

Dark Energy Survey Year 3 results: Marginalisation over redshift distribution uncertainties using ranking of discrete realisations

Juan P. Cordero^{1*}, Ian Harrison^{2,1†}, Richard P. Rollins^{3,1}, G. M. Bernstein⁴, S. L. Bridle¹, A. Alarcon^{5,6,7}, O. Alves^{8,9}, A. Amon^{10,11,12}, F. Andrade-Oliveira^{9,13}, H. Camacho^{9,13}, A. Campos¹⁴, A. Choi¹⁵, J. DeRose^{16,17}, S. Dodelson¹⁴, K. Eckert⁴, T. F. Eifler^{18,19}, S. Everett¹⁶, X. Fang¹⁹, O. Friedrich^{20,21}, D. Gruen^{10,11,12}, R. A. Gruendl^{22,23}, W. G. Hartley²⁴, E. M. Huff¹⁸, E. Krause¹⁹, N. Kuropatkin²⁵, N. MacCrann²⁶, J. McCullough¹¹, J. Myles^{10,11,12}, S. Pandey⁴, M. Raveri²⁷, R. Rosenfeld^{13,28}, E. S. Rykoff^{11,12}, C. Sánchez⁴, J. Sánchez²⁵, I. Sevilla-Noarbe²⁹, E. Sheldon³⁰, M. Troxel³¹, R. Wechsler^{10,11,12}, B. Yanny²⁵, B. Yin¹⁴, Y. Zhang²⁵, M. Aguena^{13,32}, S. Allam²⁵, E. Bertin^{33,34}, D. Brooks³⁵, D. L. Burke^{11,12}, A. Carnero Rosell^{36,37}, M. Carrasco Kind^{22,23}, J. Carretero³⁸, F. J. Castander^{6,7}, R. Cawthon³⁹, M. Costanzi^{40,41}, L. da Costa^{13,32}, M. E. da Silva Pereira⁸, J. De Vicente²⁹, H. T. Diehl²⁵, J. Dietrich⁴², P. Doel³⁵, J. Elvin-Poole^{15,43}, I. Ferrero⁴⁴, B. Flaugher²⁵, P. Fosalba^{6,7}, J. Frieman^{25,27}, J. Garcia-Bellido⁴⁵, D. Gerdes^{8,46}, J. Gschwend^{13,32}, G. Gutierrez²⁵, S. Hinton⁴⁷, D. L. Hollowood¹⁶, K. Honscheid^{15,43}, B. Hoyle^{42,48,49}, D. James⁵⁰, K. Kuehn^{51,52}, O. Lahav³⁵, M. A. G. Maia^{13,32}, M. March⁴, F. Menanteau^{22,23}, R. Miquel^{38,53}, R. Morgan³⁹, J. Muir^{11,54}, A. Palmese^{25,27}, F. Paz-Chinchon^{23,55}, A. Pieres^{13,56}, A. Plazas Malagón⁵⁷, E. Sánchez²⁹, V. Scarpine²⁵, S. Serrano^{6,7}, M. Smith⁵⁸, M. Soares-Santos⁸, E. Suchyta⁵⁹, M. Swanson²³, G. Tarle⁸, D. Thomas⁶⁰, C. To^{10,11,12}, and T. N. Varga^{48,49} (DES Collaboration)

The authors' affiliations are shown at the end of this paper.

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Cosmological information from weak lensing surveys is maximised by sorting source galaxies into tomographic redshift subsamples. Any uncertainties on these redshift distributions must be correctly propagated into the cosmological results. We present `HYPERRANK`, a new method for marginalising over redshift distribution uncertainties, using discrete samples from the space of all possible redshift distributions, improving over simple parameterized models. In `HYPERRANK` the set of proposed redshift distributions is ranked according to a small (between one and four) number of summary values, which are then sampled along with other nuisance parameters and cosmological parameters in the Monte Carlo chain used for inference. This approach can be regarded as a general method for marginalising over discrete realisations of data vector variation with nuisance parameters, which can consequently be sampled separately from the main parameters of interest, allowing for increased computational efficiency. We focus on the case of weak lensing cosmic shear analyses and demonstrate our method using simulations made for the Dark Energy Survey (DES). We show the method can correctly and efficiently marginalise over a wide range of models for the redshift distribution uncertainty. Finally, we compare `HYPERRANK` to the common mean-shifting method of marginalising over redshift uncertainty, validating that this simpler model is sufficient for use in the DES Year 3 cosmology results presented in companion papers.

Key words: methods: numerical – galaxies: distances and redshifts – large-scale structure of Universe – gravitational lensing: weak

1 INTRODUCTION

As photometric galaxy surveys begin to map large fractions of the sky at deeper magnitudes, stringent control of systematic errors and uncertainties is required to take full advantage of the statistical power of such surveys. Combining measurements of weak lensing

* E-mail: juan.cordero@postgrad.manchester.ac.uk

† E-mail: ian.harrison-2@manchester.ac.uk

and spatial clustering of distant galaxies (and cross-correlations of these two signals as galaxy-galaxy lensing) has steadily become a very competitive probe of the expansion history of the Universe and its constituents (e.g. Dark Energy Survey Collaboration 2018; Hikage et al. 2019; Hamana et al. 2020; Heymans et al. 2021). The Dark Energy Survey Year 3 (DES-Y3 DES Collaboration 2021) results, of which this work forms a part, contains information from over 100 million galaxies. One of the key required measurements in such analyses is the line of sight distribution of both the galaxies for which the shapes are measured (the *source* sample) and the generally lower redshift galaxies used to trace the massive structures acting as lenses for the source sample (the *lens* sample). In an approach known as ‘tomography’ (Hu 1999) the source sample is sub-divided into different bins of distance, allowing us to further study the evolution of massive structure across cosmic time by observing how the lensing signal changes as a function of distance. Knowledge of the distance distribution to the source sample is a crucial ingredient in this. Of particular interest for modern cosmology, the statistical properties of dark matter structures as a function of cosmic time are a promising probe of dark energy.

Cosmological redshift z is the observable most commonly used as a proxy for the distances to both galaxy samples but the methods to estimate distance via redshifts often suffer from limitations which make this one of the most difficult uncertainties to adequately model for the cosmological analysis. Estimating the redshift with high accuracy using spectroscopy is prohibitively expensive in telescope time for the large numbers of galaxies required for cosmology using weak lensing and suffers from selection effects caused by the incompleteness at fainter magnitudes (e.g. Hartley et al. 2020). Photometric redshift (photo- z) methods instead estimate the redshift based on measurements of fluxes in a number of photometric bands, and present a viable alternative in terms of sky and redshift coverage and completeness, but suffer from relatively much larger uncertainties given the highly degenerate problem of estimating z based on wide band photometry. A wide range of photo- z methods are used to estimate redshifts from band magnitudes; see Schmidt et al. (2020) and references therein for a recent review.

Current galaxy surveys rely on a combination of spectroscopic and photometric redshifts, plus clustering patterns, to train, calibrate, and validate different methods. These methods can be broadly classified into three types, based on the information and ancillary data used to estimate redshift. (i) Template fitting methods (see Section 3.1 of Schmidt et al. 2020, for a review) which rely on finding the best-fit template redshift from a library of spectral energy distributions (SED) characterising a range of galaxy types. (ii) Machine learning based techniques (see Section 3.2 of Schmidt et al. 2020, for a review) which map the colour space into redshifts. While the range of approaches used is fairly wide, the general idea consists of using a training set of secure redshifts obtained using either spectroscopy or large sets of narrow-band filter photometric observations to train the algorithm. (iii) Using spatial correlation between galaxies and a set of tracers with secure redshift information to obtain additional constraints on redshift (often known as ‘clustering redshifts’). See the introductory sections of Gatti, Giannini et al. (2020); Cawthon et al. (2020) for recent reviews.

Irrespective of the chosen method there will be an irreducible uncertainty in the galaxy distances arising from the finite number of photons received in each band, the widths of the bands and our limited knowledge of true galaxy spectral energy distributions. Where galaxies are observed only in a few ($\sim 1 - 10$) photometric bands there are also fundamental degeneracies where two galaxies at very different redshifts can produce identical observed data. This

uncertainty must be propagated through to cosmological constraints. Galaxies are conventionally grouped into a small number (~ 5 for current experiments) of tomographic redshift bins. Cosmological observables of weak lensing, galaxy clustering and galaxy-galaxy lensing formed from each of these tomographic bin sub-samples are dependent on the number density distribution of the sources as a function of redshift within each bin, $n(z)$. If each individual galaxy’s redshift were known with perfect precision and accuracy, these $n(z)$ would be non-overlapping, and their shapes would follow the true distribution in redshift of galaxies which are really in these bins. However, in real cases where the one point summary statistic used for binning is noisy, biased, or both, the $n(z)$ within different tomographic bins acquire stretched tails which often overlap across the full redshift range of the survey.

In order to constrain cosmological parameters, expected weak lensing observables for a galaxy sample with the estimated $n(z)$ and in a given cosmology are computed and compared with the data. Monte Carlo methods are then used to map the posterior for cosmological model parameters and hence constrain our physical model for the Universe. In this inference process, uncertainties on the measured $n(z)$ for each tomographic bin are marginalised over, typically widening the uncertainties on the cosmological parameters of interest. Incorrectly quantifying the uncertainty on the $n(z)$ or incorrectly marginalising over it can significantly affect cosmological parameter estimation and model selection. Indeed Joudaki et al. (2020) have argued that the adoption of different models for calibration of redshift distributions and their uncertainties for weak lensing experiments can explain the observed apparently significant difference in cosmological parameters between different weak lensing experiments and Cosmic Microwave Background (CMB) experiments.

In this paper we introduce `HYPERRANK`, a new method which allows uncertainties in galaxy redshift distributions $n(z)$ to be propagated into Monte Carlo chains generating cosmological results. `HYPERRANK` takes as input a finite set of samples of $n(z)$ drawn from the distribution implied by the redshift calibration process. It maps these onto a low-dimensional space of continuous variables which the cosmology sampler can treat as free parameters. We test that `HYPERRANK` does this both *correctly*, in that the allowed uncertainty is fully explored, and *efficiently*, in that fewer likelihood evaluations are computed than in the case where an arbitrary choice of $n(z)$ realisation is made at each step. This approach allows for the inclusion of a much wider range of types of uncertainty on $n(z)$ to be used in cosmological inference than have been included in the majority of previous analyses.

In section 2 we review methods of quantifying uncertainties on the redshift distributions of galaxy samples used for cosmology, motivating the introduction of the new `HYPERRANK` method, which is then described in section 3, in both the simplest one-dimensional case and an extended multi-dimensional case. In section 4 we then perform tests of the performance of `HYPERRANK` on a simulated version of the DES-Y3 experiment. In section 4.2 we verify that in cases where redshift distribution uncertainty is known, `HYPERRANK` correctly marginalises over this uncertainty, for four representative models of the uncertainty. In section 4.3 we also show that the use of `HYPERRANK` to explore the uncertainties results in better performance (in terms of fewer numbers of Monte Carlo steps required) than random, un-ranked exploration of realisations of possible redshift distributions. We also explore the performance of a number of different choices of variables on which to perform the ranking and find, for our fiducial case, the number of discrete samples from the possible redshift distributions which are required for the cosmolog-

ical results to converge to those of a known case where continuous sampling is possible. Section 5 describes the application of `HYPERRANK` to the real DES-Y3 data, with the results presented in Amon et al. (2021). Finally in section 6 we discuss our results and conclude.

2 MARGINALISATION OF REDSHIFT UNCERTAINTY

In general, for an inference problem in which we have a model containing parameters of interest θ (such as the cosmological parameters) and a set of nuisance parameters α (such as parameters relating to redshift distribution uncertainty), we form posterior probability distributions:

$$P(\theta, \alpha | \mathbf{x}) \propto \mathcal{L}(\mathbf{x} | \theta, \alpha) P(\theta, \alpha) \quad (1)$$

where $\mathcal{L}(\mathbf{x} | \theta, \alpha)$ is the likelihood function for the data \mathbf{x} and $P(\theta, \alpha)$ is a prior probability distribution. When generating samples from the posterior with a Monte Carlo process, the nuisance parameters are typically sampled jointly with the parameters of interest and then marginalised over, providing a marginal posterior on the model parameters θ in which the uncertainty on α is accounted for. In the particular case of redshift distributions in cosmology analyses, a common approach is to provide a fiducial tomographic redshift distribution and characterise its uncertainty using the nuisance parameter of a shift Δz along the z -axis. A different parameter Δz_i is used for each tomographic bin, with each drawn from a Gaussian prior informed by observations and/or simulations. This approach is depicted in the upper panel of figure 1 and has been used in DES SV (Bonnett et al. 2016), DES Y1 (Hoyle et al. 2018), HSC (Hikage et al. 2019) and KiDS-1000 (Joachimi et al. 2021). However, whilst convenient and capturing the uncertainty in the mean of the redshift distributions, which is strongly correlated with cosmology, it is not physically well motivated and severely restricts the possible functional forms which a proposed $n(z)$ may take.

In contrast to the Δz approach, we may wish to consider alternatives which allow for a much wider range of uncertainty in the functional forms of the $n(z)$. It is possible to take a simulation-based approach, in which realisations for the possible $n(z)$ of a survey are generated by multiple realisations of mock versions of the survey created from independent patches of cosmological simulations. Alternatively, we may explicitly parameterise the $n(z)$ as a set of histogram bin heights $n(z_i)$ which give the counts of sources within a small redshift interval and try to infer these quantities from the data. This approach creates principled models of the joint probability distribution function for all of these bin heights given the photometric data available on the observed galaxies. This is most readily done as a Bayesian Hierarchical Model and has been recently advocated in Leistedt et al. (2016), Sánchez & Bernstein (2019) and Rau et al. (2020). Outputs from this procedure are samples from the joint posterior for all of the histogram bin heights which together make up the full shape of the $n(z)$. Each sample consists of a possible realisation of what the full $n(z)$ could look like, discretised as $n(z_i)$. An ideal approach would be to treat each of these $n(z_i)$ as a model parameter and jointly infer them with the cosmological model parameters before marginalisation. In reality this is impractical; the redshift resolution required to capture important features of the model which impact cosmological inference but are not convolved with broad redshift kernels, such as intrinsic alignments, would demand hundreds of additional nuisance parameters. Current implementations of galaxy survey analysis pipelines (such as that in CosmoSIS used for DES Zuntz et al. 2015) typically take

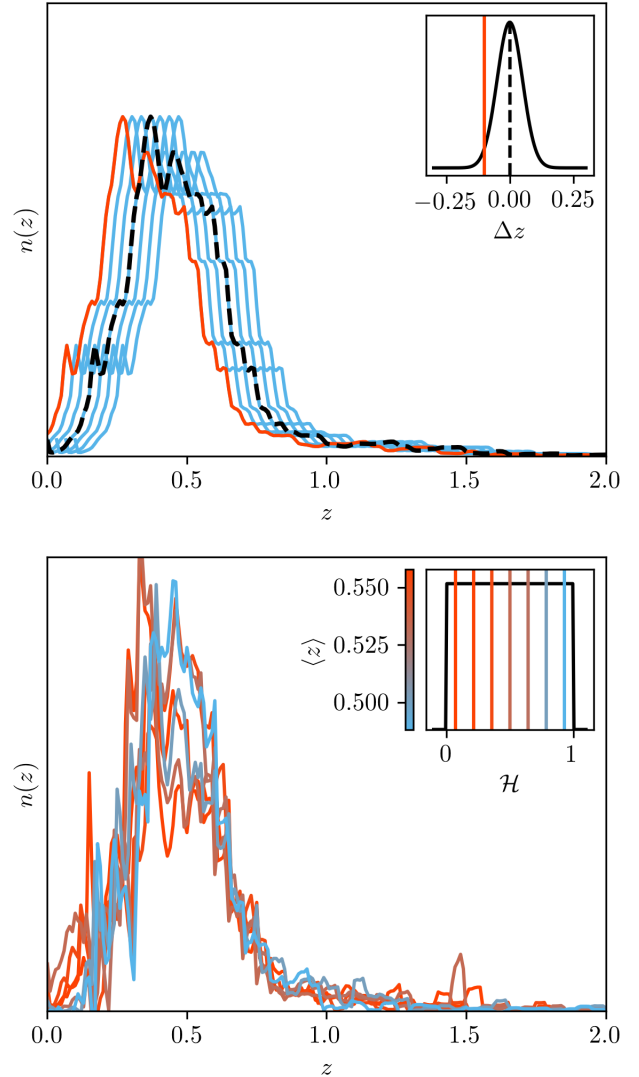


Figure 1. *Upper:* The Δz marginalisation scheme, where a fiducial redshift distribution (black dashed) is shifted horizontally at each Monte Carlo step by a value drawn from a Gaussian distribution (inset, with draw from the 2σ tail highlighted in red). *Lower:* Discrete realisations of possible $n(z)$ are shown with colours corresponding to the mean redshift of each realisation $\langle z \rangle$, which can be mapped to a ranking hyper-parameter \mathcal{H} which is then marginalised over on the Monte Carlo chain. Inset shows the uniform distribution for \mathcal{H} which is sampled from, and the centres of the regions corresponding to each coloured $n(z)$ realisation.

$\sim 1 - 10$ seconds per likelihood evaluation, meaning the addition of hundreds of parameters would mean the samplers used (MCMC such as `EMCEE` Foreman-Mackey et al. 2013 or nested sampling such as `MULTINEST` Feroz et al. 2009 or `POLYCHORD` Handley et al. 2015) could not map the full posterior in a timely manner. It should be noted that Hildebrandt et al. (2017) were able to run 750 MCMC chains in order to use a different bootstrap resampling realisation of their $n(z)$ each time, before combining these chains; we do not expect this to be feasible for the DES-Y3 pipeline. Other methods have also been proposed to address the uncertainty associated to large number of nuisance parameters, including Gaussian mixture models, which are flexible and may be analytically marginalised over (Hadzhiyska et al. 2020; Stölzner et al. 2021), and the use of

flat or Gaussian priors to characterise variations to sets of arbitrary functions used to evaluate a Gaussian likelihood (Taylor & Kitching 2010; Kitching & Taylor 2011)

Here we consider an alternative approach in which the set of samples from the $n(z)$ posterior, each consisting of a collection of histogram values for each tomographic bin, are generated outside (before) the cosmological parameter inference Markov chain. This set of realisations can then be used by choosing a new $n(z)$ in every likelihood evaluation within the cosmological parameter inference chain, allowing higher-order modes of uncertainty in $n(z)$ to be propagated into cosmological parameter constraints. It is important to note that the $n(z)$ realisations are drawn simultaneously for all tomographic bins, which also allows the propagation of uncertainty originating from correlations between tomographic bins.

A way of performing this analysis would be to randomly sample a different redshift distribution on each likelihood evaluation within the Monte Carlo chain. This has potential negative effects on the behaviour of Monte Carlo samplers which rely on the posterior function being a smooth function of the sampled parameters. A random approach can break the smoothness of the likelihood (as shown in figure 2) in the other parameter dimensions leading to unnecessarily high sample rejection rates, requiring large number of likelihood evaluations for convergence and potentially disrupting convergence criteria for different samplers.

Here, we present HYPERRANK, a way to overcome these computational limitations, whilst still exploring the space of uncertainty available from the discrete $n(z)$ realisations. In HYPERRANK we construct a mapping between the index of an ordered set of $n(z)$ realisations and a continuous parameter \mathcal{H} , such that the likelihood function $\mathcal{L}(\theta, \mathcal{H})$ is smooth on this new space and the prior $P(\mathcal{H})$ preserves an equal weighting of the $n(z)$ samples through assigning them to an evenly-spaced grid.

3 THE HYPER-RANKING METHOD

After a discrete set of realisations of tomographic bin redshift distributions $n_i(z)$ have been generated, we wish to correctly and efficiently marginalise over the uncertainty embodied by them, within a cosmological parameter inference Monte Carlo chain. We introduce the idea of HYPERRANK-ing in which the full set of realisations is mapped onto a small (in this work between one and four) number of parameters \mathcal{H} . The $n(z)$ realisations are ordered according to a set of descriptive values \mathbf{d} which are *a priori* expected to correlate strongly with values of the cosmological parameters of interest. This ordering preserves the tomographic nature of each realisation, meaning the sampling stage selects the set of all tomographic bins' distributions simultaneously, without mixing different realisations. The rank parameters \mathcal{H} become the nuisance parameters which are sampled (and subsequently marginalised over) in the cosmological analysis. Choosing descriptive values \mathbf{d} which correlate with the cosmological parameters of interest ensures the likelihood varies as smoothly as possible along each dimension of the rank parameters. The ranking parameters $\mathcal{H}_j = \mathcal{H}(\alpha_j)$ must also be such that realisations with similar descriptive values are mapped close to each other. Furthermore, the \mathcal{H}_j must be such that a uniform prior on \mathcal{H} preserves equal probability on all input $n(z)$ samples. We consider the cases below first in which we have one ranking parameter and then multiple ranking parameters. We choose to mainly use the mean redshift $\langle z \rangle$ and mean inverse comoving distance $\langle 1/\chi \rangle$ of each tomographic bin as descriptive values \mathbf{d} here, but emphasise the HYPERRANK method is not limited to these two options only.

We expect the choice of ranking method to only affect sampling efficiency and not the inferred parameter contours.

3.1 One dimensional case

We initially consider the case in which a single HYPERRANK parameter is used to rank all realisations. Since the mean redshift of the distribution $n(z)$ varies the overall amplitude of lensing expected for a given source galaxy sample, it is expected to correlate with the cosmological parameters of interest (here, the matter amplitude parameter S_8). We therefore consider a basic HYPERRANK approach in which there is only one descriptive parameter per realisation of the full $n(z)$ and it is based on the weighted mean redshift of a combination of tomographic bins,

$$\mathbf{d} = \frac{\sum w_i \langle z \rangle_i}{\sum w_i}, \quad (2)$$

where i is the index of each tomographic bin and w_i is the corresponding weight, which can embody (for instance) the number of assigned galaxies to each tomographic bin. The $n(z)$ realisations are then ranked according to their descriptive value \mathbf{d} and mapped to a continuous hyper-parameter $\mathcal{H} \in [0, 1)$, which is then sampled in the Monte Carlo chain. Each sampled value of \mathcal{H} corresponds to a stored $n(z)$ realisation which is then used in the likelihood evaluation. This approach is demonstrated in the lower panel of Figure 1 which shows a small sample of $n(z)$ realisations coloured according to their mean redshift and assigned a range of \mathcal{H} values depending on their ranked position.

An alternative set of descriptive values are the mean inverse comoving distance of sources, $\langle 1/\chi \rangle$. The correlation of this quantity with cosmological posterior value can be motivated by its relation to the lensing efficiency functions used in the calculation of the shear power spectrum, which can be written as,

$$P_\kappa(\ell) = \frac{9H_0^4 \Omega_m^2}{4c^4} \int_0^{\chi_H} g^2(\chi) \frac{P_\delta(\ell/\chi; \chi)}{a^2(\chi)} d\chi, \quad (3)$$

where χ_H , $a(\chi)$ and P_δ are the comoving horizon, scale factor and matter power spectrum, respectively and the lensing efficiency $g(\chi)$ at comoving distance χ is defined as:

$$g(\chi) = \int_\chi^{\chi_H} n(\chi') \frac{\chi' - \chi}{\chi'} d\chi', \quad (4)$$

and depends on the comoving distance distribution $n(\chi)$ of sources, or equivalently their redshift distribution $n(z)$. By evaluating at $\chi = 0$ and differentiating the above definition for the lensing efficiency we obtain

$$g(\chi)|_{\chi=0} = 1 \quad (5)$$

$$g'(\chi)|_{\chi=0} = -\langle 1/\chi \rangle n, \quad (6)$$

where $g'(\chi) = dg/d\chi$, which are boundary conditions for the lensing efficiency functions hence control their overall shape. See Tessore & Harrison (2020) for discussions of the importance of constraining $g'(\chi)$ in weak lensing studies.

The mapping of distributions is not invariant to the choice of ordering being the mean redshift or the mean inverse comoving distance. In one dimension, both are examples of ranking parameters capable of providing the smooth likelihood necessary for efficient mapping of the posterior (as can be seen in Figure 2), as well as correctly including the space of uncertainty spanned by the provided set of $n(z)$ realisations. In section 4 below we consider only the mean redshift ranking for the one dimensional case, but observed




Figure 2. The log posterior for the HYPERRANK parameter of a single tomographic bin when holding all other parameters fixed, contrasting the cases of random ranking (which gives no smooth posterior for the sampler to explore) with mean redshift ranking (which does give a smooth posterior surface).

comparable performance for the inverse comoving distance ranking in our tests.

3.2 Multi-dimensional case

While the one-dimensional approach presents a clean and simple strategy to arrange and select realisations for each likelihood evaluation, it doesn't prevent cases where two realisations with very different descriptive values are assigned a similar rank—e.g. two realisations have very distinct $\langle z \rangle_i$ in individual redshift bins, but similar when averaged over bins as per equation (2). Indeed in our initial tests with DES-Y3 simulations it was found this was often the case, leading to realisations ranked closely by a single mean redshift parameter having significantly different posterior values, hence leading to poor efficiency in the cosmology chains. To address this we describe a generalisation to rank distributions using multiple dimensions, which allows to use more than one descriptive parameter \mathbf{d} to assign the proposal $n(z)$ realisations to a space of hyper-parameters \mathcal{H} . Matching of the number of descriptive values and the dimensionality of the redshift distributions (e.g. number of tomographic bins) is not a requirement, and we find here the best performance is achieved when this is not the case.

Each of the N_p proposals for $n(z)$ is assigned a position in a uniform multi-dimensional grid, \mathbf{u} , according to a set of N_d descriptive values $\mathbf{d} = d_1, \dots, d_{N_d}$. This grid is contained inside a N_d -dimensional unit hyper-cube, and the continuous parameters $\mathcal{H}_j \in [0, 1)^{N_d}$ are sampled in the Monte Carlo chain. For each \mathcal{H} value chosen by the sampler, the method returns the closest \mathcal{H}_i in the grid, which has been assigned to one of the N_p $n(z)$ realisations.

We now need to consider how to preserve the notion of ordering the set of $n(z)$ by descriptive values in this multi-dimensional space, preserving the notion of a ‘neighbourhood’ where realisations with similar descriptive properties are grouped close together. One approach to find the optimal relative positions is to use the solution to the *Linear Sum Assignment Problem* (e.g. Burkard & Derigs 1980). Given a set of N_p workers (points in the descriptive value space) we want to find an assignment to N_p fixed jobs (i.e. fixed grid positions in the unit hyper-cube) such that the sum of the

cost to assign each worker to one and only one job (the distance from descriptive value space to hyper-cube position) is minimised:

$$\min \sum C_{ij} X_{ij},$$

where C_{ij} is the cost matrix of assigning each sample \mathbf{d}_i to each point \mathbf{u}_j of the grid, and X_{ij} is a binary matrix indicating which position is assigned to each set of descriptive values. If we use an Euclidean distance metric such that $C_{ij} = |\mathbf{d}_i - \mathbf{u}_j|^2$, the resultant assignment minimises the total distance moved by the points to the positions on the grid ensuring that any notion of neighbourhood between points in the original space of descriptive parameters is preserved in their new unit hyper-cube grid positions. We implement this technique by first linearly rescaling the \mathbf{d}_i so that they span a unit hypercube. Figure 3 shows the resultant 2-dimensional assignment for a set of 500 realisations, each comprised of a set of four tomographic bins, using as descriptive parameters the mean redshifts of tomographic bins 1 and 4, arranged in a 25×20 map. Because of the finite number of available realisations the use of additional dimensions can quickly have the undesired effect of reducing the amount of realisations available with which to fill each direction of the multidimensional grid. This can result in the exacerbation of the convergence problem, with few available samples creating large jumps in posterior as a function of the \mathcal{H} parameters. For example with 4096 realisations, double the grid size is available with $N_d = 3$ dimensions compared to $N_d = 4$.

In the case of $N_d = 1$, where a single characteristic value describes each realisation and the arrangement of points is done over a grid in the interval $[0, 1)$, the optimal distribution is the one which ranks the points in order, corresponding to the case described in Section 3.1. Analogous to the one dimensional case, we propose the use of mean redshift $\langle z \rangle$ or mean inverse comoving distance $\langle 1/\chi \rangle$ of the individual tomographic bins as sources of descriptive values to map the realisations to the hyper-cube.

Ideally, the dimensionality N_d of the hyper-ranked space is low enough to maintain an efficient cosmological sampler, but high enough that the variation in the log posterior probability from Eq. (1) in small regions of \mathcal{H} is $\ll 1$. This would allow any sampling process to smoothly traverse the full space of all $n(z)$ variations that influence the parameters of interest.

We can optimize the reduction of the nuisance-variable vector α (e.g. all of the freedom of $n(z)$) into a lower dimensional hyper-space by using the Karhunen-Loève (KL) transformation. When the observational data vector D has a Gaussian likelihood with covariance matrix C_D and mean value $\hat{D}(\theta, \alpha)$, we find the eigenvectors e^k of the matrix

$$\left(\frac{\partial D}{\partial \alpha} \right)^T C_D^{-1} \left(\frac{\partial D}{\partial \alpha} \right). \quad (7)$$

where the derivatives are taken about some reference values of θ and α . The best choice of HYPERRANK descriptive values (d_1, d_2, \dots, d_K) will be to order the eigenvectors by decreasing eigenvalues, and assign $d_k = \alpha e^k$ for each input sample. Successive d_k values have decreasing influence on the cosmological model. The sum of the eigenvalues at $k > K$ then describes the ‘roughness’ of the log-posterior in the \mathcal{H} space. Using this principal component analysis (PCA)-style approach, we can choose the first K components of the decomposition as descriptive values to inform the ranking map and assign each component to one HYPERRANK parameter each.

The main caveat is that this approach defines a set of descriptive values which are optimal only near the reference cosmology chosen to compute the KL components. While ideally one would want to

use a large number of dimensions to help construct a space where the posterior is as smooth as possible, this comes at the expense of having to construct a grid with a low number of points per dimension, if the number of input samples of $n(z)$ is held fixed. This can result in a noisy posterior as a function of the hyper-parameter \mathcal{H} if a given dimension of \mathcal{H} is sparsely sampled and has large steps between samples. While a large number of realisations can help construct a grid with a reasonably large number of realisations per side of the grid, the method to solve the linear sum assignment problem scales as $O(N_p^3)$, which quickly becomes unmanageable. In section 4.3 we explore the effects the dimension of the ranking and choice of descriptive value have on sampling efficiency, testing the mean redshift, inverse comoving distance, and KL approaches with 3 components each.

4 TESTS ON SIMULATIONS

We now test the `HYPERRANK` method for marginalising over redshift distribution uncertainty and explore its configuration, with the target of using it for the weak lensing source redshift distributions in the DES Year 3 cosmological analysis. We investigate the `HYPERRANK` method's ability to marginalise over the $n(z)$ uncertainty:

- *correctly*, in that it proportionately explores the space of possible $n(z)$ represented by the discrete realisations which are provided as an input.
- *efficiently*, in that as few likelihood evaluations as practically possible are required before the Monte Carlo process converges to the posterior.

We test the correctness by comparing the recovered posteriors on the $S_8 = \sigma_8 \sqrt{\Omega_m}/0.3$ cosmological parameter obtained from a cosmological inference pipeline. We generate sets of $n(z)$ samples using a number of well defined procedures in which the method for generating realisations involves drawing a Δz shift from a known analytic distribution. We then run analyses using `HYPERRANK` to marginalise over these uncertainties and compare the results to a set of chains in which the known analytic distributions from which the Δz were drawn are used again to marginalise over the nuisance parameter. Hence we verify that, in the case where discrete samples represent a model for uncertainty on $n(z)$, the use of `HYPERRANK` correctly explores this uncertainty. The tests show that `HYPERRANK` is capable of correctly marginalising over redshift distribution uncertainties in cases where a correct and simple model for them is known, without making assumptions on the form of the uncertainty model. This model-agnosticism represents an advantage in the case of real experiments, where it may not be known *a priori* if one or any of the simple models is adequate for obtaining small, unbiased posteriors.

We also compare the results from analyses using `HYPERRANK` to ones in which discrete $n(z)$ realisations but no ranking (or equivalently random ranking) are used, showing that the imposition of the `HYPERRANK` ranking does not bias or unduly constrain the cosmological parameter space explored.

We test the improvement in computational efficiency gained from using `HYPERRANK` by comparing one- and multi-dimensional implementations of `HYPERRANK` to a mode in which no ranking is performed and at each likelihood evaluation an $n(z)$ is chosen from the available realisations at random.

Finally, we also test the convergence of `HYPERRANK` for the configuration required for DES-Y3 cosmology, finding the number of $n(z)$ realisations which are required before systematic errors

on the cosmology parameters from the discreteness introduced by `HYPERRANK` become negligible.

Throughout these tests we use the DES-Y3 modelling choices, likelihood and pipeline software and configuration, which are described in detail in Amon et al. (2021); Secco, Samuroff et al. (2021). We only consider cosmic shear in our data vector, which reduces the dimensionality of the space of parameters to be sampled in the MC inference and enhances the effect of redshift systematics in the source sample. Nevertheless, this method can be applied when using cosmic shear in a full 3x2pt analysis including galaxy clustering and galaxy-galaxy lensing and can also be used to marginalise over systematic uncertainties of the *lens* in addition to the *source* samples described here.

4.1 Generation of fiducial redshift distribution

Here we briefly describe the method by which the cosmic shear data vector and fiducial $n(z)$ used in our tests were generated. The methodologies and simulations are described in detail in Myles, Alarcon et al. (2020), Gatti, Giannini et al. (2020) and DeRose et al. (2021).

4.1.1 Buzzard simulation

The `BUZZARD` simulations (DeRose et al. 2021) are a set of mock DES-Y3 surveys created from a suite of dark-matter N-body simulations using a memory-optimised version of `L-GADGET2` (Springel 2005). Galaxies and their main morphological properties are added using `ADDGALS` (DeRose et al. 2019), matching projected clustering statistics and color-magnitude relations observed in the Sloan Digital Sky Survey Main Galaxy Samples (SDSS MGS as described in Blanton et al. 2005; Abazajian et al. 2009). DES *ugriz* and VISTA *JHK* photometry is obtained from the simulated spectral energy distributions generated by `ADDGALS`.

4.1.2 SOMPZ redshift distributions

The simulated photometry catalogues from `BUZZARD` constitute the primary dataset to construct the fiducial $n(z)$ for our tests, using the `SOMPZ` method (fully described in Myles, Alarcon et al. 2020). This method makes use of three sets of observations: the full DES-Y3 wide field sample, the DES-Y3 Deep Fields (Hartley, Choi et al. 2020) sample, and compilation of spectroscopic redshift surveys. Galaxies from the wide sample are grouped into *phenotypes* using the Self-Organised Maps (SOM) method of dimensional reduction (see e.g. Myles, Alarcon et al. 2020; Wright et al. 2020; Masters et al. 2015). The `BALROG` machinery (which injects synthetic sources into DES data and recovers their properties, see Everett et al. 2020) is then used to quantify the probability of a given Deep Fields galaxy appearing to have a given phenotype when observed in the wide field. A second SOM dimensional reduction is then applied to the Deep Fields galaxy observations, with the spectroscopic sample used to characterise the true redshift distribution for each deep phenotype. In this way, information can effectively pass from the small, limited spectroscopic sample to the much larger wide sample through the intermediary of the deep sample.

In addition to this method of creating a best-estimate fiducial redshift distribution, we further consider realisations of possible $n(z)$ inferred from the simulated data using the method of Myles, Alarcon et al. (2020); Sánchez et al. (2020); Sánchez & Bernstein (2019). This applies a three-step Dirichlet (3sDir) sampling to model

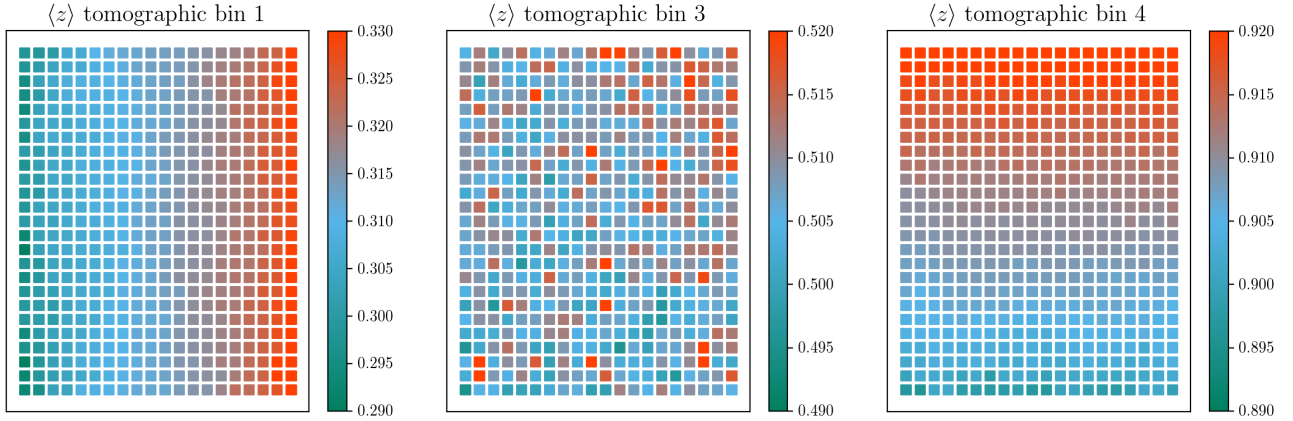


Figure 3. 25×20 ranking map generated using the mean redshift for tomographic bins 1, 2 and 4 from a set of 500 redshift distributions. Panels show distributions located in the same positions, but the color scale shows the mean redshift for the corresponding tomographic bin. It can be seen that the mapping scheme permits realisations to remain close to other realisations with similar descriptive values used for the mapping, and has a smooth variation in the directions of the hyperparameters mapped to each dimension of the grid. The arrangement does not necessarily result in a smooth ordering of all tomographic bins, as can be seen from the middle panel where the mean redshift from a bin not used of the mapping is displayed.

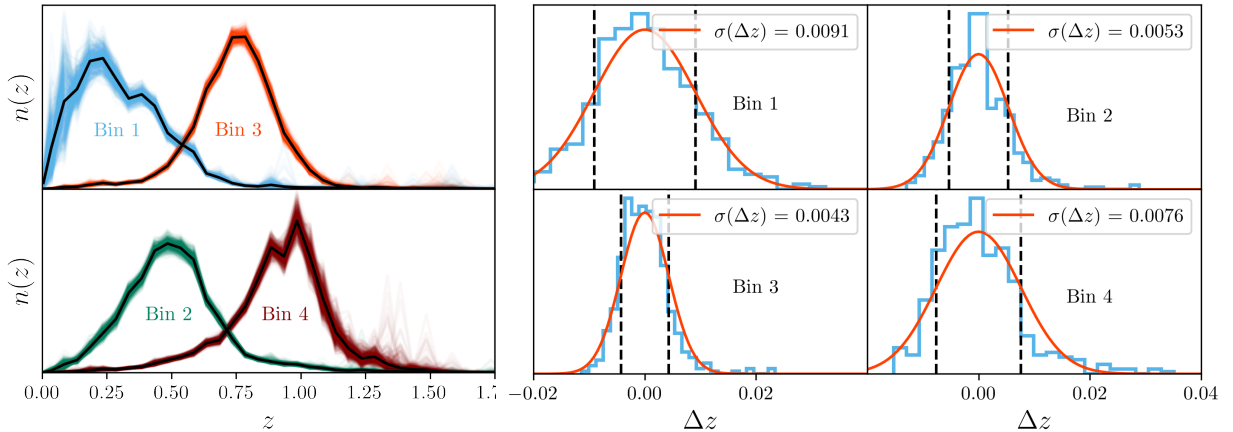


Figure 4. The BUZZARD redshift distributions. *Left:* The black lines show the redshift distribution $n_{\text{Fid}}(z)$ for each of the four redshift bins, averaged over all realisations. The light blue, red, green and brown lines show the full set of realisations for redshift bins 1 through 4, respectively, depicting the potential differences between independent samples of $n(z)$ posterior, and their peculiarities at the histogram level. *Right:* Histogram of mean redshift for each of the four tomographic bins, computed from the ensemble of distributions on the left panel. Solid orange line traces the Gaussian fit to the histogram, described by the width $\sigma(\Delta z)$ above each panel.

the uncertainties on $n(z)$ histogram bin heights from sources including shot noise, sample variance, photometric calibration uncertainty, and method errors. We use a set of 500 realisations generated this way, noting that samples are drawn jointly for all four tomographic redshift bins. The resulting estimated redshift distributions for BUZZARD are shown as the coloured lines in the left panel of figure 4. The fiducial realisation $n_{\text{Fid}}(z)$ is obtained from averaging the 500 realisations at the histogram level and re-normalising, and are shown as the black solid lines in the left panel of figure 4.

4.2 Exploration of uncertainties

As a supplement to these full SOMPZ + 3sDir realisations of the BUZZARD $n(z)$ for testing, we also now take the fiducial $n(z)$ and construct sets of realisations of potential $n(z)$ using simple parametric models for the uncertainty. We use analytic distributions to generate sets of mean redshift shifts Δz for each uncertainty model.

We then compare the posteriors on cosmological parameters (and the effective Δz nuisance parameters) recovered by two chains:

- a chain in which HYPERRANK takes these realisations as an input set of proposed $n(z)$
- a chain with Δz nuisance parameter marginalisation, using as a likelihood the same analytic distribution which was used to generate the realisations

To perform our sampling we use the MULTINEST sampler, with 500 live points, `tolerance` = 0.3, and `efficiency` = 0.01. We follow the setup for the DES-Y3 cosmic shear analysis described by Amon et al. (2021); Secco, Samuroff et al. (2021) in terms of angular scale cuts, tomographic redshift binning and modelling choices and marginalisation over other nuisance parameters such as shear calibration biases or Intrinsic Alignment model parameters. In most of the tests, and unless explicitly noted, we use the default

three dimensional `HYPERRANK` configuration described at the start of section 5.

We will describe each test in the following section, as well as the results for each one presented in figures 5 to 8. In each of these figures, the top panels show the one dimensional posterior constraints recovered on S_8 and the means of the redshift distributions in each tomographic bin $\langle z \rangle_i$. The lower panels in each figure show the two dimensional posterior constraints on these parameters. Dashed grey lines correspond to mean values of the fiducial redshift distribution in each tomographic bin, and in the S_8 panel to the values inferred from a chain run without marginalisation over redshift nuisance parameters.

4.2.1 Gaussian distributions for Δz

We begin with the simple error model in which the Δz approach used in other analyses and described above is the correct one. Within each tomographic bin we draw values of Δz from a Gaussian distribution with width $\sigma(\Delta z)$. Realisations for $n(z)$ are then generated by shifting the fiducial $n_{\text{fid}}(z)$ along the redshift axis by the drawn Δz . In order to assess performance and convergence we test this for several different levels of uncertainty, with the $\sigma(\Delta z)$ being modified by a multiplicative factor f_m . For our fiducial $\sigma(\Delta z)$ we use the values appropriate for DES-Y3 provided by BUZZARD (see right panel of figure 4). We then use values of $f_m = \{1, 5, 10\}$. We then run the full cosmological parameter estimation pipeline on the simulated data vector using these redshift distributions, once marginalising over the uncertainties using the Gaussian Δz method and once using the `HYPERRANK` method on the set of realisations.

The results of this test are shown in figure 5. In the upper panel we show the one dimensional error bars recovered on $S_8 \equiv \sigma_8 (\Omega_m/0.3)^{0.5}$ and the means of the redshift distributions in the four tomographic bins $\langle z \rangle_i$. Dashed grey lines correspond to mean values of the fiducial redshift distribution in each tomographic bin, and in the S_8 panel the 1σ region inferred from a chain run without marginalisation over redshift nuisance parameters is also shown. In the lower panel we also show the two dimensional posteriors for combinations of these parameters. We see that `HYPERRANK` gives posteriors consistent with those obtained using the standard Δz marginalisation approach. While at first glance this is a trivial example, it shows that the method is, at the very least, able to recover the same effects of redshift uncertainty when samples describe the same type of uncertainty we typically describe by means of a Δz nuisance parameter.

4.2.2 Non-Gaussian distributions for Δz

Modelling the distribution of Δz for each tomographic redshift bin as a Gaussian is a simple model choice which may not be an adequate representation of the true range and correlation structure of the Δz nuisance parameters, potentially resulting in a biased posterior and under/over-estimated uncertainties. In the right panel of figure 4 we show histograms of the Δz between the fiducial $n(z)$ and the 500 realisations generated using the full uncertainty model. These show appreciable non-Gaussianity, with skews and heavy tails which can be accentuated by the hard boundary at $z = 0$ for all distributions, especially tomographic bins at low redshift. We investigate the impact of the non-Gaussianity in the distribution of Δz by sampling Δz values from a highly skewed Gamma distribution:

$$f(\Delta z; k, \theta) = \frac{\Delta z^{k-1} e^{-\Delta z/\theta}}{\theta^k \Gamma(k)} \quad (8)$$

(where $\Gamma(k)$ is the integral Gamma function evaluated at k) to shift our fiducial distribution $n_{\text{fid}}(z)$. We use scale parameters θ such that the $\sigma(\Delta z)$ for each tomographic bin is equal to that of the prior with the largest uncertainty in section 4.2.1 ($f_m = 10$, $\sigma(\Delta z) \sim 0.05$). We fix the shape parameter k of the Gamma distribution to a set of values $k = 1, 2, 3$ to ensure the distribution of mean shifts of all tomographic bins have a positive skewness with a long tail to high values, and to explore the effect of different degrees of non-Gaussianity. The distribution of values is then centred so that the mean shift value is equal to zero, which generates a set of Gamma distributed Δz with the same variance and mean to that of the $f_m = 10$ prior, but with a skewness that cannot be captured by the use of a standard Gaussian prior. We then again run two chains, one marginalising over redshift uncertainty using the Gamma function Δz model, and one using `HYPERRANK` on the generated realisations.

The result of these chains is shown in figure 6. The differences on the S_8 parameter remain comparable to the typical dispersion seen for this number of distributions. As in figure 5 and figure 7, in the case of the $\langle z \rangle$ sampled values, small differences appear between `HYPERRANK` and Δz chains, but they are distributed very similarly as seen in the lower panels of figure 6, deviating in the same way from the reference values of the $n_{\text{fid}}(z)$ distribution and no marginalisation run. One aspect of the way Δz values are reported by `COSMOSIS` can be responsible for the differences, as the $\langle z \rangle$ values shown here are the sampled Δz plus the means of the fiducial distribution $n_{\text{fid}}(z)$. Because of the cut imposed at $z = 0$, this can result in a slightly inaccurate mean redshift value being computed here.

4.2.3 Correlations between tomographic bins

Another aspect of uncertainty the simplest Δz scheme does not directly account for is the potential correlation between the uncertainty from different tomographic bins (though see Appendix A of Hoyle et al. 2018, in which the diagonal elements of the covariance matrix are inflated to account for potential off-diagonal elements). Since each tomographic bin is shifted independently, combinations of Δz values which would not be expected to appear in multiple realisations of the survey or photo- z analysis are equally sampled. In addition to this, the use of a single fiducial shifted $n(z)$ blurs the potential effect of correlation at the histogram bin level. Correlation can come from the binning of galaxies and from how the shapes of the distributions and their moments can change when galaxies are re-assigned to another histogram or tomographic bin in a different realisation of a photo- z analysis. Depending on the nature of the colour-redshift degeneracy, correlation can also appear between non contiguous tomographic bins.

In this case, the standard Δz scheme can not be expected to preserve the effects of such correlations, as the set of $N_{\text{tom}} \Delta z$ nuisance parameters are sampled independently from their corresponding priors in the Monte Carlo chain. By contrast, drawing a value of the `HYPERRANK` parameter(s) in a chain jointly specifies the $n(z)$ to be used in all tomographic bins and preserves these correlations, which can potentially lead to tighter contours on the cosmological parameters since the space of Δz values is restricted to those allowed by the samples. Depending on the sign of the correlation, this can also result on a shift of the contours if the Δz values favour a combination of high or low mean redshift only (positive correlation), instead of a combination of low and high mean redshift (negative correlation). To explore the potential effects of these correlations at the tomographic bin level on inferred cosmological parameters, we generate three sets of mean-shifted realisations of

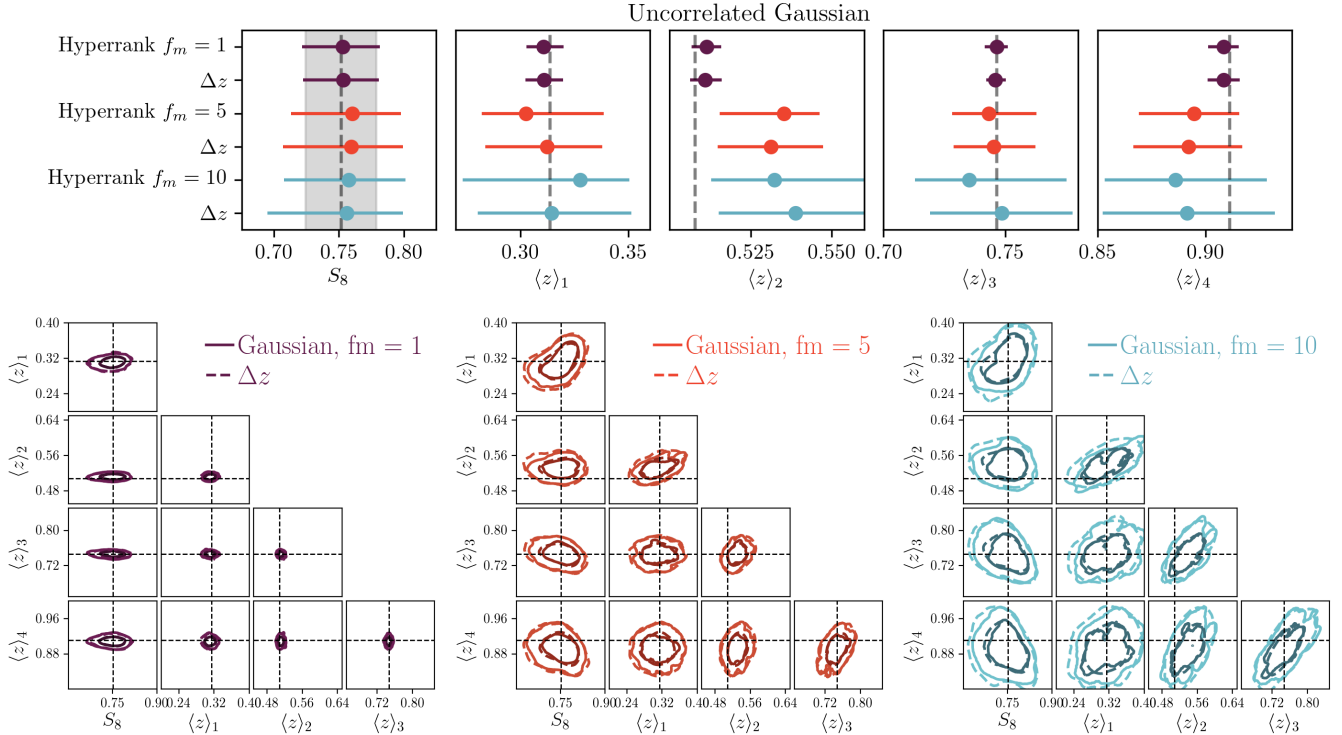


Figure 5. *Upper:* Marginalised one- σ confidence regions for the S_8 , and $\langle z \rangle_i$ parameters for the uncorrelated Gaussian model of $n(z)$ uncertainty described in section 4.2.1. Different f_m values refer to different overall amplitudes of uncertainty, and for each value of f_m we show both the posterior from chains using HYPERRANK and the Δz marginalisation schemes. *Lower:* Two dimensional posteriors on the same parameters. Dashed grey lines correspond to mean values of the fiducial redshift distribution in each tomographic bin, and in the S_8 panel the 1σ region inferred from a chain run without marginalisation over redshift nuisance parameters is also shown.

the fiducial BUZZARD $n_{\text{fid}}(z)$ by values of Δz sampled from a covariance matrix with increasing correlation between tomographic bin pairs (1,2) and (3,4). This is intended to be a simple model of leakage of galaxies between adjacent tomographic bins, with more complicated models for bin correlation also possible. We generate the samples so their Pearson correlation coefficients take the values $\rho = \{0.25, 0.5, 0.75\}$, and employ the same coefficient for both bin pairs while leaving all other bin pairs uncorrelated ($\rho = 0$). To better visualise the effects of these correlations once again we use an amplified $\sigma(\Delta z)$ prior to describe the diagonal of the covariance matrix, equal to the $f_m = 10$ prior described in section 4.2.1. We again run two chains, one in which a correlated Gaussian Δz marginalisation is used by drawing values from a correlated prior with the same correlation matrix to that used to generate the proposal samples, and one in which a HYPERRANK marginalisation is used.

The result of this test is shown in figure 7. We can again see that HYPERRANK correctly recovers the uncertainty in Δz and S_8 represented by the two dimensional posteriors.

4.2.4 Higher order modes of uncertainty

The above tests show that HYPERRANK is capable of correctly marginalising over redshift distribution uncertainties in cases where a correct and simple model for them is known. Finally, in this section, we use a set of realisations of $n(z)$ which represent a fully flexible model of the uncertainty in $n(z)$, following the approach of Myles, Alarcon et al. (2020) and summarised in section 4.1.2, as applied to the BUZZARD simulation.

As above for the cases of different values of f_m , we apply a procedure to these realisations to artificially increase the level of uncertainty they represent. Starting from the set of 500 realisations, we amplify the difference between each of the $n(z_i)$ values and the value of the fiducial distribution, $n^f(z_i)$, such that $n'(z_i) = n(z_i) + \lambda [n(z_i) - n^f(z_i)]$. Hence, we decide to use the typical dispersion values found for 500 realisations for that amplification as the reference to evaluate the contours obtained with HYPERRANK. For this test we generate three sets of distributions: one with no amplification, $\lambda = 0$; and two with amplified peculiarities, $\lambda = \{1.5, 3\}$. While the average $n(z)$ obtained from the amplified realisations remains unaltered, this procedure can result in a slightly wider equivalent Gaussian prior $\sigma(\Delta z)$ to those of the un-amplified realisations. Thus, we also obtain the $\sigma(\Delta z)$ values for each set of distributions and use them to compare HYPERRANK to the standard Δz marginalisation.

The results from this test are shown in figure 8. As can be seen, the HYPERRANK and Δz chains again recover highly consistent contours on the S_8 and $\langle z \rangle$ parameters. For the $\lambda = 0$ case the Δz posterior on S_8 is 2.6% narrower than the HYPERRANK one, for $\lambda = 1.5$ the Δz posteriors are 16% wider, and for $\lambda = 3.0$ the Δz posterior is 18% wider. This follows the idea that HYPERRANK is capable of better modelling of these more complex uncertainties, but that in the $\lambda = 0$ (DES-Y3-like) regime, Δz is an acceptable approximation.

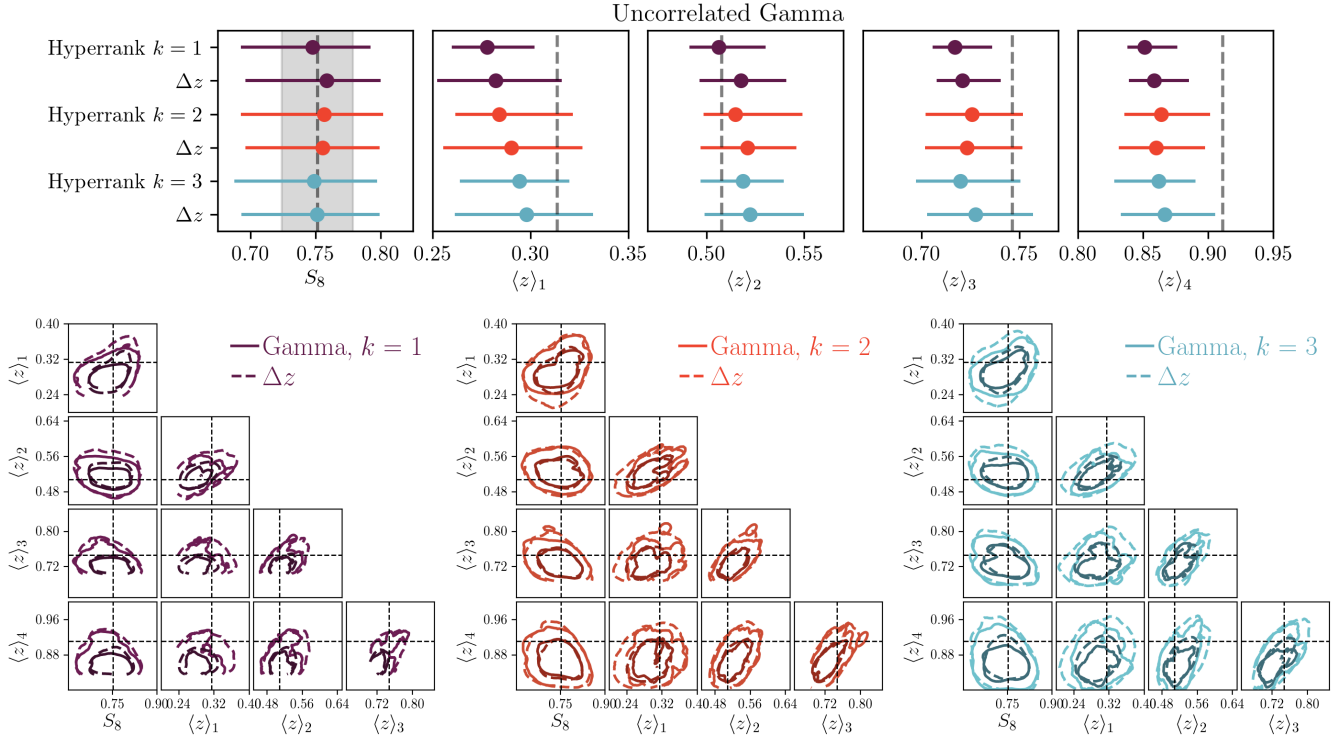


Figure 6. *Upper:* Marginalised one- σ confidence regions for the S_8 , and $\langle z \rangle_i$ parameters for the uncorrelated Gamma distribution model of $n(z)$ uncertainty described in section 4.2.2. Different k values refer to different amounts of skewness in the uncertainty distributions and for each value of k we show both the posterior from chains using HYPERRANK and the Δz marginalisation schemes. *Lower:* Two dimensional posteriors on the same parameters. Dashed grey lines correspond to mean values of the fiducial redshift distribution in each tomographic bin, and in the S_8 panel the 1σ region inferred from a chain run without marginalisation over redshift nuisance parameters is also shown.

4.3 Sampling efficiency and ranking mode

As well as the correct exploration of the uncertainties, we also wish to see the effect of the HYPERRANK procedure on the efficiency of mapping the posterior of cosmological and nuisance parameters. For a randomly sampled set of distributions the likelihood is not a smooth function of the parameters being sampled (see figure 2). Therefore, the parameter space volume cannot be sampled consistently in higher likelihood regions since there is no correlation between the sampled nuisance parameter and cosmology posterior. Any proposal step in the Monte Carlo algorithm typically does not have the intended effect, since proposed jumps in the redshift nuisance parameters are now across a random, discontinuous likelihood. This leads to the sampler requiring many more likelihood evaluations to find new samples of the posterior. We define sampling efficiency η as the number of replacements (samples of the posterior) made by MULTINEST over the total number of likelihood evaluations required for convergence, with higher η representing better performance.

We test the different mapping schemes described in section 3 comparing 1D and 3D $\langle z \rangle$, 3D $\langle 1/\chi \rangle$ and a KL approach where the first $K=3$ components are used. We compare the sampling efficiency between them and against a naive sampling where realisations are chosen at random on each likelihood evaluation. To reduce the effect of sampling noise due to the stochastic nature of the sampler, we repeat each run five times with different initial random seeds for the sampler.

Figure 9 shows the sampling efficiencies η at different f_m values as a function of different choices for descriptive values \mathbf{d} , all

compared to the average efficiency from five runs obtained using the Δz approach (dashed horizontal lines). The different colours used represent this test for different values of the f_m parameter. In all cases it is clear that the more complex choices of \mathbf{d} using multiple dimensions are more efficient at exploring the space of uncertainties, with 3D $\langle z \rangle$ and 3D $\langle 1/\chi \rangle$ showing better performance at all f_m values. This is expected since the addition of more dimensions helps breaking the degeneracy of the posterior values present when a single parameter is used and all the information of the $n(z)$ realisations is compressed into a single value.

The KL approach, also tested in three dimensions, provides an improvement over random and 1D sampling, but does not reach the same levels of efficiency for methods of equal dimensionality. with respect to a reference data vector obtained at a fixed cosmology, and the relative importance of each $n(z)$ element can change as the sampler moves in cosmology space.

Perhaps one surprising result occurs when comparing the random approach against 1D $\langle z \rangle$ in the un-amplified case ($f_m = 1$), in which the former appears $\sim 10\%$ more efficient. We believe this is caused by the relatively small contribution of $n(z)$ uncertainty to the posterior in the $f_m = 1$ regime, as all realisations have very similar mean values across all tomographic bins. This can lead to a very small change of smoothness of the posterior at a fixed cosmology when moving from a random ordering to a 1D ordering, resulting in similar efficiencies. While we do not show the effect of additional dimensions for a similar type of descriptive value \mathbf{d} (i.e. 4D $\langle z \rangle$), some test runs suggest their efficiency is not noticeably better than a 3D approach, at the expense of noisier posteriors on the \mathcal{H} parameters. This is likely caused by the larger discontinuities

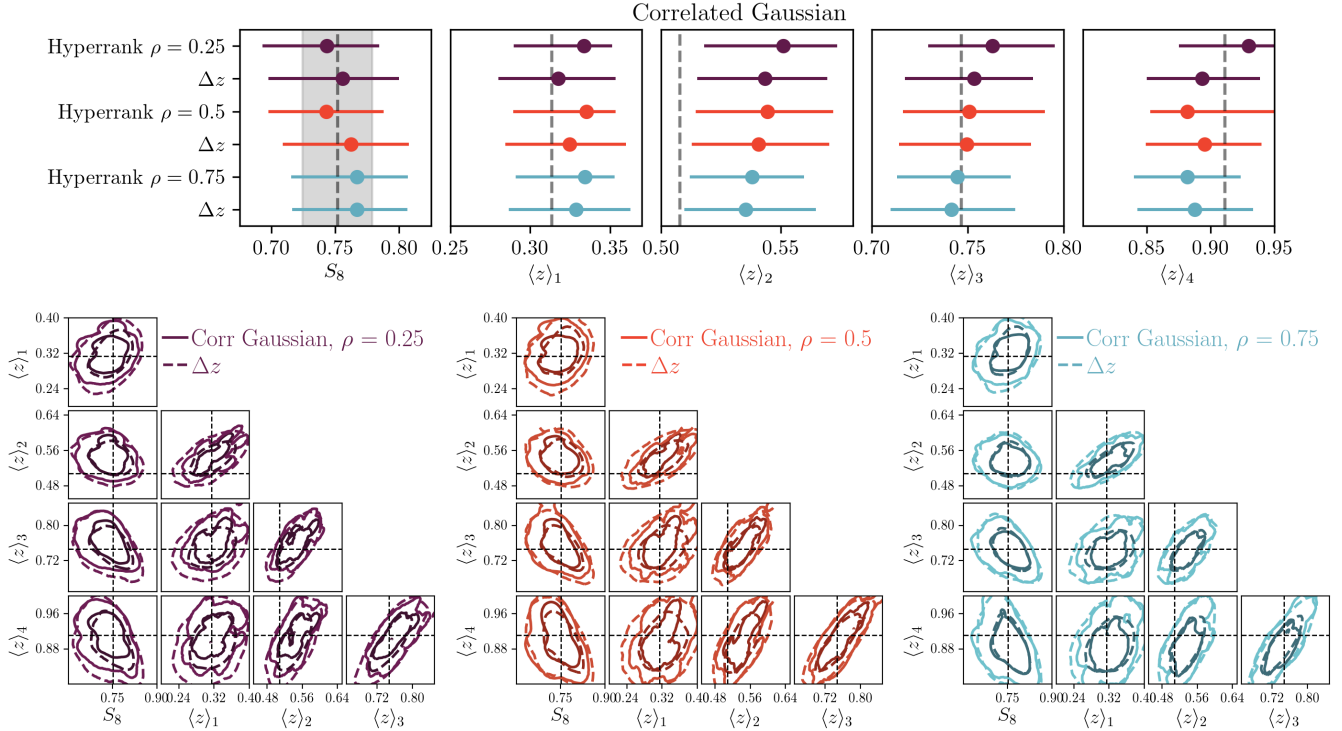


Figure 7. *Upper:* Marginalised one- σ confidence regions for the S_8 , and $\langle z \rangle_i$ parameters for the correlated Gaussian distribution model of $n(z)$ uncertainty described in section 4.2.3. Different ρ values refer to different amounts of correlation between tomographic bins in the uncertainty distributions and for each value of ρ we show both the posterior from chains using `HYPERRANK` and the Δz marginalisation schemes. *Lower:* Two dimensional posteriors on the same parameters. Dashed grey lines correspond to mean values of the fiducial redshift distribution in each tomographic bin, and in the S_8 panel the 1σ region inferred from a chain run without marginalisation over redshift nuisance parameters is also shown.

in the posterior surface as a function of \mathcal{H} , which is a consequence of the lower resolution of the multidimensional grid in higher N_d (as discussed in section 3.2).

Based on these results we consider a 3D approach an appropriate default configuration, with a preference for $\langle z \rangle$ since its computation does not involve the use of a fiducial cosmology, unlike $\langle 1/\chi \rangle$ (which requires calculation of $\chi(z)$ for equation (4), done here and typically elsewhere at a fixed cosmology). When considering which three of the four tomographic bins to choose to use as `HYPERRANK` dimensions, we recommend that tomographic bins should be ordered in terms of the variance in the descriptive value (e.g. the spread of different mean redshifts across different realisations), with the tomographic bin with lowest variance in the descriptive value not used as a `HYPERRANK` dimension.

In figure 10 we show posterior contours recovered for each of the different ranking schemes, including the ‘Random’ scheme in which no ranking is performed. The consistency of these contours confirms that the `HYPERRANK` procedure does not affect the cosmology recovered whilst improving the efficiency of a chain with respect to un-ranked, random sampling of $n(z)$ realisations.

4.3.1 Convergence

In `HYPERRANK`, discrete samples from the posterior over the subset of redshift nuisance parameters are generated outside of the main chain used to sample over the cosmological and other nuisance parameters. This means a limited and discrete set of values of the nuisance parameters are available to the main sampling, as opposed to the continuous range of parameters within a specified prior which

would be available otherwise. Whilst it could be possible to form an interpolation between closely-ranked samples to boost their number, such interpolated samples would no longer carry the intended property of being true samples from the posterior for $n(z)$, with the interpolation scheme effectively becoming part of the model and correlation structure for histogram bin heights, with a set of hidden hyperparameters. It is also unclear how such an interpolation could include the correlations across tomographic bins which we have found to be an important describing factor of the samples.

Without interpolation, there will be a transition from the regime in which there are too few realisations of $n(z)$ available to effectively explore the redshift distribution uncertainty, and the limit where infinitely many realisations would be available, corresponding to the continuous case. Here we investigate the convergence of `HYPERRANK` marginalisation with respect to the number of $n(z)$ samples generated, for the case of our DES-Y3 simulated data set.

We first generate several sets of distributions where each realisation is a shifted version of a fiducial $n_{\text{fid}}(z)$, and the shifts are drawn from a Gaussian prior, following a similar approach to the Δz method described in Section 2. We generate 8 sets of redshift distributions, each containing $3^3, 4^3, 5^3, 6^3, 7^3, 8^3, 9^3, 10^3$ realisations which are then ranked using the 3 dimensional default configuration described at the end of section 3.2.

Since we expect the approximate minimum number of realisations required for this convergence to depend on the level of uncertainty in the $n(z)$, we generate two additional sets of proposal distributions by multiplying the $\sigma(\Delta z)$ obtained above, by a factor $f_m = 5, 10$. We then repeat the generation of proposal realisations with five different random seeds for each of the three f_m values, and

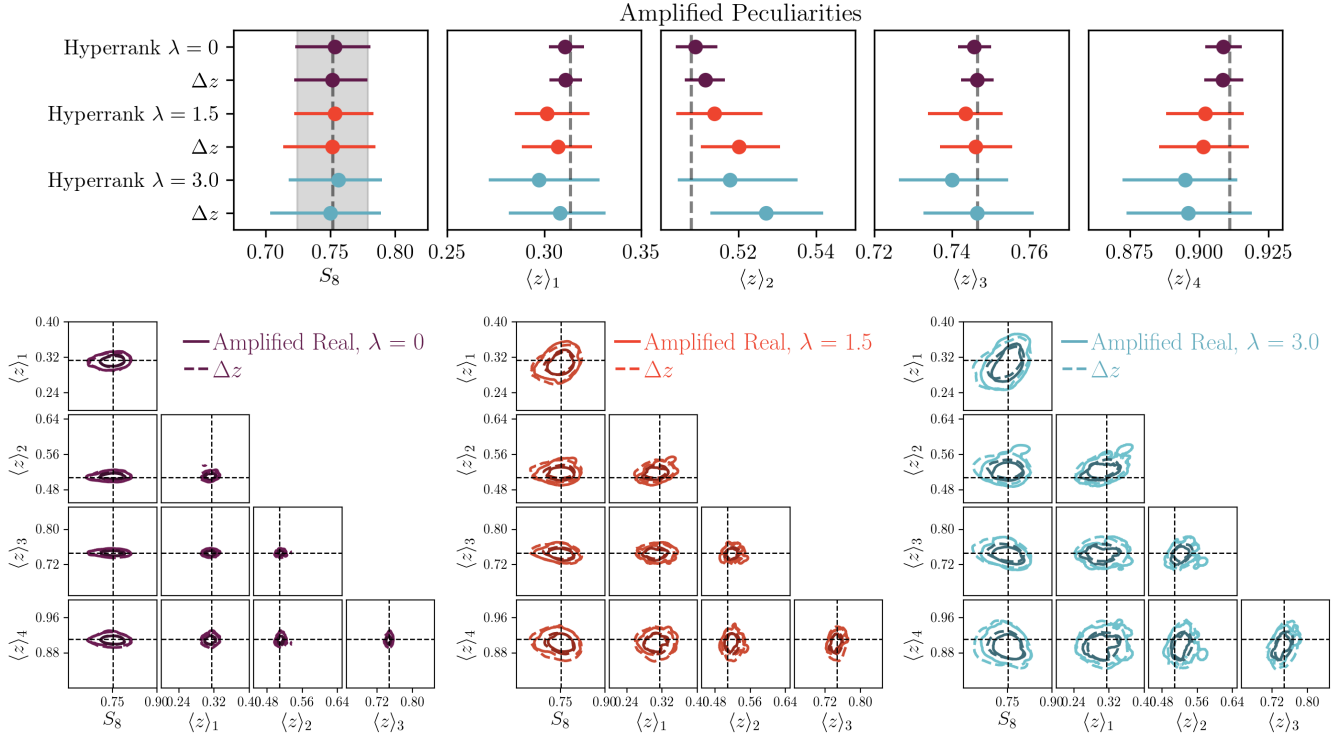


Figure 8. *Upper:* Marginalised one- σ confidence regions for the S_8 , and $\langle z \rangle_i$ parameters for the amplified deviations model of $n(z)$ uncertainty described in section 4.2.4. Different λ values refer to different amplifications uncertainty in the $n(z)$ distributions and for each value of λ we show both the posterior from chains using HYPERRANK and the Δz marginalisation schemes. *Lower:* Two dimensional posteriors on the same parameters. Dashed grey lines correspond to mean values of the fiducial redshift distribution in each tomographic bin, and in the S_8 panel the 1σ region inferred from a chain run without marginalisation over redshift nuisance parameters is also shown.

for each of the 8 sets of realisations containing different number of proposals. By comparing the standard deviation on the central, lower and upper confidence values for S_8 as a function of the number of realisations, we can find an approximate minimum number of realisations required for the standard deviation of error bars from HYPERRANK to converge to that obtained using the Δz approach (which is formally correct for this set of realisations). In figure 11 we observe that for all three levels of uncertainty, described by the amplification factor f_m , 1000 realisations yield standard deviation of the error bars obtained using HYPERRANK comparable to the ones using the Δz approach.

5 APPLICATION TO DES YEAR 3

Based on the above tests, we derive an appropriate configuration for using HYPERRANK on DES-Y3 (or similar) data:

- $\langle z \rangle$ ranking
- Three HYPERRANK parameter dimensions
- Ranking according to tomographic bins 1, 2 and 4
- At least 10^3 $n(z)$ samples available to HYPERRANK

We then run a full shear-only cosmology chain on the BUZZARD simulation of the DES-Y3 data set, with model parameterisations and priors as discussed in the main cosmology papers (Amon et al. 2021; Secco, Samuroff et al. 2021), and with 1000 realisations of possible $n(z)$ generated using the full procedure of Myles, Alarcon et al. (2020). This as closely as possible mimics the experimental data and setup of the DES-Y3 analysis. We also run a chain with this

setup, but with the HYPERRANK marginalisation of redshift uncertainties replaced by the Δz approach. This results of these two chains are shown in figure 12. The left panel shows the posteriors on mean redshift within the four tomographic bins, produced directly by the Δz analysis and by taking the posterior weighted means within the HYPERRANK analysis. Good consistency can be seen in the space of mean redshifts explored. The right panel of figure 12 shows posteriors on the cosmological parameters, and the HYPERRANK parameters for each tomographic bin. For the cosmological parameters we also show the recovered posterior from the Δz analysis, finding highly consistent results between the two approaches. This suggests that for the uncertainties which are modelled as part of the DES-Y3 analysis, the Δz approach is adequate to fully explore their effect on cosmic shear cosmological parameters. The Δz approach is hence adopted as fiducial in Amon et al. (2021) and Secco, Samuroff et al. (2021) and subsequent DES-Y3 analyses, with the validation test between HYPERRANK and Δz shown here on the BUZZARD simulation repeated for the real data vector in section E.1. of Amon et al. (2021).

We also show the recovered posteriors on HYPERRANK ranking parameters, showing that different subspaces of the ranked $n(z)$ realisations are indeed favoured in a systematic way, indicating the cosmological data are in turn helping constrain the space of plausible redshift distributions.




Figure 9. The sampling efficiency η for the different mapping schemes described in section 3, for three prior amplifications values f_m ($f_m = 1$ is the original BUZZARD redshift distributions). The spread of points at each location shows the differences due to different initial random seeds. The rankings are ordered from left to right as a function of perceived complexity, with a random ranking being the most naive approach and a 3D KL corresponding to the most complex to implement. Horizontal dashed lines show the efficiencies obtained by marginalising the same equivalent uncertainty $\sigma(\Delta z)$ using the Δz method, obtained after averaging 5 runs with each equivalent prior.

6 CONCLUSIONS

We have presented HYPERRANK, a new approach to marginalise over redshift distribution uncertainties in weak lensing and galaxy clustering experiments by ranking and mapping a set of proposal redshift distributions to a set of continuous hyper-parameter which are then sampled in the Monte Carlo chain.

To test the accuracy of the method we generated a series of $n(z)$ ensembles to describe different types of uncertainty, and compared the obtained S_8 error estimates and sampled uncertainty to those obtained by only marginalising over a shift Δz along the redshift direction for each tomographic bin.

We showed that this approach provides equivalent results to the ones obtained marginalising over Δz , when the realisations of the ensemble are obtained by shifting the tomographic bins of a fiducial distribution by a set of values drawn from the same prior used to describe the Δz uncertainty.

We generated additional ensembles to represent types of uncertainty which cannot be fully characterised by a set of uncorrelated Gaussian shifts Δz , and if unaccounted for, can lead to an incorrect estimation of the marginalised cosmological parameters posteriors. These included samples with Δz shifts drawn from non-Gaussian distributions, drawn from highly correlated multivariate Gaussian distributions and from a set of realistic distributions with amplified peculiarities, based on the estimates obtained from the SOMPZ scheme on the BUZZARD simulations. In all cases, HYPERRANK correctly explores the uncertainty described by the input distribution ensemble, providing posteriors on the cosmological S_8 and redshift tomographic bin means $\langle z \rangle$ which are highly consistent with those

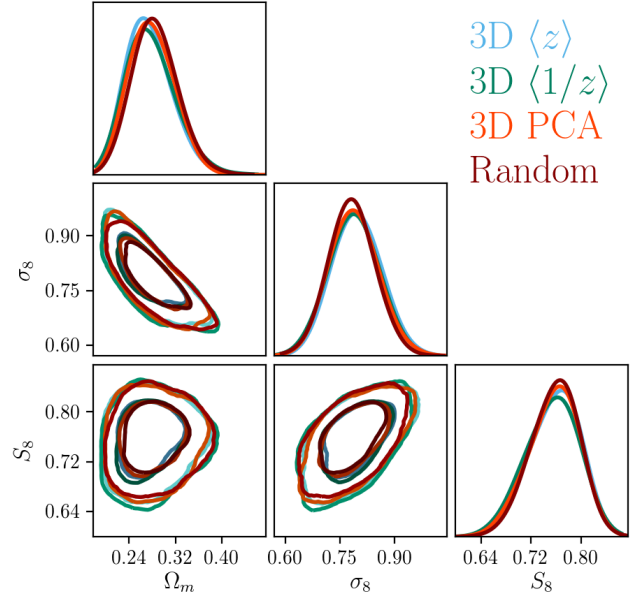


Figure 10. Recovered cosmological parameters for four different ranking approaches described in the text: Three dimensional rankings by mean redshift and inverse comoving distance of tomographic bins, three dimensional ranking by principal components of the data vector, and random sampling of realisations. The contours shown correspond to the case where realisations sampled with hyperrank describe an uncorrelated Gaussian distribution with an amplification factor of the uncertainty $f_m = 5$. For $f_m = 1, 10$ the contours are also very similar between the different ranking schemes.

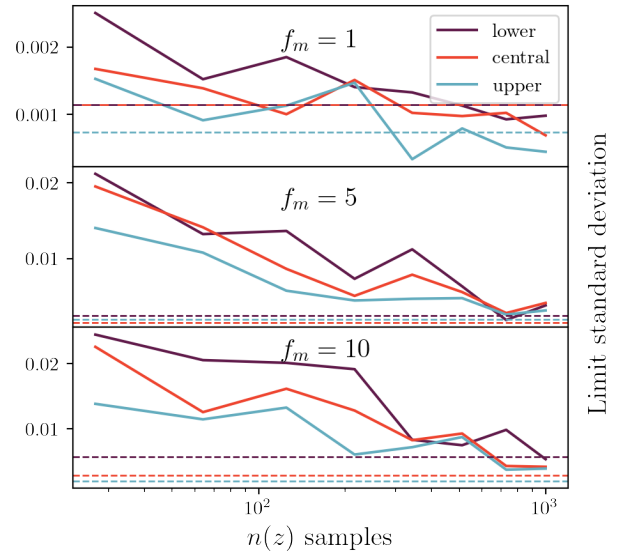


Figure 11. Standard deviation of the lower (purple), central (red), and upper (cyan) values for the S_8 parameter obtained using HYPERRANK for 5 realisations of the ensemble of $n(z)$ samples, as a function of the total number of distributions to form the ensemble. From top to bottom, the equivalent $\sigma(\Delta z)$ width is amplified by a progressively larger number, f_m , with respect to the original distributions of BUZZARD samples. Horizontal dashed lines indicate the typical standard deviation for runs using the traditional Δz marginalisation approach.

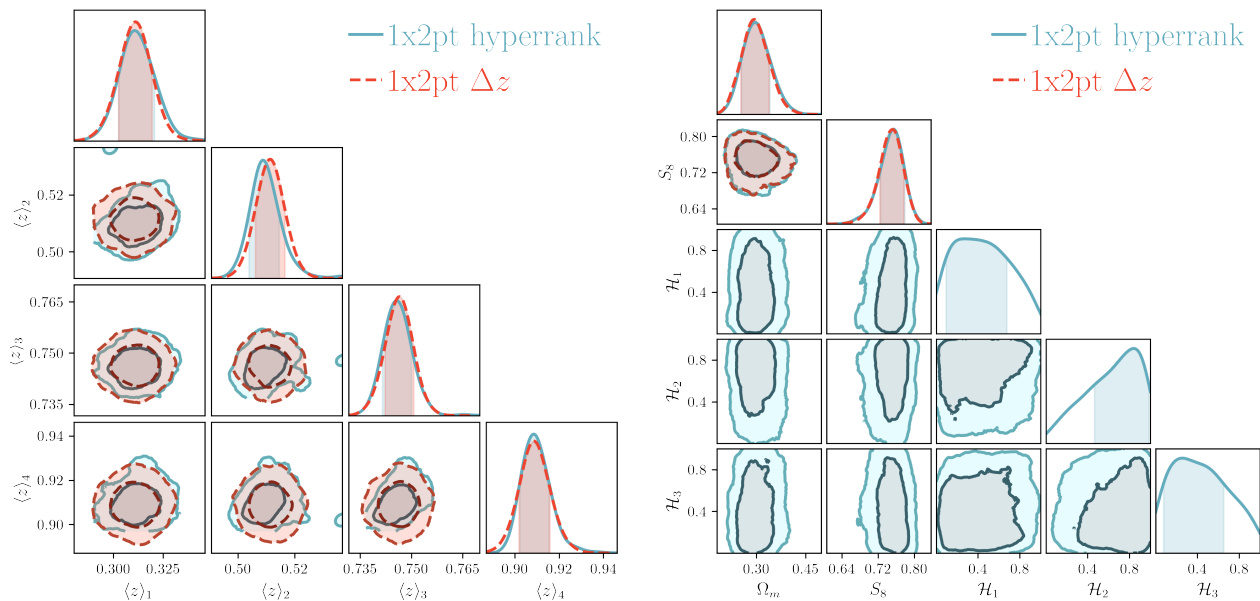


Figure 12. Correlation between redshift distribution uncertainty nuisance parameters in the BUZZARD simulated DES-Y3 analysis, comparing the standard Δz approach (red) with the HYPERRANK approach presented in this work (blue). *Left* shows the recovered posteriors on mean redshifts of redshift distributions within the tomographic bins considered. *Right* shows the recovered cosmological parameters for both approaches, and the HYPERRANK ranking parameters. Both show good agreement between the two approaches for the modelled uncertainty expected in DES-Y3.

from the estimates obtained using marginalisation with Δz nuisance parameters which are distributed according to the input model (and hence are the correct model for the uncertainty).

A set of tests were conducted to obtain an approximately optimal configuration for the choice of descriptive values used to rank the distributions and the subsequent effect on sampling efficiency, resulting in the use of mean redshift of a subset of tomographic bins, $\langle z \rangle_n$ being the choice of ranking parameter which gives the best efficiency (lowest number of likelihood evaluations per posterior sample required for convergence of the chain). As estimation of the minimum number of samples required for posterior estimates to become less noisy than the typical sampling noise in standard Δz marginalisation is also provided for the expected photometric redshift uncertainties of source distributions of the DES-Y3 analysis.

Tests were conducted simulating a cosmic shear analysis where only a subset of cosmological and systematic parameters are inferred, compared to a full cosmic shear plus galaxy clustering case. Despite this, HYPERRANK is not limited to cosmic shear analysis and can be used without significant modifications on cosmic shear plus galaxy clustering (3x2pt) analysis. We do not expect our conclusions to vary significantly for 3x2pt analysis. Similarly, while tests here focus on the propagation of uncertainty from *source* galaxy redshift distributions, HYPERRANK can be used to simultaneously and independently propagate uncertainties from the *lens* sample of galaxies for galaxy clustering plus tangential shear (2x2pt) and 3x2pt.

For the particular levels of uncertainty expected for the DES-Y3 analysis we showed that the difference in obtained confidence contours between the standard approach using Δz shifts and HYPERRANK are small and hence concluded that Δz was sufficient for the requirements of DES-Y3. For the level of uncertainties present in DES-Y3 we have demonstrated that it is satisfactory to use the Δz approach which, whilst not as accurate as the HYPERRANK approach, typically allows for faster convergence of the Monte Carlo inference

chains, as can be seen as the dashed horizontal lines in figure 9 which show the efficiencies for Δz .

HYPERRANK provides a well-motivated approach for marginalising over the redshift distribution uncertainty affecting cosmological galaxy clustering and weak lensing surveys. It is nominally capable of marginalising over any potential form of such an uncertainty, subject to the ability to generate realisations samples of possible $n(z)$ using a model for the uncertainty. It thus also provides a much more complete and flexible approach to the commonly used and ad-hoc Δz approach, whilst still being able to contain that particular model and replicate findings made using it.

ACKNOWLEDGEMENTS

This research manuscript made use of Astropy (Astropy Collaboration 2013; Astropy Collaboration & Astropy Contributors 2018), ChainConsumer¹ (Hinton 2016) and Matplotlib (Hunter 2007), and has been prepared using NASA’s Astrophysics Data System Bibliographic Services.

IH, RR, SB acknowledge support from the European Research Council in the form of a Consolidator Grant with number 681431.

IH also acknowledges support from the Beecroft Trust.

JPC acknowledges support granted by Agencia Nacional de Investigación y Desarrollo (ANID) DOCTORADO BECAS CHILE/2016 - 72170279.

Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education

¹ samreay.github.io/ChainConsumer

Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministério da Ciência, Tecnologia e Inovação, the Deutsche Forschungsgemeinschaft and the Collaborating Institutions in the Dark Energy Survey.

The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zürich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l'Espai (IEEC/CSIC), the Institut de Física d'Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, NSF's NOIRLab, the University of Nottingham, The Ohio State University, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, Texas A&M University, and the OzDES Membership Consortium.

Based in part on observations at Cerro Tololo Inter-American Observatory at NSF's NOIRLab (NOIRLab Prop. ID 2012B-0001; PI: J. Frieman), which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

The DES data management system is supported by the National Science Foundation under Grant Numbers AST-1138766 and AST-1536171. The DES participants from Spanish institutions are partially supported by MICINN under grants ESP2017-89838, PGC2018-094773, PGC2018-102021, SEV-2016-0588, SEV-2016-0597, and MDM-2015-0509, some of which include ERDF funds from the European Union. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya. Research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) including ERC grant agreements 240672, 291329, and 306478. We acknowledge support from the Brazilian Instituto Nacional de Ciência e Tecnologia (INCT) do e-Universo (CNPq grant 465376/2014-2).

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

DATA AVAILABILITY

The DES Y3 data products used in this work, full ensemble of DES Y3 source galaxy redshift distributions, chains and data products of the tests conducted here will be made publicly available following publication, at the URL <https://des.ncsa.illinois.edu/releases>.

The HYPERRANK code is available in the COSMOSIS standard library repository <https://bitbucket.org/joezuntz/cosmosis-standard-library/src/des-y3/>.

REFERENCES

- Abazajian, K. N., Adelman-McCarthy, J. K., Agüeros, M. A., et al., 2009, *ApJS*, 182, 2, 543, arXiv:0812.0649
- Amon, A., et al., 2021, Submitted to PRD, arXiv:2105.13543
- Astropy Collaboration, 2018, *AJ*, 156, 3, 123, arXiv:1801.02634
- Astropy Collaboration, 2013, *A&A*, 558, A33, arXiv:1307.6212
- Blanton, M. R., Schlegel, D. J., Strauss, M. A., et al., 2005, *AJ*, 129, 6, 2562, arXiv:astro-ph/0410166
- Bonnett, C., Troxel, M. A., Hartley, W., et al., 2016, *Phys. Rev. D*, 94, 4, 042005, arXiv:1507.05909
- Burkard, R. E., Derigs, U., 1980, *The Linear Sum Assignment Problem*, 1–15, Springer Berlin Heidelberg, Berlin, Heidelberg
- Cawthon, R., et al., 2020, Submitted to MNRAS, arXiv:2012.12826
- Dark Energy Survey Collaboration, 2018, *Phys. Rev. D*, 98, 4, 043526, arXiv:1708.01530
- DeRose, J., Wechsler, R. H., Becker, M. R., et al., 2019, arXiv e-prints, arXiv:1901.02401
- DeRose, J., Wechsler, R. H., Becker, M. R., et al., 2021, arXiv e-prints, arXiv:2105.13547
- DES Collaboration, 2021, arXiv e-prints, arXiv:2105.13549
- Everett, S., et al., 2020, Submitted to *ApJS*, arXiv:2012.12825
- Feroz, F., Hobson, M. P., Bridges, M., 2009, *MNRAS*, 398, 4, 1601, arXiv:0809.3437
- Foreman-Mackey, D., Hogg, D. W., Lang, D., Goodman, J., 2013, *PASP*, 125, 925, 306, arXiv:1202.3665
- Gatti, M., Giannini, G., et al., 2020, Submitted to MNRAS, arXiv:2012.08569
- Hadzhiyska, B., Alonso, D., Nicola, A., Slosar, A., 2020, *J. Cosmology Astropart. Phys.*, 2020, 10, 056, arXiv:2007.14989
- Hamana, T., Shirasaki, M., Miyazaki, S., et al., 2020, *PASJ*, 72, 1, 16, arXiv:1906.06041
- Handley, W. J., Hobson, M. P., Lasenby, A. N., 2015, *MNRAS*, 450, L61, arXiv:1502.01856
- Hartley, W. G., Chang, C., Samani, S., et al., 2020, *MNRAS*, 496, 4, 4769, arXiv:2003.10454
- Hartley, W. G., Choi, A., et al., 2020, Submitted to MNRAS, arXiv:2012.12824
- Heymans, C., Tröster, T., Asgari, M., et al., 2021, *A&A*, 646, A140, arXiv:2007.15632
- Hikage, C., Oguri, M., Hamana, T., et al., 2019, *PASJ*, 71, 2, 43, arXiv:1809.09148
- Hildebrandt, H., Viola, M., Heymans, C., et al., 2017, *MNRAS*, 465, 2, 1454, arXiv:1606.05338
- Hinton, S. R., 2016, *The Journal of Open Source Software*, 1, 4, 00045
- Hoyle, B., Gruen, D., Bernstein, G. M., et al., 2018, *MNRAS*, 478, 1, 592, arXiv:1708.01532
- Hu, W., 1999, *ApJ*, 522, 1, L21, arXiv:astro-ph/9904153
- Hunter, J. D., 2007, *Computing in Science & Engineering*, 9, 3, 90
- Joachimi, B., Lin, C. A., Asgari, M., et al., 2021, *A&A*, 646, A129, arXiv:2007.01844
- Joudaki, S., Hildebrandt, H., Traykova, D., et al., 2020, *A&A*, 638, L1, arXiv:1906.09262
- Kitching, T. D., Taylor, A. N., 2011, *MNRAS*, 410, 3, 1677, arXiv:1005.2063
- Leistedt, B., Mortlock, D. J., Peiris, H. V., 2016, *MNRAS*, 460, 4, 4258, arXiv:1602.05960
- Masters, D., Capak, P., Stern, D., et al., 2015, *ApJ*, 813, 1, 53, arXiv:1509.03318

Myles, J., Alarcon, A., et al., 2020, Submitted to MNRAS, arXiv:2012.08566
 Rau, M. M., Wilson, S., Mandelbaum, R., 2020, MNRAS, 491, 4, 4768, arXiv:1904.09988
 Sánchez, C., Bernstein, G. M., 2019, MNRAS, 483, 2, 2801, arXiv:1807.11873
 Sánchez, C., Raveri, M., Alarcon, A., Bernstein, G. M., 2020, MNRAS, 498, 2, 2984, arXiv:2004.09542
 Schmidt, S. J., Malz, A. I., Soo, J. Y. H., et al., 2020, MNRAS, 499, 2, 1587, arXiv:2001.03621
 Secco, L. F., Samuroff, S., Krause, E., et al., 2021, arXiv e-prints, arXiv:2105.13544
 Springel, V., 2005, MNRAS, 364, 4, 1105, arXiv:astro-ph/0505010
 Stözlner, B., Joachimi, B., Korn, A., Hildebrandt, H., Wright, A. H., 2021, A&A, 650, A148, arXiv:2012.07707
 Taylor, A. N., Kitching, T. D., 2010, MNRAS, 408, 2, 865, arXiv:1003.1136
 Tessore, N., Harrison, I., 2020, The Open Journal of Astrophysics, 3, 1, 6, arXiv:2003.11558
 Wright, A. H., Hildebrandt, H., van den Busch, J. L., et al., 2020, A&A, 640, L14, arXiv:2005.04207
 Zuntz, J., Paterno, M., Jennings, E., et al., 2015, Astronomy and Computing, 12, 45, arXiv:1409.3409

AFFILIATIONS

¹ Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK
² Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK
³ Institute for Astronomy, University of Edinburgh, Edinburgh EH9 3HJ, UK
⁴ Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA
⁵ Argonne National Laboratory, 9700 South Cass Avenue, Lemont, IL 60439, USA
⁶ Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain
⁷ Institut d'Estudis Espacials de Catalunya (IEEC), E-08034 Barcelona, Spain
⁸ Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA
⁹ Instituto de Física Teórica, Universidade Estadual Paulista, São Paulo, Brazil
¹⁰ Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA
¹¹ Kavli Institute for Particle Astrophysics & Cosmology, P. O. Box 2450, Stanford University, Stanford, CA 94305, USA
¹² SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
¹³ Laboratório Interinstitucional de e-Astronomia - LIneA, Rua Gal. José Cristino 77, Rio de Janeiro, RJ - 20921-400, Brazil
¹⁴ Department of Physics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15312, USA
¹⁵ Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, OH 43210, USA
¹⁶ Santa Cruz Institute for Particle Physics, Santa Cruz, CA 95064, USA
¹⁷ Department of Astronomy, University of California, Berkeley,

501 Campbell Hall, Berkeley, CA 94720, USA
¹⁸ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109, USA
¹⁹ Department of Astronomy/Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721-0065, USA
²⁰ Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK
²¹ Churchill College, University of Cambridge, CB3 0DS Cambridge, UK
²² Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 W. Green Street, Urbana, IL 61801, USA
²³ Center for Astrophysical Surveys, National Center for Supercomputing Applications, 1205 West Clark St., Urbana, IL 61801, USA
²⁴ Département de Physique Théorique and Center for Astroparticle Physics, Université de Genève, 24 quai Ernest Ansermet, CH-1211 Geneva, Switzerland
²⁵ Fermi National Accelerator Laboratory, P. O. Box 500, Batavia, IL 60510, USA
²⁶ Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, UK
²⁷ Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA
²⁸ ICTP South American Institute for Fundamental Research Instituto de Física Teórica, Universidade Estadual Paulista, São Paulo, Brazil
²⁹ Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain
³⁰ Brookhaven National Laboratory, Bldg 510, Upton, NY 11973, USA
³¹ Department of Physics, Duke University Durham, NC 27708, USA
³² Departamento de Física Matemática, Instituto de Física, Universidade de São Paulo, CP 66318, São Paulo, SP, 05314-970, Brazil
³³ CNRS, UMR 7095, Institut d'Astrophysique de Paris, F-75014, Paris, France
³⁴ Sorbonne Universités, UPMC Univ Paris 06, UMR 7095, Institut d'Astrophysique de Paris, F-75014, Paris, France
³⁵ Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK
³⁶ Instituto de Astrofísica de Canarias, E-38205 La Laguna, Tenerife, Spain
³⁷ Universidad de La Laguna, Dpto. Astrofísica, E-38206 La Laguna, Tenerife, Spain
³⁸ Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona) Spain
³⁹ Physics Department, 2320 Chamberlin Hall, University of Wisconsin-Madison, 1150 University Avenue Madison, WI 53706-1390
⁴⁰ INAF-Osservatorio Astronomico di Trieste, via G. B. Tiepolo 11, I-34143 Trieste, Italy
⁴¹ Institute for Fundamental Physics of the Universe, Via Beirut 2, 34014 Trieste, Italy
⁴² Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstr. 1, 81679 Munich, Germany
⁴³ Department of Physics, The Ohio State University, Columbus, OH 43210, USA
⁴⁴ Institute of Theoretical Astrophysics, University of Oslo. P.O. Box 1029 Blindern, NO-0315 Oslo, Norway
⁴⁵ Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, 28049 Madrid, Spain
⁴⁶ Department of Astronomy, University of Michigan, Ann Arbor,

MI 48109, USA

⁴⁷ School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia

⁴⁸ Max Planck Institute for Extraterrestrial Physics, Giessenbachstrasse, 85748 Garching, Germany

⁴⁹ Universitäts-Sternwarte, Fakultät für Physik, Ludwig-Maximilians Universität München, Scheinerstr. 1, 81679 München, Germany

⁵⁰ Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA

⁵¹ Australian Astronomical Optics, Macquarie University, North Ryde, NSW 2113, Australia

⁵² Lowell Observatory, 1400 Mars Hill Rd, Flagstaff, AZ 86001, USA

⁵³ Institució Catalana de Recerca i Estudis Avançats, E-08010 Barcelona, Spain

⁵⁴ Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada

⁵⁵ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

⁵⁶ Observatório Nacional, Rua Gal. José Cristino 77, Rio de Janeiro, RJ - 20921-400, Brazil

⁵⁷ Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA

⁵⁸ School of Physics and Astronomy, University of Southampton, Southampton, SO17 1BJ, UK

⁵⁹ Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831

⁶⁰ Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth, PO1 3FX, UK

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.