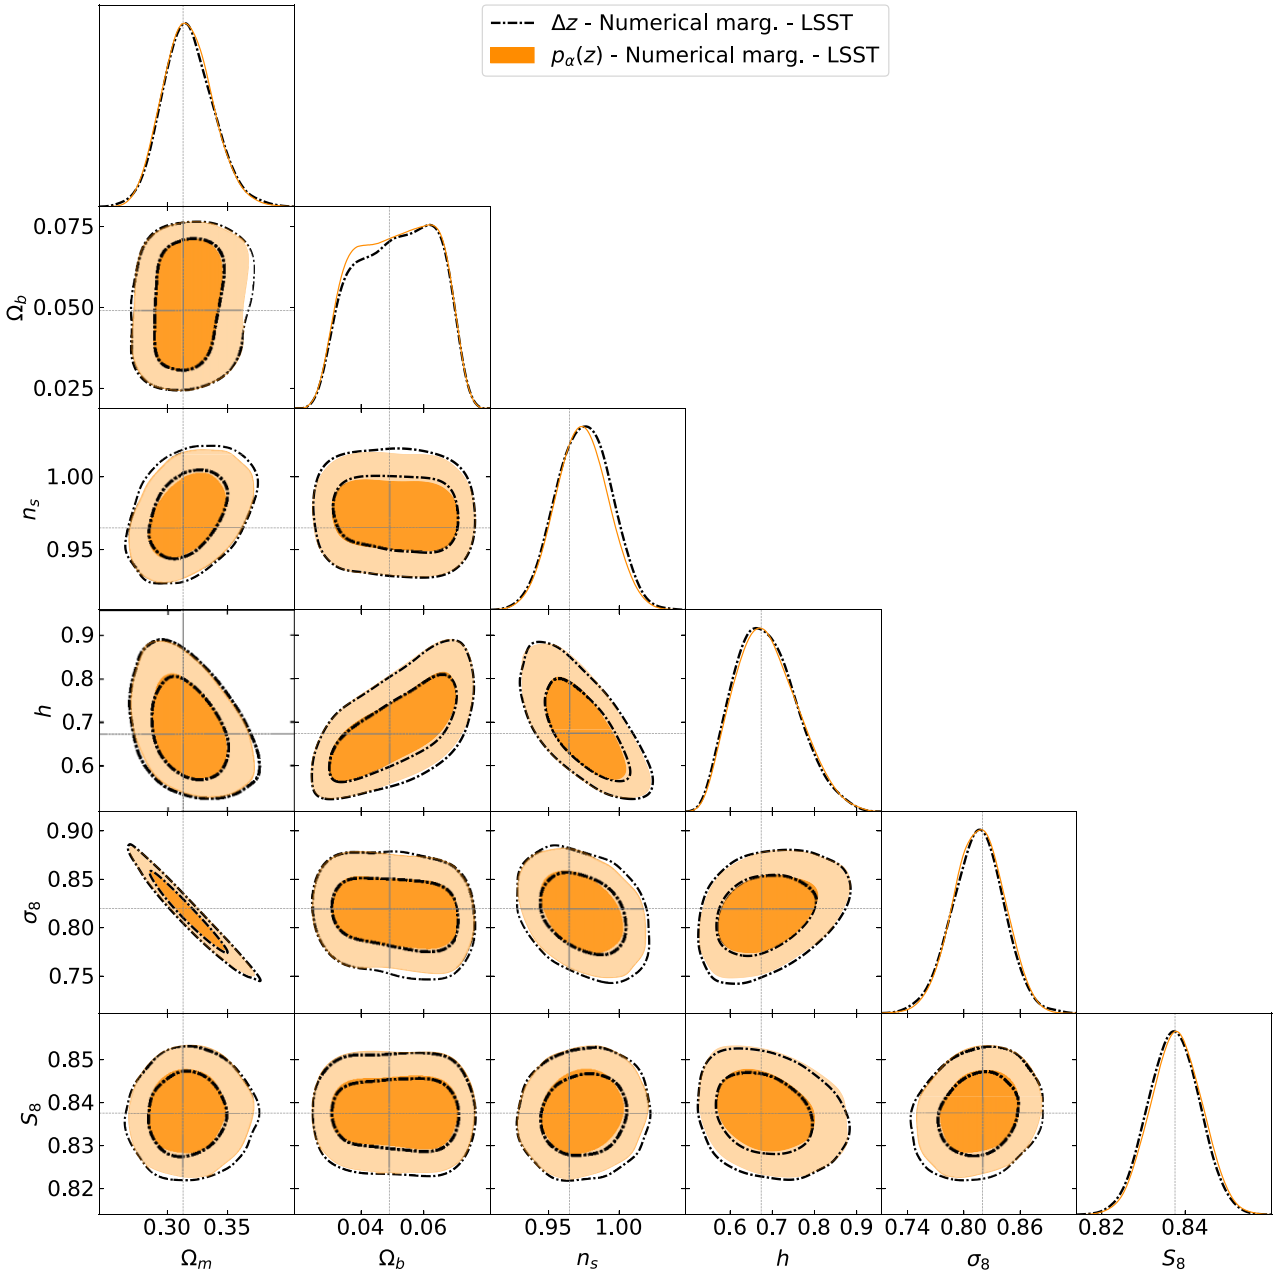

Figure 7. Comparison between the obtained marginalized posterior distributions of the cosmological parameters when numerically marginalizing over the Δz (black dash–dotted) and $p_\alpha(z)$ (orange) photometric uncertainties models when applied to LSST-like futuristic data. Dotted lines mark the values of the fiducial cosmology used to generate the data. We can observe that both parametrizations of the photometric redshift uncertainties return identical posteriors for the cosmological parameters.

constraints on cosmological parameters of up to a factor of ∼2. Capturing this effect for an arbitrarily complex parametrization of the redshift distribution uncertainties is an a priori difficult task without resorting to a full exploration of the parameter space. Nevertheless, we find that the analytical approximate scheme explored here is still able to recover the marginalized constraints on cosmological parameters to high fidelity, even after marginalizing over more than 100 nuisance parameters. This means that, while future surveys will certainly have to account for these uncertainties, they will be able to do so using fast marginalization methods without increasing the dimensionality of their astrophysical and cosmological models.

Our results have also shown that simple parametrizations of the redshift distribution for cosmic shear samples, in terms of shifts in the mean redshift, are, surprisingly, able to reproduce the impact of the full uncertainty on $p(z)$ on the final constraints to high precision. Although this result will likely not hold for other probes (e.g. tomographic galaxy clustering), it should certainly simplify the analysis of future cosmic shear data.

It is worth emphasizing that our work has focused exclusively on the case of cosmic shear data, and that our conclusions only apply in this context. The validity of the analytical approximation employed here for general tomographic tracers of structure with uncertain radial kernels is not guaranteed, and future work should

quantify its performance on photometric clustering data – the other key probe of the flagship '3 × 2pt' analysis of imaging surveys – and its cross-correlation with cosmic shear and CMB lensing data (García-García et al. 2021; Heymans et al. 2021; Abbott et al. 2022; White et al. 2022).

ACKNOWLEDGEMENTS

We would like to thank Aně Slosar and Marius Millea for useful discussions. DA was supported by the Science and Technology Facilities Council through an Ernest Rutherford Fellowship, grant reference no. ST/P004474. PGF, CGG, and AM were supported by European Research Council grant no. 693024 and the Beecroft Trust. JRZ was supported by an STFC doctoral studentship. We made extensive use of computational resources at the University of Oxford Department of Physics, funded by the John Fell Oxford University Press Research Fund. We made extensive use of the NUMPY (Oliphant 2006; Van Der Walt, Colbert & Varoquaux 2011), SCIPY (Virtanen et al. 2020), ASTROPY (Astropy Collaboration 2013, 2018), HEALPY (Zonca et al. 2019), GETDIST (Lewis 2019), and MATPLOTLIB (Hunter 2007) PYTHON packages. We also make use of the JULIA packages FORWARDDIFF.JL (Revels, Lubin & Papamarkou 2016) and TURING.JL (Ge et al. 2018).

DATA AVAILABILITY

The code developed for this work as well as the derived data sets produced (power spectra and covariances) are available upon request. The catalogues and maps used were made publicly available by the authors of the relevant papers, as described in the text.

REFERENCES

Abbott T. M. C. et al., 2018a, Phys. Rev. D, 98, 043526
Abbott T. M. C. et al., 2018b, ApJS, 239, 18
Abbott T. M. C. et al., 2022, Phys. Rev. D, 105, 023520
Afshordi N., Loh Y.-S., Strauss M. A., 2004, Phys. Rev. D, 69, 083524
Alsing J., Handley W., 2021, MNRAS, 505, L95
Astropy Collaboration, 2013, A&A, 558, A33
Astropy Collaboration, 2018, AJ, 156, 123
Bartelmann M., Schneider P., 2001, Phys. Rep., 340, 291
Betancourt M., 2017, preprint (arXiv:1701.02434)
Bonnett C. et al., 2016, Phys. Rev. D, 94, 042005
Brown M. L., Taylor A. N., Hambly N. C., Dye S., 2002, MNRAS, 333, 501
Chisari N. E. et al., 2019, ApJS, 242, 2
Cordero J. P. et al., 2022, MNRAS, 511, 2170
DES Collaboration, 2021, ApJS, 254, 24
DES Collaboration, 2022, Phys. Rev. D, 105, 023515
Efstathiou G., 2004, MNRAS, 349, 603
Ferreira P. G., 2019, ARA&A, 57, 335
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
García-García C., Ruiz-Zapatero J., Alonso D., Bellini E., Ferreira P. G., Mueller E.-M., Nicola A., Ruiz-Lapuente P., 2021, J. Cosmol. Astropart. Phys., 2021, 030
García-García C., Alonso D., Ferreira P. G., Hadzhiyska B., Nicola A., Sánchez C., Slosar A., 2023, J. Cosmol. Astropart. Phys., 2023, 025
Ge H., Xu K., Ghahramani Z., 2018, in Turing: A Language for Flexible Probabilistic Inference, Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics, AISTATS 2018. Storkey Amos, Perez-Cruz Fernando, 84, 1682, http://proceedings.mlr.press/v84/ge18b/ge18b.pdf, https://proceedings.mlr.press/v84/ge18b.html
Hadzhiyska B., Alonso D., Nicola A., Slosar A., 2020, J. Cosmol. Astropart. Phys., 2020, 056
Hadzhiyska B. et al., 2023, Cosmology with 6 parameters in the Stage-IV era: efficient marginalisation over nuisance parameters, Astrophysics - Cosmology and Nongalactic Astrophysics, Astrophysics - Instrumentation and Methods for Astrophysics
Hastings W. K., 1970, Biometrika, 57, 97
Heymans C. et al., 2021, A&A, 646, A140
Hildebrandt H. et al., 2020a, A&A, 633, A69
Hildebrandt H. et al., 2020b, A&A, 633, A69
Hirata C. M., Seljak U., 2004, Phys. Rev. D, 70, 063526
Hoffman M. D., Gelman A., 2011, preprint (arXiv:1111.4246)
Hoyle B. et al., 2018, MNRAS, 478, 592
Hunter J. D., 2007, Comput. Sci. Eng., 9, 90
Kilbinger M. et al., 2017, MNRAS, 472, 2126
Koukoufilippas N., Alonso D., Bilicki M., Peacock J. A., 2020, MNRAS, 491, 5464
Krause E., Eifler T., 2017, MNRAS, 470, 2100
Krause E. et al., 2017, preprint (arXiv:1706.09359)
Kuijken K. et al., 2015, MNRAS, 454, 3500
Laigle C. et al., 2016, ApJS, 224, 24
Lewis A., 2019, preprint (arXiv:1910.13970)
Lima M., Cunha C. E., Oyaizu H., Frieman J., Lin H., Sheldon E. S., 2008, MNRAS, 390, 118
Limber D. N., 1953, ApJ, 117, 134
LSST Dark Energy Science Collaboration, 2012, preprint (arXiv:1211.0310)
Ma Z., Hu W., Huterer D., 2006, ApJ, 636, 21
MacKay D. J. C., 2002, Information Theory, Inference and Learning Algorithms. Cambridge Univ. Press, Cambridge
Matthews D. J., Newman J. A., 2010, ApJ, 721, 456
Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, J. Chem. Phys., 21, 1087
Miyazaki S. et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 8446, Ground-Based and Airborne Instrumentation for Astronomy IV. SPIE, Bellingham, p. 84460Z
Newman J. A., 2008, ApJ, 684, 88
Newman E. T., Penrose R., 1966, J. Math. Phys., 7, 863
Nicola A. et al., 2020, J. Cosmol. Astropart. Phys., 2020, 044
Nicola A., García-García C., Alonso D., Dunkley J., Ferreira P. G., Slosar A., Spergel D. N., 2021, J. Cosmol. Astropart. Phys., 2021, 067
Oliphant T. E., 2006, A Guide to NumPy, Vol. 1. Trelgol Publishing, USA
Planck Collaboration, 2020, A&A, 641, A6
Prat J. et al., 2018, Phys. Rev. D, 98, 042005
Reischke R., 2023, preprint (arXiv:2301.04085)
Revels J., Lubin M., Papamarkou T., 2016, preprint(arXiv:1607.07892)
Riess A. G. et al., 2022, ApJ, 934, L7
Sánchez C., Bernstein G. M., 2019, MNRAS, 483, 2801
Sánchez C. et al., 2022, Phys. Rev. D, 105, 083529
Schmidt S. J., Ménard B., Scranton R., Morrison C., McBride C. K., 2013, MNRAS, 431, 3307
Schneider M., Knox L., Zhan H., Connolly A., 2006, ApJ, 651, 14
Scott D., 2018, preprint (arXiv:1804.01318)
Smith R. E. et al., 2003, MNRAS, 341, 1311
Spergel D. et al., 2015, preprint (arXiv:1503.03757)
Stölzner B., Joachimi B., Korn A., Hildebrandt H., Wright A. H., 2021, A&A, 650, A148
Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, ApJ, 761, 152
The LSST Dark Energy Science Collaboration, 2018, preprint (arXiv:1809.01669)
Troxel M. A. et al., 2018, Phys. Rev. D, 98, 043528
Van Der Walt S., Colbert S. C., Varoquaux G., 2011, Comput. Sci. Eng., 13, 22
Virtanen P. et al., 2020, Nat. Methods, 17, 261
White M. et al., 2022, J. Cosmol. Astropart. Phys., 2022, 007
Wright A. H., Hildebrandt H., van den Busch J. L., Heymans C., 2020, A&A, 637, A100
Yao J., Shan H., Zhang P., Kneib J.-P., Jullo E., 2020, ApJ, 904, 135
Zhang T., Rau M. M., Mandelbaum R., Li X., Moews B., 2023, MNRAS, 518, 709

Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, J. Open Source Softw., 4, 1298

APPENDIX A: STRESS-TESTING THE APPROXIMATION

As described in Section 2, the approximation used here to analytically marginalize over the redshift calibration parameters assumes a sufficiently tight prior on these parameters, such that the dependence of the theory prediction on them can be linearized. Testing whether this assumption might break in a realistic scenario, is therefore essential. This is important in the context of Stage-IV, since even though it is expected that spectroscopic samples and the associated calibration techniques will improve over time, the increase in depth that LSST-like surveys will represent may make the calibration of the faintest samples in the survey particularly challenging.

To further stress-test our approximate method, we repeat our analysis of the LSST-like futuristic data using the Δz model for redshift uncertainties with priors four times larger than the one used in our fiducial analysis (which themselves were based on existing calibration samples). The result of this test is shown in Fig. A1. Reassuringly, the results show that, despite quadrupling the uncertainty in the redshift nuisance parameters, the analytical

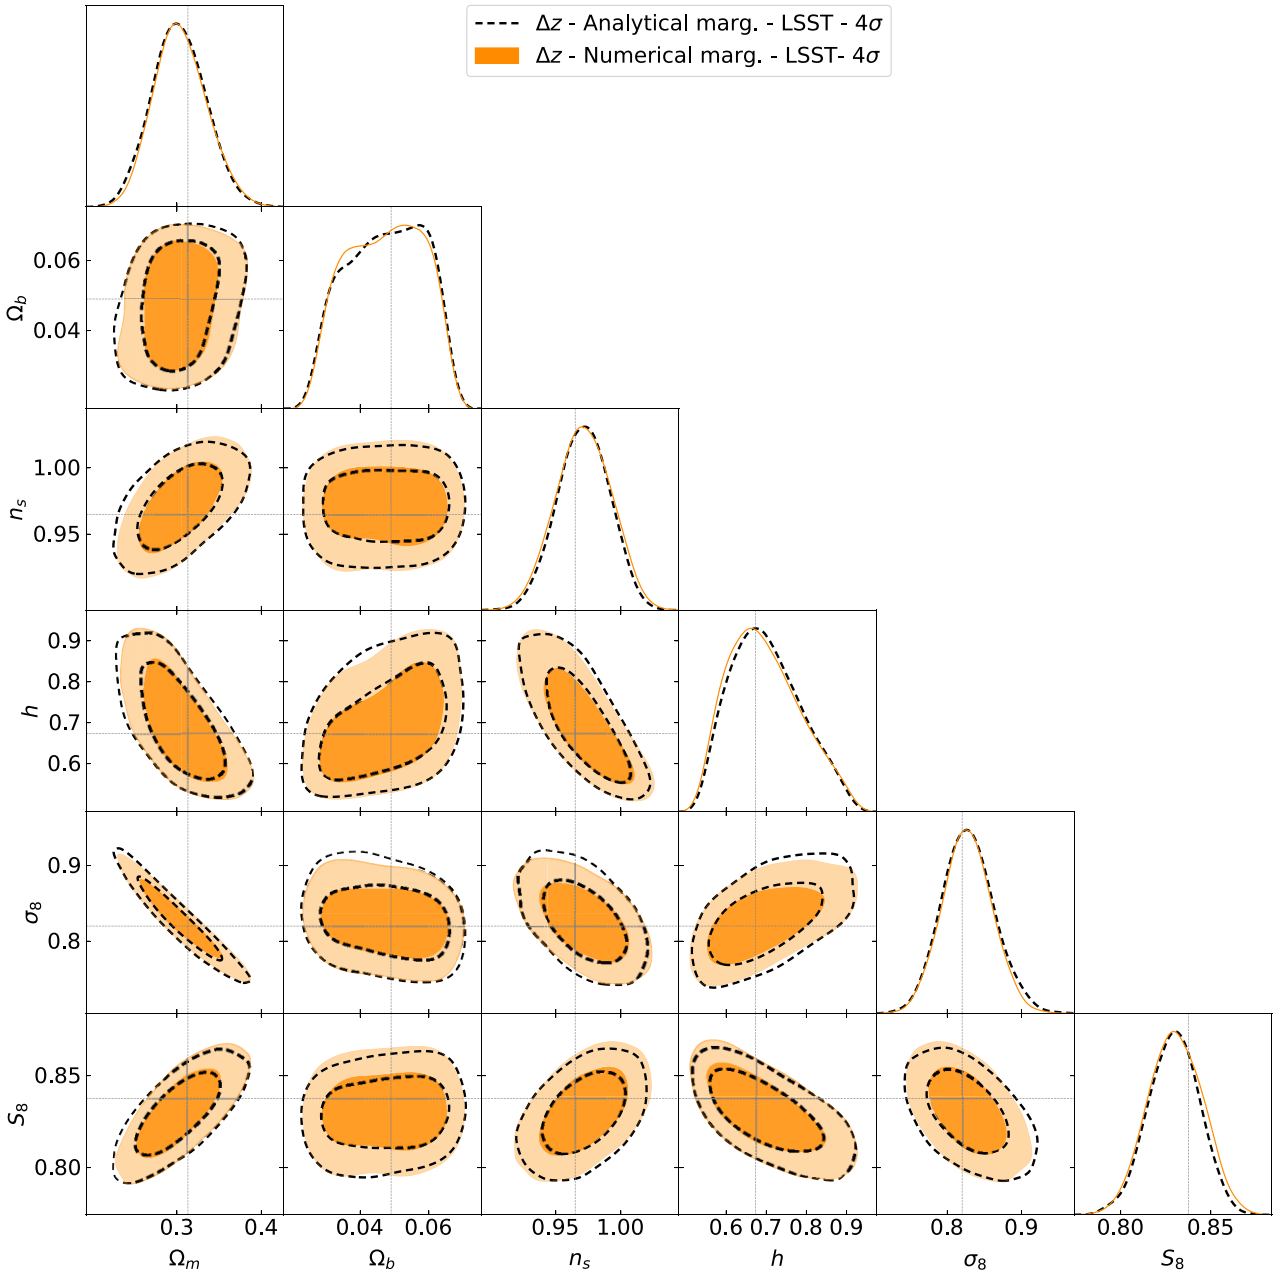

Figure A1. Shows a comparison between the obtained marginalized posterior distributions of the cosmological parameters when analytically marginalizing over the Δz (black dashed) and when performing the full numerical marginalization (orange) when analysing LSST-like data. In both cases, the Δz prior distributions were made four times wider. Dotted lines mark the values of the fiducial cosmology used to generate the data. We can observe that despite significantly broadening the prior distributions, the analytical marginalization returns virtually identical posteriors for the cosmological parameters.

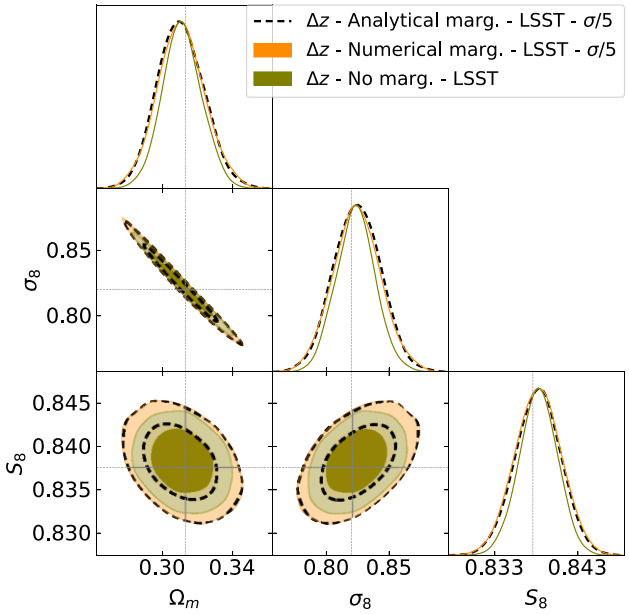

Figure A2. Marginalized posterior distributions for the combination of parameters Ω_m, σ_8, and S_8 obtained when considering the Δz model for photometric uncertainties for futuristic LSST-like data. In this analysis, the priors on the Δz parameters were made five times tighter. The green contours correspond to the case where the Δz parameters are fixed. The orange contours are obtained when numerically marginalizing over the Δz parameters. The black dashed contours are obtained when analytically marginalizing over the Δz parameters. Finally, the dotted lines mark the values of the fiducial cosmology used to generate the data. We can observe that the analytical and numerical marginalizations return nearly identical posteriors.

marginalization method yields virtually the same constraints on the cosmological parameters as the brute-force marginalization, in spite of the significantly broader posterior contours. This implicit validates the approximation that a first-order expansion of the theory data vector with respect to a change in redshift distribution is sufficient over a conservative range of calibration priors.

Finally, the calibration plan for Stage-IV surveys like LSST (LSST Dark Energy Science Collaboration 2012) and Euclid (Spergel et al. 2015) describes a clear path towards improving the calibration of the $p(z)$ by a factor of 5 compared to e.g. DESY1. In order to verify the approximation in this regime, we repeat our analysis of LSST-like futuristic data while reducing the uncertainty in the $p(z)$ priors by a factor of 5. As we can see in Figs A2 and A3, the marginalizing over the $p(z)$ uncertainties analytically or numerically returns virtually indistinguishable posteriors even for much tight $p(z)$, this is true for

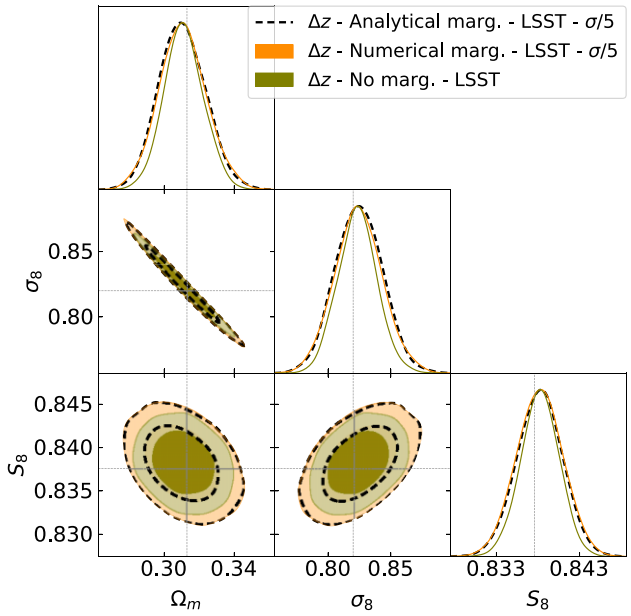

Figure A3. Marginalized posterior distributions for the combination of parameters Ω_m, σ_8, and S_8 obtained when considering the $p_\alpha(z)$ model for photometric uncertainties for LSST-like futuristic data. In this analysis, the priors on the $p(z)$ parameters were made five times tighter. The green contours correspond to the case where the $p_\alpha(z)$ parameters are fixed. The orange contours were obtained when numerically marginalizing over the $p_\alpha(z)$ parameters. The black dashed contours were obtained when analytically marginalizing over the $p_\alpha(z)$ parameters. Finally, the dotted lines mark the values of the fiducial cosmology used to generate the data. We can observe that the analytical and numerical marginalization return nearly identical posteriors.

the two parametrization considered for said uncertainty in this work. Moreover, we also observe that as we reduce the uncertainty on $p(z)$, the impact of marginalizing over its uncertainty is drastically reduced as one would expect. These results are consistent with the fact that the approximation becomes progressively more exact as the priors on the marginalized parameters become tighter. None the less, this exercise confirms that the methods developed in this work are ready to be applied to Stage-IV data once it becomes available.

This paper has been typeset from a TEX/LATEX file prepared by the author.