

KiDS-SBI: Simulation-based inference analysis of KiDS-1000 cosmic shear

Maximilian von Wietersheim-Kramsta^{1,2,3,*}, Kiyam Lin¹, Nicolas Tessore¹, Benjamin Joachimi¹,
Arthur Loureiro^{4,5}, Robert Reischke^{6,7}, and Angus H. Wright⁷

¹ Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

² Department of Physics, Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK

³ Department of Physics, Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK

⁴ The Oskar Klein Centre, Department of Physics, Stockholm University, AlbaNova University Centre, SE-106 91 Stockholm, Sweden

⁵ Astrophysics Group and Imperial Centre for Inference and Cosmology (ICIC), Blackett Laboratory, Imperial College London, London SW7 2AZ, UK

⁶ Argelander-Institut für Astronomie, Auf dem Hügel 71, D-53121 Bonn, Germany

⁷ Ruhr University Bochum, Faculty of Physics and Astronomy, Astronomical Institute (AIRUB), German Centre for Cosmological Lensing, D-44780 Bochum, Germany

Received 23 April 2024 / Accepted 10 December 2024

ABSTRACT

We present a simulation-based inference (SBI) cosmological analysis of cosmic shear two-point statistics from the fourth weak gravitational lensing data release of the ESO Kilo-Degree Survey (KiDS-1000). KiDS-SBI efficiently performs non-Limber projection of the matter power spectrum via Levin's method and constructs log-normal random matter fields on the curved sky for arbitrary cosmologies, including effective prescriptions for intrinsic alignments and baryonic feedback. The forward model samples realistic galaxy positions and shapes, based on the observational characteristics of KiDS-1000. It incorporates shear measurement and redshift calibration uncertainties, as well as angular anisotropies due to variable survey depth and point spread function (PSF) variations. To enable direct comparisons with standard inference, we limited our analysis to pseudo-angular power spectra as summary statistics. Here, the SBI is based on neural density estimation of the likelihood with active learning to infer the posterior distribution of spatially flat Λ CDM cosmological parameters from 18 000 realisations. We inferred a mean marginal for the growth of the structure parameter of $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5} = 0.731 \pm 0.033$ (68%). We present a measurement of the goodness-of-fit for SBI, determining that the forward model fits the data well, with a probability-to-exceed of 0.42. For a fixed cosmology, the learnt likelihood is approximately Gaussian, while its constraints are wider, compared to a Gaussian likelihood analysis due to the cosmology dependence in the covariance. Neglecting variable depth and anisotropies in the point spread function in the model can cause S_8 to be overestimated by $\sim 5\%$. Our results are in agreement with previous analyses of KiDS-1000 and reinforce a 2.9σ tension with early Universe constraints from cosmic microwave background measurements. This work highlights the importance of forward-modelling systematic effects in upcoming galaxy surveys, such as *Euclid*, *Rubin*, and *Roman*.

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

1. Introduction

The weak gravitational lensing effect on distant galaxies due to matter in the foreground, known as cosmic shear, is a powerful tool for studying the distribution of matter in the Universe and to probe its large-scale structure. By measuring the distortions in the shapes of galaxy images caused by the gravitational influence of intervening matter, we can infer a combination of the matter density and the amplitude of the matter power spectrum within the framework of a spatially flat cold dark matter cosmology with a cosmological constant, Λ CDM. Recent cosmic shear analyses of the data taken in stage-III galaxy surveys, such as the Kilo-Degree Survey¹ (KiDS; Kuijken et al. 2019; Asgari et al. 2021; Heymans et al. 2021; Loureiro et al. 2022; van den Busch et al.

2022; Li et al. 2023a), the Hyper Suprime-Cam survey² (HSC; Sugiyama et al. 2022; Aihara et al. 2022; Li et al. 2023b), and the Dark Energy Survey³ (DES; Gatti et al. 2021; Amon et al. 2022; Secco et al. 2022), have constrained these cosmological parameters with unprecedented precision. Upcoming stage-IV galaxy surveys, such as *Euclid*⁴ (Laureijs et al. 2011), *Rubin*⁵ (LSST Science Collaboration 2009), and *Roman*⁶ (Spergel et al. 2015), will further improve upon these constraints, as these surveys will be wider, deeper, and more precise. However, many analytical challenges still remain.

Current stage-III galaxy survey measurements of the root-mean square of the matter overdensity field at 8 Mpc h^{-1} ,

² <https://hsc.mtk.nao.ac.jp/ssp/>

³ <https://www.darkenergysurvey.org/>

⁴ <https://www.euclid-ec.org/>

⁵ <https://www.lsst.org/>

⁶ <https://roman.gsfc.nasa.gov/>

* Corresponding author; maximilian.von-wietersheim-kramsta@durham.ac.uk

¹ <https://kids.strw.leidenuniv.nl/>

σ_8 , or its analogue $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5}$, are in good agreement with each other (Hikage et al. 2019; Asgari et al. 2021; van den Busch et al. 2022; Amon et al. 2022; Secco et al. 2022; Li et al. 2023b; Dalal et al. 2023), despite the use of largely independent methodologies and survey volumes. However, the constraints from observations of the late Universe (Asgari et al. 2021; Heymans et al. 2021; van den Busch et al. 2022; Amon et al. 2022; Abbott et al. 2022; Li et al. 2023a; Dark Energy Survey and Kilo-Degree Survey Collaboration 2023) are in disagreement by up to $\sim 3.4\sigma$ with the σ_8 value consistent with the early Universe observations based on the cosmic microwave background (CMB; Planck Collaboration VI 2020). Recently, cluster abundances measured by SRG/eROSITA (Ghirardini et al. 2024) have provided another measurement of the fluctuation amplitude even higher than those of the CMB, although similar recent cluster studies have obtained S_8 constraints more consistent with weak lensing and CMB analyses (Mantz et al. 2015; Garrel et al. 2022; Chiu et al. 2023; Bocquet et al. 2024). This discrepancy might be pointing at new physical phenomena, but it could also be caused by unconsidered systematic effects when modelling the low-redshift large-scale structure or the CMB signal and noise (see for example Handley & Lemos 2021; Dark Energy Survey and Kilo-Degree Survey Collaboration 2023). Such systematic effects may be physical, such as baryonic feedback or the intrinsic alignments of galaxies (Mandelbaum 2018; Amon & Efstathiou 2022; Li et al. 2023b; Miyatake et al. 2023; Dalal et al. 2023), or observational.

We aim to shed some light on this issue by conducting a full cosmic shear analysis using neural likelihood estimation simulation-based inference (SBI) characterised by the same complexity as current stage-III analyses. SBI, also known as ‘likelihood-free inference’ or implicit ‘likelihood inference’, is a Bayesian inference method that does not require an explicit formulation for the likelihood function of the data given the parameters of interest. Instead, the likelihood is implicitly calculated by evaluating the joint probability of data and parameters of interest from forward simulations, which map the parameters to the corresponding mock data vectors. This comes with multiple advantages with respect to other standard approaches that require an explicit form for the likelihood. Firstly, the likelihood is allowed to take an arbitrary form, so we can avoid the typical assumption of a Gaussian likelihood or avoid having to define a complex analytical expression for the likelihood. Secondly, for some models and measurements, it may not even be possible to define an analytical likelihood or it would be too resource-intensive; for instance, the case of covariance among n^{th} -order statistics depending on up to $2n^{\text{th}}$ -order correlation functions, which become geometrically more expensive to compute. In such cases, as long as the observables can be simulated, an effective or implicit likelihood may be found using SBI. Lastly, in standard analyses, it often is required to validate the signal and noise modelling with forward simulations, so that we can make use of existing tools to perform SBI.

When running a Markov chain Monte Carlo (MCMC) to sample the posterior distributions from a Gaussian likelihood, we either have to compute a numerical covariance matrix from forward simulations or create an analytical model for the covariance, which ought to be validated with a numerical one. In both cases, for a data vector, \mathbf{d} , we would require a number of simulations $>|\mathbf{d}|$ to get an accurate sample covariance matrix. This is already the approximate amount of forward simulations needed for an SBI analysis (Alsing et al. 2018; Lin et al. 2023). Therefore, SBI allows us to perform a full Bayesian uncertainty prop-

agation from the data to obtain parameters for any model that can be simulated, without substantial additional computational costs when compared to standard approaches at a comparable accuracy.

With respect to cosmic shear, we know that the commonly used two-point statistics to quantify the shear-shear correlations of galaxies on the sky have an approximately Gaussian likelihood, but this assumption has its limitations (Schneider & Hartlap 2009; Sellentin & Heavens 2018; Sellentin et al. 2018; Taylor et al. 2019; Upham et al. 2021; Hall & Taylor 2022). Additionally, when observational and physical systematics are included, the likelihood may become even less Gaussian (Jeffrey et al. 2021). In such cases, like the one presented in this paper, SBI allows us to capture potential non-Gaussianities and propagate them to the posterior distributions.

The SBI approach can also enable us to compute effective likelihoods for summary statistics for which analytical methods may be intractable, such as higher order cosmic shear statistics (Fluri et al. 2022; Euclid Collaboration: Ajani et al. 2023; Lu et al. 2023). In this work, we limit the scope to only angular power spectra of cosmic shear and leave higher order statistics or field-level statistics within SBI as an avenue for future work. This allows us to make a direct comparison to previous standard cosmic shear analyses (such as Asgari et al. 2021; Loureiro et al. 2022), while also harnessing the efficiency of log-normal forward models of matter density fields which are mostly accurate for two-point statistics (Hall & Taylor 2022; Piras et al. 2023).

A potential cause for deviations from Gaussianity in the likelihood of two-point statistics/angular power spectra is the anisotropy in the observational depth of a galaxy survey across the sky (Guzik & Bernstein 2005; Shirasaki et al. 2019). As a survey such as KiDS images the sky, each pointing and each pixel are subject to variable environmental and background conditions that may affect the selection of galaxies (for example atmospheric conditions, thermal expansion of the telescope, moonlight, zodiacal light, Galactic extinction, etc.).

Due to the given survey strategy and/or dithering to compensate for gaps in the telescope’s charge-coupled device (CCD), it is possible that there is spatial variation in the overlap between pointings. When the images of each pointing are subsequently combined, some galaxies are observed more often than others that can cause variations in the signal-to-noise ratio (S/N), which are spatially correlated.

In effect, all these variations lead to fluctuations in the effective galaxy density, which are observed across the footprint, while also selecting galaxies according to their shape. This has been found to be a percent-level effect on the cosmic shear signal measured by KiDS (Heydenreich et al. 2020), while also affecting the standard deviation of the KiDS-1000 cosmic shear signal (by $\lesssim 20\%$; Joachimi et al. 2021). For KiDS-1000, this is within the statistical uncertainty, however, it will become important for upcoming galaxy shear surveys such as *Euclid*, *Rubin*, and *Roman*. Additionally, the effect of spatial variability on the non-Gaussianity of the likelihood has not been explored in previous works. To address these needs, in this work, our forward model includes the effects of the observational depth variations in KiDS, so that we may propagate the effects of spatial variations in the depth to the uncertainty of inferred cosmological parameters.

To conduct the SBI, we chose to make use of the density estimation likelihood-free inference (DELFI, Alsing et al. 2019) method, as it offers good performance and scales better when increasing the dimensionality of the parameter space of interest

than other SBI methods (Leclercq 2018; Alsing et al. 2018), such as approximate Bayesian computation (ABC; Rubin 1984; Pritchard et al. 1999; Lin & Kilbinger 2015; Lin et al. 2016; Beaumont 2019). DELFI discards less information from the forward simulations than ABC, whereas the latter only keeps the forward simulations which have a certain level of agreement with the desired data vector. In contrast, DELFI learns a probability density distribution of the data vectors as a function of the model parameters based on all the forward simulations. To achieve this, DELFI employs ensembles of neural density estimators to learn the sampling distribution of the data and the parameters from the forward simulations.

Of course, as with any inference method, we need to be mindful of any intrinsic biases that might come with SBI, as well as how to mitigate and avoid them. As the likelihood is not known analytically, it is important that the chosen SBI method, sampling scheme, and number of simulations are sufficient to ensure that the learnt likelihood has converged to the model's true likelihood function. To avoid biasing the inference, this requires careful robustness and convergence tests of the implicit likelihood, which may not always be necessary for a standard Gaussian likelihood analysis. Hence, we dedicate a considerable effort in this work to ensure the inference pipeline is validated, the neural density estimators of DELFI have converged, the inferred posterior is accurate, and the inferred parameters give a good fit to the data.

Applications of SBI are increasing in popularity, in major part because it requires about an order of magnitude fewer evaluations of the likelihood than a standard Gaussian likelihood MCMC analysis to constrain the posterior (Papamakarios & Murray 2016; Alsing et al. 2018; Gerardi et al. 2021; Jeffrey et al. 2021; Lemos et al. 2021; Legin et al. 2021; Mishra-Sharma & Cranmer 2022; Spurio Mancini et al. 2023; Lin et al. 2023; Chen et al. 2023; Hu et al. 2022; Lemos et al. 2023b,c; Gatti et al. 2024; Moser et al. 2024; Franco-Abellán et al. 2024). We highlight the methodology of Gatti et al. (2024), which used a similar analysis to the one outlined in this paper, with the DES Year 3 weak lensing data. We refer to Jeffrey et al. (2025) for a full description of the results. While the DES analysis focuses on SBI of two-point and higher-order shear statistics, along with cosmologies beyond Λ CDM based on a fixed number of N-body simulations, while incorporating standard systematics, the analysis in this paper exclusively considers two-point statistics within the context of Λ CDM. We highlight that in our analysis, the data likelihood is learned by actively sampling as many efficient forward simulations as needed. This allows us to capture many physical and observational systematics, including anisotropic systematics that are typically not considered in weak lensing analyses. With this, we aim to give more insight into the robustness against systematics of standard weak lensing analyses.

In a previous paper (Lin et al. 2023, L23 hereafter), we introduced a new inference pipeline for estimating the cosmological parameters from the cosmic shear data of KiDS-1000 (Kuijken et al. 2019) using DELFI. In L23, the SBI pipeline for KiDS is shown to be robust, accurate and efficient, even when constraining a 12-dimensional (12D) posterior distribution. However, in that work, the simulated vectors were based on random samples from the covariance matrix used in the fiducial Gaussian KiDS-1000 analysis, so the likelihood was Gaussian by construction to allow for direct comparison with the standard formalism. In this paper, we present a novel suite of physically motivated and realistic forward simulations of the cosmic shear observables, as seen by KiDS. Thus, the neural density estimators can learn any non-Gaussianities in the likelihood induced

by our model, plus any complex relations between the data and the model parameters. We then used these simulations to create two inference pipelines: a fiducial one which makes the same modelling choices as the KiDS-1000 analysis (Joachimi et al. 2021) and a more realistic one, which explicitly models systematic effects that would be prohibitively difficult to incorporate into a standard likelihood analysis. From each pipeline, we obtained a separate posterior estimate for S_8 , which we then scrutinised with a novel measure of the goodness-of-fit for SBI posteriors.

This paper has the following structure. Section 2 presents the KiDS data used for this analysis. Section 3 gives a detailed description of the forward simulations developed for the SBI analysis. Section 4 details how the galaxy shapes and summary statistics are measured in both the simulations and the data. Section 5 describes the setup of the simulation-based inference pipeline using DELFI based on L23. Section 6 presents the novel measure of goodness-of-fit for SBI we use in this paper. In Sect. 7, we describe our consistency tests to ensure the SBI pipeline is unbiased and robust, while in Sect. 8, we discuss the validity of a Gaussian likelihood assumption. Section 9 presents the cosmological inference results from the KiDS-1000 weak lensing data. Lastly, we give our concluding remarks in Sect. 10. We provide additional details about the analysis in the appendix.

2. KiDS-1000 data

The Kilo-Degree Survey (KiDS) is a large public galaxy survey conducted by the European Southern Observatory using the OmegaCAM CCD mosaic camera (Kuijken 2011), attached to the 2.6 m VLT Survey Telescope (VST). The survey covers approximately 1350 deg² between two distinct fields known as KiDS North and South, which are fields located in the northern and southern galactic caps (with Dec \approx 0° and Dec \approx -30°, respectively). This area is the same as the one covered by the VISTA Kilo-degree Infrared Galaxy survey (VIKING; Edge et al. 2013), which means that both surveys together observe every object with a total of nine photometric bands: *ugriZYJHK_s*. The analysis presented in this paper makes use of DR4 (Kuijken et al. 2019), also known as KiDS-1000, which covers approximately 1000 deg². The effective area covered by KiDS-1000 shape measurements (shown in Fig. 1) extends over 773.3 deg² (Joachimi et al. 2021). This area is calculated directly from the mosaic mask defined at the native OmegaCAM pixel scale of 0.213 arcsec.

The KiDS images were processed using Astro-WISE (McFarland et al. 2013) and the methodology outlined in Wright et al. (2019) to combine the VST and VIKING photometric measurements. To obtain the shape measurements, Theli (Erben et al. 2005) was used to process the *r*-band images, while the gravitational shear estimates are then obtained from lensfit (Miller et al. 2007, 2013; Fenech Conti et al. 2017). Depending on the shape noise variance and the ellipticity measurement noise variance, lensfit assigns a weight to the shear measurement of each galaxy, w_i , which scales the shear signal such that the signal-to-noise ratio is optimal. Although KiDS-SBI currently does not model lensfit weights (all galaxies are weighted equally by default), the weights are carried through the measurement pipeline when measuring the cosmic shear summary statistics from the real KiDS-1000 data. If the lensfit weight is mischaracterised, this may lead to shear bias. For this purpose, the shape measurements in the KiDS-1000 gold sample are calibrated to determine any shear biases due to selection




Fig. 1. Spatial map of the KiDS-1000 footprint. The top panel shows a Mollweide projection of the full KiDS-1000 footprint, while the two panels at the bottom show zoomed-in Cartesian projections of KiDS-North and KiDS-South fields, respectively.

biases, noise, weight bias, point spread function (PSF) residuals, or otherwise (Giblin et al. 2021). The resulting estimate of the linear multiplicative shear bias per tomographic bin is then used to unbiase the galaxy shear measurements in KiDS-1000. All details on the reduction of the images and the shape calibration are given in Kuijken et al. (2019), Wright et al. (2019), Giblin et al. (2021).

The photometric redshifts were estimated from the KiDS 9-band photometry using Bayesian template-fitting as incorporated within the BPZ code (Benítez 2000; Wright et al. 2019, 2020). The photometric redshifts are calibrated with spectroscopic redshift measurements made of galaxies in the KiDS sample (Hildebrandt et al. 2021) based on a subsample of 99% completeness. The mapping from photometric to spectroscopic redshift, $P(z|z_{\text{ph}})$, is learnt through a self-organising map (Wright et al. 2020; Hildebrandt et al. 2021). As is the case for previous KiDS-1000 analyses, we use the KiDS-1000 ‘gold-sample’, which only considers galaxies with reliable redshift and shape measurements. The maximum posterior photometric redshift is subsequently used to split the source galaxy catalogue into five tomographic bins at the following boundaries: (0.1, 0.3, 0.5, 0.7, 0.9, 1.2]. These give the redshift distributions shown in Fig. 2. As the cosmological constraints are sensitive to shifts in the redshift distributions, the KiDS-1000 analysis also allows the shift in the mean of each tomographic bin to vary freely in the cosmological inference as a nuisance parameter while incorporating for the correlations between them (Asgari et al. 2021; Heymans et al. 2021; van den Busch et al. 2022). This not only avoids biases in the cosmological parameter estimates, but also provides an additional test on the redshift calibration.

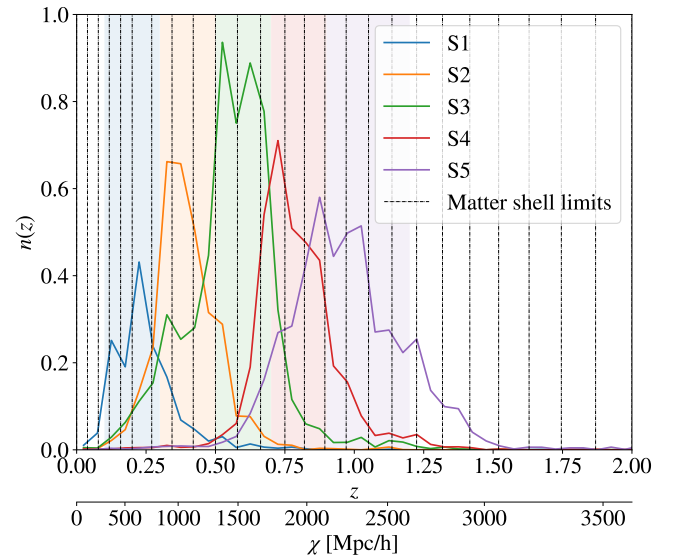


Fig. 2. Redshift distributions of the five KiDS-1000 tomographic bins. The shaded areas show the limits of each tomographic bin, while the solid lines show the $n(z)$ of the source galaxies in each tomographic bin as a function of both redshift, z , and comoving distance, χ (the latter is derived assuming a Planck 2018 cosmology; Planck Collaboration VI 2020). The black dashed lines show the limits of the spherical matter shells in our forward simulations.

We also note that KiDS-SBI treats the KiDS-1000 catalogue as a population of point sources unaffected by source blending, since the KiDS data already accounts for overlapping sources in

multiple ways. Firstly, the KiDS photometry integrates Gaussian Aperture and PSF photometry (GAaP; Kuijken et al. 2015, 2019; de Jong et al. 2017), where a Gaussian aperture is applied to the images of all sources that is optimised for S/N in the fluxes while also avoiding source blending. An aperture size of $2''$ is found to be optimal for these purposes and applied to all sources (Kuijken et al. 2015, 2019; de Jong et al. 2017). Secondly, at the source detection step with SExtractor (Bertin & Arnouts 1996), explicit flagging of overlapping sources occurs. Approximately 5% of sources (Kuijken et al. 2019; Giblin et al. 2021) are excluded from the sample due to this selection. Additionally, lensfit weights also account for any biases due to irregular objects which are not affected by the selection. We note that deblending of sources is not attempted at any point. Lastly, at the calibration stage, the KiDS-1000 shape measurements are validated against image simulations from Kannawadi et al. (2019). The multiplicative shear bias values obtained in this analysis (Giblin et al. 2021) incorporate any residual biases due to correlated noise in the shape measurements, including source blending.

3. Forward simulations

The forward simulations for cosmic shear analyses described in this paper are fully implemented and available in the KiDS-SBI⁷ module. This code is built within the Cosmological Survey Inference System (CosmoSIS⁸; Zuntz et al. 2015) and based upon the KiDS Cosmology Analysis Pipeline (KCAP⁹; Joachimi et al. 2021; Asgari et al. 2021; Heymans et al. 2021; Tröster et al. 2021). In addition, the KiDS-SBI pipeline is generalisable to other cosmic shear analyses.

The overall outline of the forward simulations is shown in Fig. 3. Firstly, the simulation is fed a set of cosmological and astrophysical parameters, Θ , from a sampler of our choice. All of these parameters are passed into CAMB (Lewis et al. 2000; Lewis & Challinor 2002; Howlett et al. 2012) which calculates the 3D matter power spectrum, $P_{\delta, \text{nl}}$, at a given cosmology (see also Sect. 3.1 for more details). We incorporate baryonic feedback using HMCode (Mead et al. 2016). Next, the simulation splits the light cone along the line of sight into concentric spherical shells centred at the observer, as required by our large-scale structure simulator of choice, GLASS (Tessore et al. 2023). Such an approach is computationally efficient as all observables are simulated with a coarse resolution along the line of sight. Within each shell, we sample log-normal random matter fields to efficiently model the two-point statistics with high accuracy, while also incorporating some higher-order fluctuations. As an added benefit, this simplifies the implementation of effects such as redshift evolution as well as survey characteristics, as they usually also follow a spherical geometry. On the other hand, the concentric shells also allow for the Universe to be split into shells, which are of a scale large enough such that a log-normal field becomes an accurate description of the matter over-densities within it (see Sect. 3.2 for details).

Once we know the geometry of our matter shells, we can project the 3D matter power spectrum using the redshift distributions of these shells in order to obtain their 2D angular power spectra, $C_{\delta\delta}$. The latter describes the correlations between the matter over-densities in each shell. Since we also require accu-

rate cross-correlations between different shells to compute the log-normal fields, the Limber approximation is not applicable here, as the cross-correlations are zero in that approximation with the shells being non-overlapping along the line of sight. To avoid this, we project the 3D matter power spectrum with a non-Limber integral. This is usually an expensive computation, so we show here a novel implementation of the Levin method of integration (Levin 1996), which allows for the efficient integration of the spherical Bessel functions found in non-Limber integrals (more details are given in Sect. 3.3).

We can then construct log-normal matter fields which are consistent with these correlations and their associated cosmology. To do so, we used the Generator for Large Scale Structure (GLASS¹⁰, Tessore et al. 2023), which efficiently generates correlated log-normal random matter fields, δ , for each shell. Then, it accurately computes the respective convergence field, κ , (see Sect. 3.4) while also incorporating an effective convergence field due to intrinsic galaxy alignments, κ_{IA} (see Sect. 3.5).

Upon the construction of the matter, δ , and convergence fields, κ , we used them to Poisson-sample galaxies that reside within the matter fields as well as their properties, namely, intrinsic ellipticities and shear. During this step, with the aid of the SALMO module (Joachimi et al. 2021), we also consider survey characteristics, such as the survey footprint, along with its spatial variability in depth and in its redshift distributions, as well as in the shape and shot noise, as described in Sect. 3.6. Thus, we obtained a full galaxy catalogue with shape information consistent with the input cosmology as well as with the previously mentioned survey properties.

The last steps of the forward simulation pipeline involve the post-processing and compression of the simulated data catalogue. In principle, any data vector could be obtained here, such as two-point shear statistics (Asgari et al. 2021; van den Busch et al. 2022), the shear fields (Porqueres et al. 2022; Loureiro et al. 2023), or even the full catalogues. However, in order to reduce the dimensionality of the data vector efficiently, while still retaining most information about the cosmology, similar to previous KiDS-1000 analyses (Loureiro et al. 2022; Tröster et al. 2022), we chose the observed angular power spectra, also known as ‘pseudo-Cls’ (Peebles 1973), as our data vector of choice. This allowed us to compare our SBI analysis to the results from the standard analysis of KiDS-1000 data using pseudo-Cls in Loureiro et al. (2022). The KiDS-SBI pipeline is, in principle, set up to make cosmological inferences from cosmic shear based on any statistic which could be derived from a galaxy catalogue, but the underlying statistical random fields are currently only capable of achieving percent-level accuracy for two-point statistics. Any higher-order statistics measured from the current forward simulations are not necessarily accurate enough to conduct inference. The statistics and their post-processing are defined in Sect. 4.2.

The forward simulations are designed to model weak gravitational lensing observations on the level of galaxy catalogues. The galaxy populations are sampled such that they trace the underlying log-normal random matter fields, while the galaxy shapes are lensed in accordance with the lensing potential of the matter fields along the line of sight. As the simulations do not model images, any image-level systematic effects, such as shear biases or variable depth, are included as spatially varying probability density functions.

⁷ Kilo-Degree Survey – Simulation-Based Inference; https://github.com/mwiet/kids_sbi

⁸ <https://github.com/joezuntz/cosmosis>

⁹ <https://github.com/KiDS-WL/kcap>

¹⁰ Generator for Large Scale Structure; <https://github.com/glass-dev/glass>

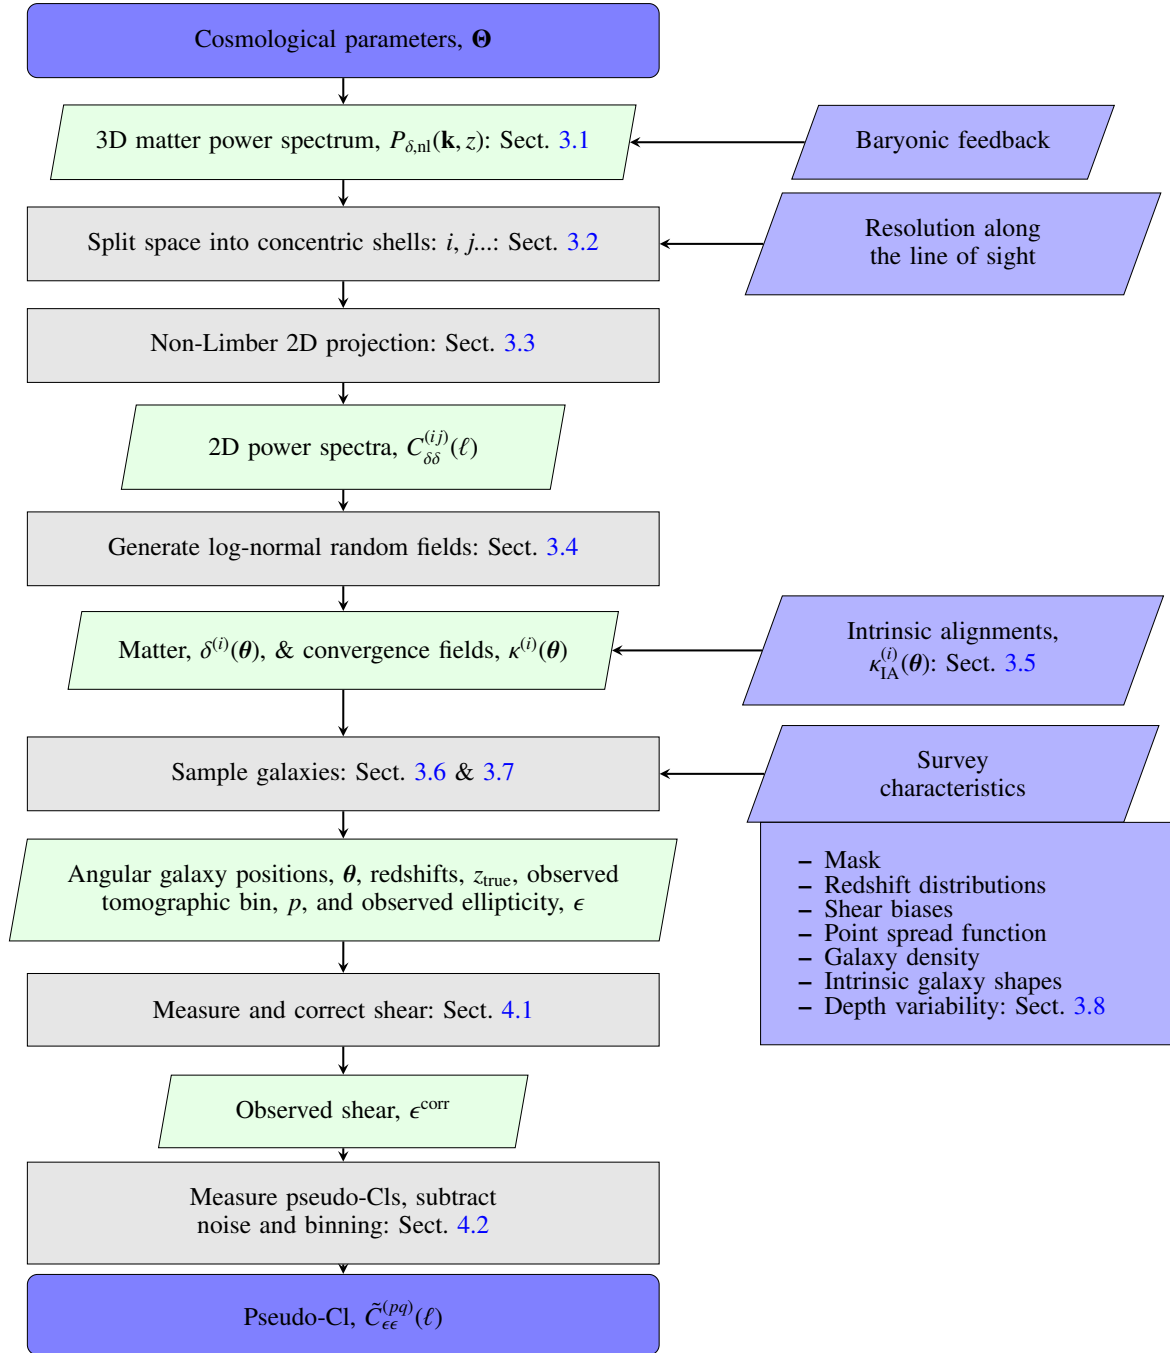


Fig. 3. Flowchart describing the steps in a single forward simulation of cosmic shear observables from cosmological parameters. The dark blue rounded boxes represent the inputs and outputs which are given to the simulation-based inference pipeline. The green slanted boxes represent relevant quantities which are calculated during the simulation. The grey rectangular boxes show steps in the calculations, while the blue slanted boxes show any (systematic) effects which are included. All variables are defined within the respective sections quoted in the diagram.

3.1. Cosmology dependence: 3D matter power spectrum

The basis of the forward simulations is the 3D matter power spectrum, $P_{\delta, \text{nl}}(k, z; \Theta)$. In the linear regime, the equal-time 3D matter power spectrum's cosmology dependence comes from the following relation:

$$P_{\delta, \text{nl}}(k, z; \Theta) = T^2(k; \Theta) D^2(z; \Theta) P_{\mathcal{R}}(k; \Theta), \quad (1)$$

where k is a wavenumber, z is redshift, Θ is the set of cosmological parameters, $P_{\mathcal{R}}(k; \Theta)$ is the primordial density fluctuation power spectrum, D is the growth factor of structure, T

is the transfer function between the matter overdensity field, $\delta(\mathbf{k}, z)$, and the primordial curvature fluctuation field, $\mathcal{R}(\mathbf{k}; \Theta)$, such that $\delta(\mathbf{k}, z) = T(\mathbf{k}; \Theta) D(z; \Theta) \mathcal{R}(\mathbf{k}; \Theta)$. To compute this, we make use of CAMB (Lewis et al. 2000; Lewis & Challinor 2002; Howlett et al. 2012). To stay in line with the main KiDS-1000 analysis (Joachimi et al. 2021; Asgari et al. 2021; van den Busch et al. 2022), we assume a normal neutrino hierarchy, while also assuming a fixed sum of neutrino masses of $\sum m_\nu = 0.06 \text{ eV}/c^2$.

We can compute the non-linear matter power spectrum via a non-perturbative model, HMCode-2016 (Mead et al. 2015, 2016). Although updated iterations of this model exist

(Mead et al. 2021), we limited our study to HMCode-2016, so our model remains comparable to the KiDS-1000 cosmic shear analysis. HMCode uses a halo model approach to incorporate the effects of baryonic feedback on the matter distribution. The main driving factor in this is the baryonic matter expelled by active galactic nuclei (AGNs). This suppresses the power at small scales, namely, large values of k , as a function of the amplitude of the halo mass-concentration relation, A_{bary} , and the halo bloating parameter, η_0 . We also fixed the relation between these two parameters to $\eta_0 = 0.98 - 0.12A_{\text{bary}}$ in accordance with Joudaki et al. (2018) and Joachimi et al. (2021), and treated A_{bary} as the only free parameter related to baryonic feedback.

We note that this treatment may lead to systematic biases at small scales, as the galaxies will be sampled from a matter power spectrum which is already modified by baryonic feedback. It would be more physically accurate to sample the galaxies from a matter power spectrum without baryonic non-linearities and subsequently add a perturbation to the matter density contrast fields due to baryonic feedback (similar to the approach used for the intrinsic alignments described in Sect. 3.5). Nevertheless, we do not expect this to cause a large discrepancy for the scales which we are probing in the cosmic shear signal ($\ell < 1500$). Field-level implementations of baryonic feedback already exist (Schneider & Teyssier 2015; Schneider et al. 2019, 2022; Porqueres et al. 2023) and they would constitute interesting avenues for future extensions to KiDS-SBI.

3.2. Working on the sphere

Astronomical observations can be viewed as images of a 3D space projected onto the inside of a 2D surface of a sphere at a given distance with the observer at its centre. Therefore, galaxy position measurements, most observational biases and some systematics are naturally expressed in spherical coordinates on the sky. For this reason, the forward model presented in this work is expressed in the same geometry.

The KiDS-SBI forward simulations use GLASS (Tessore et al. 2023) to model the underlying large-scale structure. Within GLASS, we chose to model the large-scale structure through log-normal random matter fields within non-overlapping concentric shells centred at the observer. Therefore, a given shell is a comoving volume which spans the full sky and has a finite width along the line of sight from one redshift, z_i , to another, z_{i+1} , such that $i \in \{1, 2, 3, \dots, N_{\text{shells}}\}$. We define the matter weight function $W^{(i)}(z)$ describing the distribution along the line of sight of the comoving volume spanned by a given shell as (Tessore et al. 2023),

$$W^{(i)}(z; \Theta) = \begin{cases} f_k^2(z; \Theta)/E(z; \Theta), & \text{if } z_i \leq z < z_{i+1}, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where $E(z; \Theta)$ is the dimensionless Hubble function and $f_k(z; \Theta)$ is the transverse comoving distance for a given redshift and cosmology, which is defined as

$$f_k(z; \Theta) = \begin{cases} \frac{c}{H_0 \sqrt{\Omega_k}} \sinh\left(\frac{\sqrt{\Omega_k} H_0}{c} \chi(z; \Theta)\right), & \text{if } \Omega_k > 0, \\ \chi(z; \Theta), & \text{if } \Omega_k = 0, \\ \frac{c}{H_0 \sqrt{|\Omega_k|}} \sin\left(\frac{\sqrt{|\Omega_k|} H_0}{c} \chi(z; \Theta)\right), & \text{if } \Omega_k < 0, \end{cases} \quad (3)$$

where H_0 is the Hubble constant, Ω_k is the curvature density parameter and c is the speed of light in a vacuum. As noted in Tessore et al. (2023), different choice of window function may

induce percent-level biases in the resulting shear power spectrum. We test for this in Appendix B, and find sub-percent level agreement between the measured shear from the simulations with this choice of window function and the shear signal expected from theory (see Appendix D).

A given shell should not be thicker than a few hundred Mpc of comoving distance along the line of sight for log-normal random fields to be an accurate enough representation of the distribution of large-scale structure (Xavier et al. 2016; Hall & Taylor 2022; Tessore et al. 2023). Otherwise, discretisation effects would smooth away large-scale structure (Tessore et al. 2023). Additionally, we need to be mindful of computational resources, because the number required of angular power spectra between the matter fields within each shell scales as $2N_{\text{shells}}$. Simultaneously, the simulation run-time for a given set of cosmological parameters increases with N_{shells} . To balance these factors, we aim to have as many shells along the line of sight as computationally feasible.

For an accurate weak lensing simulation at the relevant redshifts, $0.6 \lesssim z \lesssim 1.2$, we require shells that finely sample the lowest redshifts. The lensing signal in a given matter shell depends on the weighted sum of the matter fields of the shells with lower redshifts. At redshifts where lensing becomes more important as the lensing efficiency increases, prioritising thinner shells helps to reduce discretisation effects. We find that the following set of redshift limits for 19 shells are sufficient for the resolution needed for KiDS-1000, while still being efficient: {0, 0.04, 0.08, 0.12, 0.16, 0.2, 0.27, 0.34, 0.42, 0.5, 0.58, 0.66, 0.75, 0.88, 1.03, 1.19, 1.36, 1.55, 1.76, 2} (see Fig. 2).

3.3. Non-Limber projection

Although the 3D matter power spectrum characterises all cosmological dependence of the matter fields in our simulation, for it to define the cosmology dependence in the simulated matter fields, it has to be projected according to the geometry of the shells considered in Sect. 3.2. To determine the correlations between the matter fields of comoving spherical shells, we define the 2D angular matter power spectra for a given cosmology, Θ , and a given shell combination (ij) as follows

$$C_{\delta\delta, \ell}^{(ij)}(\Theta) = \langle \tilde{\delta}_{\ell m}^{(i)}(\Theta) \tilde{\delta}_{\ell m}^{(j)*}(\Theta) \rangle, \quad (4)$$

where $\tilde{\delta}_{\ell m}^{(i)}(\Theta)$ represents the harmonic coefficients defined through a spherical harmonic transform of the projection of the spin-0 matter field, $\delta^{(i)}$, as given by (Zaldarriaga & Seljak 1997; Reinecke 2011),

$$\tilde{\delta}_{\ell m}^{(i)}(\Theta) = \int d^2\theta \delta^{(i)}(\theta; \Theta) {}_0Y_{\ell m}^*(\theta), \quad (5)$$

where ${}_0Y_{\ell m}(\theta)$ are the spherical harmonics for a spin value of 0 as a function of the 2D spatial sky position θ . The projected matter field, $\delta^{(i)}(\theta; \Theta)$, can in turn be defined with respect to the underlying 3D matter overdensity field, $\delta(\theta, z; \Theta)$, as

$$\delta^{(i)}(\theta; \Theta) = \sum_{\ell m} \tilde{\delta}_{\ell m}^{(i)}(\Theta) {}_0Y_{\ell m}(\theta) = \int dz W^{(i)}(z; \Theta) \delta(\theta, z; \Theta), \quad (6)$$

where $W^{(i)}(z; \Theta)$ is the weight function of the given shell as defined in Eq. (2).

The 3D matter overdensity field's cosmology dependence can be quantified as follows

$$\langle \delta(\mathbf{k}, z; \Theta) \delta^*(\mathbf{k}', z'; \Theta) \rangle = (2\pi)^3 \delta_{\text{D}}(\mathbf{k} - \mathbf{k}') P_{\delta, \text{nl}}(k, z, z'; \Theta), \quad (7)$$

where $\delta(\mathbf{k}, z; \Theta)$ is the Fourier transform of the 3D matter overdensity field at z , $\delta(\theta, z, \Theta)$, while $\delta_D(\mathbf{k} - \mathbf{k}')$ is the Dirac delta function.

If we combine the relations in Eqs. (4)–(7) and, for the sake of simplicity, assume that $\Omega_k = 0$, specifically $f_k(z; \Theta) = \chi(z; \Theta)$, we can characterise the cosmology dependence of $C_{\delta\delta}^{(ij)}(\ell; \Theta)$ directly with respect to the 3D matter power spectrum, $P_{\delta, \text{nl}}(k, z, z'; \Theta)$, using the following relation

$$C_{\delta\delta}^{(ij)}(\ell; \Theta) = \frac{2}{\pi} \int d\chi W^{(i)}(z[\chi]; \Theta) \int d\chi' W^{(j)}(z[\chi']; \Theta) \int dk k^2 P_{\delta, \text{nl}}(k, z[\chi], z[\chi']; \Theta) j_\ell(k\chi) j_\ell(k\chi'), \quad (8)$$

where $\chi' \equiv z[\chi']$, $j_\ell(k\chi)$ are spherical Bessel functions of order ℓ , and $\ell \in \mathbb{Z}^{0+}$. We can simplify Eq. (8) by making the geometric approximation, that is taking the geometric mean of two equal-time power spectra (Castro et al. 2005; Kitching & Heavens 2017; Kilbinger et al. 2017; de la Bella et al. 2021), such that

$$P_{\delta, \text{nl}}(k, z[\chi], z[\chi']; \Theta) = \sqrt{P_{\delta, \text{nl}}(k, z[\chi]; \Theta) P_{\delta, \text{nl}}(k, z[\chi']; \Theta)}, \quad (9)$$

which has been shown to be an accurate approximation both in the linear and non-linear regimes (Kitching & Heavens 2017).

Additionally, it is common to make the so-called Limber approximation (Limber 1953; Kaiser 1992) which consists in approximating a spherical Bessel function as a single peak or delta function at $k\chi = \ell + 1/2$. This approximation is reasonably accurate in the auto-correlations for broad window functions and high angular modes, $C_{\delta\delta}^{(ii)}(\ell; \Theta)$, but depending on the thickness of the shells, the Limber approximation can lead to substantial biases, especially, at large scales (Lemos et al. 2017). However, when making the Limber approximation, all off-diagonal angular power spectra, $C_{\delta\delta}^{(ij)}$, where $i \neq j$, will be zero as long as their associated weight functions, $W^{(i)}(z[\chi]; \Theta)$, are non-overlapping, as is the case for the weights constructed in Eq. (2). This can be seen by looking at the full Taylor expansion of Eq. (8) as given by (LoVerde & Afshordi 2008)

$$C_{\delta\delta}^{(ij)}(\ell) = \int \frac{d\chi}{\chi^2} W^{(i)}(z[\chi]) W^{(j)}(z[\chi]) P_{\delta, \text{nl}}\left(\frac{\ell + 1/2}{\chi}, z[\chi]\right) \left\{ 1 - \frac{1}{(\ell + 1/2)^2} \left[\chi^2 \left(\frac{W^{(i)''}(z[\chi])}{W^{(i)}(z[\chi])} + \frac{W^{(j)''}(z[\chi])}{W^{(j)}(z[\chi])} \right) + \frac{\chi^3}{6} \left(\frac{W^{(i)'''}(z[\chi])}{W^{(i)}(z[\chi])} + \frac{W^{(j)'''}(z[\chi])}{W^{(j)}(z[\chi])} \right) \right] + \mathcal{O}[(\ell + 1/2)^{-4}] \right\}, \quad (10)$$

where each apostrophe, $'$, denotes a partial derivative with respect to χ , specifically $\partial/\partial\chi$, while we also omit the explicit cosmological dependence of the power spectrum and the weight functions here for the sake of clarity. The non-Limber integral therefore also contributes to the autocorrelations, where $|i - j| = 0$, particularly for small ℓ and the narrow weight functions considered here. If not considered this can induce a large bias in the angular power spectra of the shells which could significantly bias cosmological constraints from KiDS-1000. At the same time, we find that for the cross-correlations of adjacent shells, such that $|i - j| = 1$, the angular power can be of the order of $10^{-1} C_{\delta\delta}^{(ii)}$ or less. In this case, the simulated $\delta^{(i)}(\theta; \Theta)$ would have a percent-level bias while being void of any large-scale structure along the line of sight beyond the scales of individual shells. To avoid this, we chose to calculate the full non-Limber projection as given by Eq. (8). The correlations between shell pairs which are not

immediate neighbours, that is $|i - j| > 1$, contribute less than 0.1% to the overall angular power of a shell, so their effect is undetectable for KiDS.

Calculating the full non-Limber integral is a computationally expensive endeavour as it involves three numerical integrals over highly oscillatory Bessel functions. To address this, we present the Levin module¹¹. This is a novel Python module which implements the Levin integration method for oscillatory functions (Levin 1996) in the context of non-Limber integrals for weak lensing and galaxy clustering which is discussed in more detail in Appendix A.

The module allows us to divide the domain of the angular power spectra into three regions: a non-Limber domain, a second-order extended Limber domain and a Limber domain. In this work, we chose to split the calculation of $C_{\delta\delta}^{(ij)}(\ell)$ into two domains: non-Limber ($1 < \ell \leq \ell_{\text{max, nL}}$) and second-order extended Limber given in Eq. (10) for $\ell_{\text{max, nL}} < \ell \leq 30\,000$. Over the former domain, Levin numerically integrates the expression shown in Eq. (8) for the given matter power spectrum and the weights. Over the latter domain, we perform the numerical integration of Eq. (10) up to the second-order term in $(\ell + 1/2)^{-1}$. This reduces the computational resources needed for the integration without any substantial loss of accuracy as long as $\ell_{\text{max, nL}}$ is sufficiently large, since the residuals scale with $(\ell + 1/2)^{-4}$ (LoVerde & Afshordi 2008). With the shells shown in Fig. 2 weighted with Eq. (2), we obtain angular power spectra, $C_{\delta\delta}^{(ij)}(\ell; \Theta)$, shown in Fig. 4. We note that, in Fig. 4, the Limber extension to second order from Eq. (10) calculated for $\ell > 1200$ reduces roughly to the Limber approximation as the higher-order terms have a vanishing contribution to the overall angular power spectrum. In fact, in the cross-correlations between shells where $|i - j| > 0$, the second order term in Eq. (10) becomes vanishingly small for these values of ℓ . Consequently, we effectively apply a scale cut in the matter shell cross-correlations for $\ell > 1200$. We find that this does not bias our analysis, as the relative contribution cross-correlations to the total $C_{\delta\delta}^{(ij)}(\ell)$ of a given shell is less than 0.1% for $\ell \geq 1200$ for all redshifts. This ensures sub-percent level accuracy of the input matter fields across all scales considered in this analysis.

With this configuration, we find that we can obtain computationally efficient and accurate (consistent with CAMB within 0.1%) angular power spectra which describe the correlations between the matter fields within each spherical shell of the simulation. This is in line with testing conducted on the Levin module in previous applications (Zieser & Merkel 2016; Spurio Mancini et al. 2018a,b; Baleato Lizancos & White 2023).

3.4. Log-normal matter field simulations

With the input angular power spectra for a given cosmology computed as described in Appendix A, the next step in the forward simulations involves sampling random matter fields consistent with these correlations. These fields are generated with GLASS (Tessore et al. 2023). This framework allows us to create efficient and accurate random field simulations of large-scale structure matter fields as well as the associated weak gravitational lensing signals. Within the GLASS suite, many different choices can be made, while it also allows us to sample the angular positions, redshifts, shapes and shears of galaxies. Nevertheless, in this analysis, we only chose to use it to simulate the

¹¹ Levin; https://github.com/rreichke/nonLimber_matter_shells




Fig. 4. 2D angular matter power spectra, $C_{\delta\delta}^{(ij)}(\ell)$, projected with Levin describing the correlations between the large-scale structure within a set of 19 concentric shells spaced along the line of sight as shown in Fig. 2. The top panel shows the autocorrelations of all shells, i.e. $|i - j| = 0$. The bottom panel shows the correlations of each shell with its nearest neighbour, i.e. $|i - j| = 1$ ($|i - j| > 1$ are not shown as they do not have a large effect on matter fields within each shell, see Tessore et al. 2023). The colour of each line is given by the mean redshift of the i^{th} bin, $\bar{z}^{(i)}$. The underlying linear 3D matter power spectrum is based on flat Λ CDM assuming $\Omega_c = 0.05$, $\Omega_b = 0.28$, $\sigma_8 = 0.79$, $S_8 = 0.84$ and $H_0 = 67 \text{ km s}^{-1} \text{ Mpc}^{-1}$ calculated using CAMB (Lewis et al. 2000; Lewis & Challinor 2002; Howlett et al. 2012), while the non-linear contribution is calculated using HMcode-2016 (Mead et al. 2015, 2016) assuming $A_{\text{bary}} = 3.1$. The non-Limber projection done by Levin assumes $\ell_{\text{max,nL}} = 1200$ which causes the $C_{\delta\delta}^{(ij)}$ with $|i - j| > 0$ go to a value of zero for $\ell > \ell_{\text{max,nL}}$.

underlying matter overdensity fields, $\delta^{(i)}(\theta; \Theta)$, as well as the associated convergence fields for each shell, $\kappa^{(i)}(\theta; \Theta)$. These are then integrated along the line of sight to give the convergence field for a source within a given shell, $\kappa^{(i)}$.

It has been found that a good approximation to such fields is a log-normally distributed random field (which can be derived from a Gaussian random field). In fact, log-normal fields are even able to produce reasonable approximations to the three-point and four-point statistics measured from N-body simulations (Hall & Taylor 2022; Piras et al. 2023). For this reason, log-normal random fields are a common approximation used

for matter overdensity fields (Coles & Jones 1991; Böhm et al. 2017; Abramo et al. 2016, 2022) as well as for convergence fields (Hilbert et al. 2011; Clerkin et al. 2017; Giocoli et al. 2017; Gatti et al. 2020).

Log-normal fields within GLASS are based on Gaussian random fields. We can obtain these by taking the angular power spectra of the shells and assuming that they describe the correlations between log-normal fields, $\delta_{\log}^{(i)}(\theta; \Theta)$, as follows

$$\langle \delta_{\log}^{(i)}(\theta; \Theta) \delta_{\log}^{(j)*}(\theta; \Theta) \rangle = C_{\delta\delta}^{(ij)}(\theta; \Theta). \quad (11)$$

By defining the mean and the variance for $G_{\delta\delta}^{(ij)}(\theta; \Theta)$ as

$$\langle \delta_G^{(i)}(\theta; \Theta) \delta_G^{(j)*}(\theta; \Theta) \rangle = G_{\delta\delta}^{(ij)}(\theta; \Theta), \quad (12)$$

we can then sample $\delta_G^{(i)}(\theta; \Theta)$ for each shell. To achieve this with numerical methods, the 2D spherical shell of each layer is discretised using HEALPix pixels (Górski et al. 2005). This means that θ is discretised such that it can be mapped to a linearised variable given by θ_m where $m \in \{1, 2, 3, \dots, 12N_{\text{side}}^2\}$ and N_{side} is the HEALPix resolution parameter. To obtain the final matter overdensity fields, we use the following expression (Coles & Jones 1991; Kayo et al. 2001; Hilbert et al. 2011; Xavier et al. 2016),

$$\delta_{\log}^{(i)}(\theta_m; \Theta) = e^{\delta_G^{(i)}(\theta_m; \Theta)} - \lambda, \quad (13)$$

where λ is the shift of the log-normal distribution which is assumed to be $\lambda = 1$ for matter fields. For more details about this calculation, (see Tessore et al. 2023).

The layers of concentric log-normal matter overdensity fields within each shell give a full description of the large-scale structure. The resolution of this structure will of course be limited by the size of the shells along the line of sight, while being limited by the resolution of the discretised 2D pixels on the sphere. However, as discussed with regards to the resolution of shells in Sect. 3.2, if we chose a sufficiently large ℓ_{max} value up to which the input angular power spectra are calculated, we can accurately sample $\delta_{\log}^{(i)}(\theta_m; \Theta)$ up to an $N_{\text{side}} \sim 0.5\ell_{\text{max}}$ (Leistedt et al. 2013; Alonso et al. 2019).

As these simulations were performed prior to carrying out the weak gravitational lensing study, the next step in the forward simulations is to model how the image of galactic sources within a given shell will be distorted by weak gravitational lensing due to the matter overdensities the light encounters along the line of sight to the observer. Here, the geometry of the simulations aids us again, since the concentric volumes in comoving distance that each shell makes up intrinsically allow us to trace all possible light cones emanating from the observer.

We start from the definition of the convergence for a source located at θ and z along a continuous line of sight under the Born approximation (Bartelmann & Schneider 2001),

$$\kappa(\theta, z; \Theta) = \frac{3\Omega_m}{2} \int_0^z dz' \frac{f_k(z'; \Theta) [f_k(z'; \Theta) - f_k(z; \Theta)]}{f_k(z; \Theta)} \times \frac{1 + z'}{E(z'; \Theta)} \delta(\theta, z'; \Theta), \quad (14)$$

where Ω_m is the matter density fraction at $z = 0$, and $\Omega_m \in \Theta$. We then discretise the continuous 2D matter overdensity field along over θ using HEALPix pixels, while also discretising along the line of sight using Eq. (6) and defining a mean redshift, $\bar{z}^{(i)}$, for a given weighted shell as follows

$$\bar{z}^{(i)}(\Theta) = \frac{\int dz z W^{(i)}(z; \Theta)}{\int dz W^{(i)}(z; \Theta)}. \quad (15)$$

To accelerate the calculation of the discretised convergence fields for each shell, $\kappa^{(i)}$, within GLASS, it makes use of the fact that for all Robertson-Walker space-times, transverse comoving distances will scale in such a way that we can write down a recurrence relation from one interval to the next along the line of sight (Schneider 2016). This gives a recurrence relation for the convergence field, $\kappa^{(i)}$, within shells where $i \geq 2$, which only depends on $\kappa^{(i-1)}$, $\kappa^{(i-2)}$ and $\delta^{(i-1)}$ (see Tessore et al. 2023 for details). Such a relation allows us to calculate the convergence fields for all shells, while only holding three fields in memory at a given time. This greatly reduces the amount of computational resources needed per realisation, so the large amount of simulations at different cosmologies, Θ , needed for this analysis can be obtained with fewer resources when compared to forward models in previous analysis (see Sect. 5.1). Once we have all the ingredients to sample galaxies and to shear their shapes, while simultaneously having the flexibility to add any relevant field-level systematics, we can proceed to the next phase of the study.

3.5. Intrinsic alignments: Non-linear alignment model

An important effect to consider when modelling weak gravitational lensing is intrinsic alignments (IAs). This refers to the fact that correlations between source galaxy shapes in different parts of the sky may not only be caused by weak gravitational lensing due to the foreground matter distribution. This happens due to two local processes which occur irrespective of weak gravitational lensing: tidal alignments and tidal torquing. The latter is thought to arise when galaxy discs form perpendicular to the angular momentum axis of the tidal field, while tidal alignments occur when the gravitational tidal forces from neighbouring matter distributions cause the galaxies to align with the surrounding matter density field. This alignment leads to a coherent stretching or compression of the galaxy shapes as a function of the local matter overdensity field. Hence, the intrinsic ellipticities of source galaxies are systematically aligned with the underlying large-scale structure within which they form. Thus, the measured cosmic shear signal will be biased by these intrinsic alignments (Heavens et al. 2000a; King & Schneider 2002; Heymans & Heavens 2003; Bridle & King 2007).

To model this effect, we followed the methodology set out in Tessore et al. (2023). In this prescription, for a given intrinsic alignment model, we define an associated effective convergence field, $\kappa_{\text{IA}}^{(i)}$. This $\kappa_{\text{IA}}^{(i)}$ is a useful construct which describes the contribution to the observed weak lensing signal from IA under the Born approximation. Alternatively, an IA model could also directly be implemented into the dark matter distribution of the matter fields which could be an interesting avenue for future work.

Based on the fact that for the IA model, we take into consideration $\kappa_{\text{IA}}^{(i)} \propto \delta^{(i)}$, we can assume that, for a given shell, this field can be added linearly to the underlying convergence field due to weak lensing, as given by

$$\kappa^{(i)}(\theta_m; \Theta) \rightarrow \kappa^{(i)}(\theta_m; \Theta) + \kappa_{\text{IA}}^{(i)}(\theta_m; \Theta). \quad (16)$$

For simplicity and consistency with typical modelling assumptions in weak lensing surveys such as KiDS-1000 (Joachimi et al. 2021), we chose to model $\kappa_{\text{IA}}^{(i)}$ with the Non-Linear Alignment (NLA) model (Catelan et al. 2001; Hirata et al. 2004; Bridle & King 2007), so we consider tidal alignments and do not consider tidal torquing. The NLA model assumes that the bias in the shear signal from IAs is linearly

dependent on the projected local tidal field. The ‘non-linear’ part of the NLA model then simply refers to the fact that it has been found that the modelling of IAs is more accurate when modelling the underlying large-scale structure using a non-linear matter power spectrum, rather than a linear one (Bridle & King 2007, similar to our approach described in Sect. 3.1). This means that it is also proportional to the local matter overdensity field, which is given by (Hirata et al. 2004),

$$\kappa_{\text{IA}}^{(i)}(\theta_m; \Theta) = -A_{\text{IA}} \frac{C_1 \Omega_m \bar{\rho}_{\text{cr}}(\bar{z}^{(i)}; \Theta)}{D(\bar{z}^{(i)}; \Theta)} \delta^{(i)}(\theta_m; \Theta), \quad (17)$$

where A_{IA} is the intrinsic alignments amplitude (which we are treating as a nuisance parameter that is also sampled, see Sect. 5), C_1 is a normalisation constant which we set to $C_1 = 5 \times 10^{-14} h^{-2} M_{\odot}^{-1} \text{Mpc}^3$ in agreement with the IA measurements at low redshifts by SuperCOSMOS (Brown et al. 2002), $\bar{\rho}_{\text{cr}}(\bar{z}^{(i)}; \Theta)$ is the mean critical matter density as a function of redshift and $D(\bar{z}^{(i)}; \Theta)$ is the linear growth factor normalised to be unity at $z = 0$. We note that in Eq. (17), we do not consider an explicit redshift dependence which is often expressed as a power-law term (Joachimi et al. 2011). We omit this term, as it is not considered in the fiducial KiDS-1000 analysis (Asgari et al. 2021; van den Busch et al. 2022), and it has been found that at least for $z < 1$, the IA signals do not vary largely with redshift (Fortuna et al. 2021).

Nevertheless, the NLA model still comes with some caveats. Mainly, it neglects non-linearities which may not be captured in the underlying matter power spectrum, such as source density weighing, expressly the next-to-leading order perturbations proportional to the source density considered in a full tidal alignment model (Blazek et al. 2015). This effect can in fact produce a signal in the shape measurements comparable to the correction due to non-linear structure growth (Krause et al. 2016). Despite these limitations, the NLA model has been found to be accurate enough to model the IA signal seen by past and current weak lensing surveys (Joachimi et al. 2011; Blazek et al. 2011; Heymans et al. 2013; Krause et al. 2016; Hilbert et al. 2017), including KiDS-1000 (Fortuna et al. 2021).

3.6. Galaxy positions and redshifts

With the cosmological dependence of the underlying large-scale structure characterised and astrophysical effects such as intrinsic alignments and baryonic feedback included, the forward simulations still need to include observational biases from our instruments. We chose to apply these on the level of galaxies, namely, by biasing the sampled galaxy positions, redshifts, and shapes and shears in accordance with our models for the observational systematic effects.

When sampling the galaxies’ positions on the sky, we consider the survey footprint, Ω_{survey} , the galaxy number density which the survey can observe, n_{gal} , and its variability as a function of redshift and observational depth across the sky, namely, as a function of the local observing conditions. When sampling the galaxies’ redshifts, we take into consideration the redshift distributions calibrated from photometry for a given survey, $n(z)$. Lastly, when determining the observed galaxy shapes, we consider the intrinsic ellipticity dispersion of the galaxies, σ_e , the multiplicative and additive shear bias, while also including the effects on the measured galaxy shapes caused by the variation in the point spread function of the instrument as a function of the position in the sky.

Since we are interested in simulating a photometric survey, we must model the resolution along the line of sight realistically. This means that rather than modelling large-scale structure observables across N_{shells} concentric shells, they ought to be modelled across N_{tomo} tomographic bins. The tomographic bins of a given survey are defined by finite usually non-overlapping domains in photometric redshift, z_{ph} , in which all observed galaxies are grouped, such that $z_p < z_{\text{ph}} \leq z_{p+1}$, $p \in \{1, 2, 3, \dots, N_{\text{tomo}}\}$ ¹². The measured photometric redshift of a galaxy does not necessarily agree with its true/spectroscopic redshift, z . This means that galaxies which are grouped within a given photometric redshift bin, p , may not be located in this redshift range. To account for this, we can characterise a conditional redshift distribution for each tomographic bin $P(z|z_{\text{ph}})$ by mapping the measured z_{ph} to known spectroscopic redshift measurements which have a comparatively low bias in the redshift estimate (Hildebrandt et al. 2021; van den Busch et al. 2022). Therefore, a given galaxy is in an associated pixel, m , shell, i , and tomographic bin, p . The probability of the galaxy being in tomographic bin, p , while it is sampled in shell, i , then depends on $P(z|z_{\text{ph}})$.

In line with the KiDS-1000 analysis, we define the effective galaxy number density as follows

$$n_{\text{eff}}^{(p)} \equiv \frac{1}{|\Omega_{\text{survey}}|} \frac{\left[\sum_{i \in (p)} w_i \right]^2}{\sum_{i \in (p)} w_i^2}, \quad (18)$$

where $|\Omega_{\text{survey}}|$ is the effective area covered by the survey footprint, and w_i is the *lensfit* weight for a given galaxy, i . We note that the effective galaxy density within the KiDS-1000 gold sample is not the same as the galaxy density, n_{gal} , when not correcting for the *lensfit* weights for each galaxy. However, as KiDS-SBI does not model galaxy weights, in the forward simulations $n_{\text{eff}} = n_{\text{gal}}$ since $w_i = 1 \forall i$.

Similarly, we define the galaxy intrinsic ellipticity dispersion, σ_ϵ , within a given tomographic bin as follows

$$\sigma_\epsilon^{(p)2} \equiv \frac{\sum_{i \in (p)} w_i^2 (\epsilon_{\text{obs},i,1}^2 + \epsilon_{\text{obs},i,2}^2)}{(1 + \overline{M}^{(p)})^2 \sum_{i \in (p)} w_i^2}, \quad (19)$$

where $\epsilon_{\text{obs},i,1}$ and $\epsilon_{\text{obs},i,2}$ are the two components of the observed galaxy ellipticity, and $\overline{M}^{(p)}$ is the mean multiplicative shear bias in tomographic bin p as calibrated in (Giblin et al. 2021). Applying Eqs. (18) and (19) to the KiDS-1000 gold sample, we obtain the values shown in Table 1. We note that due to the way σ_ϵ is defined in Eq. (19), it incorporates both the noise due to intrinsic galaxy shapes as well as due to image noise in the survey.

Although it does not have detectable effects on the cosmic shear signal, the KiDS-SBI forward simulations sample galaxy positions such that they trace the underlying matter fields, so spatially varying systematics may be accurately applied. Using the SALMO¹³ framework (Joachimi et al. 2021) for simulating galaxy catalogues, we Poisson-sample galaxies within each HEALPix pixel, m , on the sky using the following expectation for the galaxy counts

¹² Note: we use the indices p and q for tomographic bins and the indices i and j for underlying shells.

¹³ Speedy Acquisition for Lensing and Matter Observables: <https://github.com/Linc-tw/salmo>

Table 1. Parameters used to sample galaxies and their shapes in line with the expectations for the KiDS-1000 gold sample for each tomographic bin (from S1 to S5).

Bin	\bar{n}_{eff} [arcmin ⁻²]	$\bar{\sigma}_\epsilon / \sqrt{2}$	$a_{n_{\text{eff}}}$	$b_{n_{\text{eff}}}$	a_{σ_ϵ} $\times 10^3$	b_{σ_ϵ}
S1	0.62	0.27	-0.035	0.72	1.81	0.267
S2	1.18	0.26	-0.042	1.30	1.30	0.257
S3	1.85	0.28	-0.243	2.58	-0.86	0.280
S4	1.26	0.25	-0.250	2.01	-0.62	0.255
S5	1.31	0.27	-0.416	2.56	2.39	0.261

Notes. \bar{n}_{eff} is the mean galaxy number density for a given tomographic bin, $\bar{\sigma}_\epsilon / \sqrt{2}$ is the mean per-component shape dispersion, $a_{n_{\text{eff}}}$ and $b_{n_{\text{eff}}}$ are the slope and y-intercept, respectively, for the linear interpolation of the galaxy density as a function of the root-mean-square of the background noise in the KiDS catalogue, σ_{rms} , according to Eq. (28), while a_{σ_ϵ} and b_{σ_ϵ} are the parameters to linearly interpolate σ_ϵ from σ_{rms} according to Eq. (29).

$$\langle N_m^{(i)(p)} \rangle(\Theta) = w_m(\Omega_{\text{survey}}) \left[1 + b^{(i)} \delta^{(i)}(\theta_m; \Theta) \right] P_m(p|i) n_{\text{gal},m}^{(p)} A_{\text{pix},m}, \quad (20)$$

where $w_m(\Omega_{\text{survey}})$ is the weight for a given pixel m which we take to be unity if a pixel is within the survey's mask/footprint, Ω_{survey} , and zero if it is not, $b^{(i)}$ is galaxy bias parameter within a given shell which we assume to be $b^{(i)} = 1 \forall i$ for the purposes of this cosmic shear analysis, $P_m(p|i)$ is the probability of a galaxy in pixel m being detected within tomographic bin p while being within shell i (in other words, a discretised version of $P_m(z|z_{\text{ph}})$), $n_{\text{gal},m}^{(p)}$ is the observed galaxy number per unit area for a given tomographic bin p and a given pixel m (here the pixel dependence accounts for the spatial variability of the detection limit of galaxies as a function of observing conditions, the number of exposures, etc.; see Sect. 3.8), and $A_{\text{pix},m}$ is the area of pixel m . As $n_{\text{gal},m}^{(p)}$ changes due to variations in observational depth, it is more or less likely to observe high-redshift galaxies. Hence, a change in depth from one pixel to the next can change the observed redshift distribution (see Sect. 3.8).

Taking $w_m(\Omega_{\text{survey}})$, $P_m(p|i)$ and $n_{\text{gal},m}^{(p)}$ from the photometric survey's measurements, we can Poisson sample $N_m^{(i)(p)}(\Theta)$ given the expectation value shown in Eq. (20) from which we can get the expected number of galaxy counts within each pixel of a tomographic bin p by taking

$$N_m^{(p)}(\Theta) = \sum_{i=1}^{N_{\text{shells}}} N_m^{(i)(p)}(\Theta). \quad (21)$$

With that, we then randomly sample each galaxy's right ascension and declination within each pixel m assuming a uniform distribution. This means that below the scales of pixels, the simulations do not have any information on galaxy clustering.

To sample each galaxy's redshift, we can use the fact that we know exactly how many galaxies are in each shell i for a given tomographic bin p and pixel m from $N_m^{(i)(p)}(\Theta)$. We can therefore randomly assign an index i and p to each sampled galaxy in accordance with that number. We can then randomly sample a specific z_{spec} for each galaxy by assuming that the galaxies are uniformly distributed within a given shell. Again, this means that weak lensing and galaxy clustering measurements from the simulated galaxies do not contain any information

about the large-scale structure along the line of sight below the scales of shells.

3.7. Galaxy shapes

To determine the galaxy shapes and shears, we begin by assigning intrinsic ellipticities. We define a given galaxy's shape through a complex ellipticity, ϵ , and any galaxy ellipticity in this work as follows

$$\epsilon = \epsilon_1 + i \epsilon_2, \quad (22)$$

where $i \equiv \sqrt{-1}$. The intrinsic ϵ_1 and ϵ_2 of each galaxy are then sampled within SALMO as two independent normal random variates with zero mean and a variance equal to $\sigma_{\epsilon,m}^{(p)2}$. As with the galaxy density, $n_{\text{gal},m}^{(p)}$, the shape dispersion, $\sigma_{\epsilon,m}^{(p)}$, can also vary with tomographic bin and as a function of the position on the sky due to the effects of the spatial variability of observational depth. See Sect. 3.8 for a more detailed discussion. The ellipticity of each source galaxy is then altered through weak gravitational lensing by the matter along the line of sight. The exact shape distortion is given by (Seitz & Schneider 1996; Hu 2000),

$$\epsilon_{\text{lensed}}(\Theta) = \frac{\epsilon_{\text{int}} + g(\Theta)}{1 + g^*(\Theta)\epsilon_{\text{int}}}, \quad (23)$$

where $\epsilon_{\text{lensed}}(\Theta)$ is the lensed galaxy ellipticity, ϵ_{int} is the intrinsic galaxy ellipticity and $g(\Theta)$ is the reduced shear ($g \in \mathbb{C}$), given by

$$g(\Theta) = \frac{\gamma(\Theta)}{1 - \kappa(\Theta)}, \quad (24)$$

where $\gamma(\Theta)$ is the shear factor ($\gamma \in \mathbb{C}$). The use of Eq. (23) is an extension to the KiDS-1000 analysis which assumed the reduced shear approximation, where $\epsilon_{\text{lensed}} \approx \epsilon_{\text{int}} + g$. In any case, this should cause a negligible change in the measured shear considering the sensitivity of KiDS-1000 (Joachimi et al. 2021). To calculate the shear field, it is useful to decompose the spin-0 convergence field into spherical harmonics as follows:

$$\kappa_{\ell m}^{(i)}(\Theta) = \int d^2\theta \kappa^{(i)}(\theta; \Theta)_0 Y_{\ell m}^*(\Theta), \quad (25)$$

where $\kappa_{\ell m}^{(i)}(\Theta)$ are the harmonic coefficients of the convergence field of a given shell i . We can then use the following relation to define the shear field's harmonic coefficients, $\gamma_{\ell m}^{(i)}$, given by (Peebles 1973),

$$\gamma_{\ell m}^{(i)}(\Theta) = -\left(\frac{(\ell+2)(\ell-1)}{\ell(\ell+1)}\right)^{1/2} \kappa_{\ell m}^{(i)}(\Theta). \quad (26)$$

The $\gamma_{\ell m}^{(i)}(\Theta)$ coefficients can then be used to calculate the spin-2 discrete shear field, $\gamma^{(i)}(\theta_m, \Theta)$. Having defined the effective $\kappa^{(i)}(\theta_m, \Theta)$ in Sect. 3.4 and 3.5, while also defining $\gamma^{(i)}(\theta_m, \Theta)$, we can assign a value of κ and $\gamma = \gamma_1 + i \gamma_2$ to each galaxy sampled according to the procedure described in Sect. 3.6. This is done by taking the value of $\kappa^{(i)}(\theta_m, \Theta)$ and $\gamma^{(i)}(\theta_m, \Theta)$ within the shell i and pixel m in which a given galaxy is located. Between neighbouring pixels, the values of $\kappa^{(i)}(\theta_m, \Theta)$ and $\gamma^{(i)}(\theta_m, \Theta)$ are linearly interpolated. This choice may not be optimal (Tessore et al. 2023), but we find that at the precision level of KiDS-1000, it does not lead to any significant biases (see Appendix B). We then use Eq. (24) to calculate the associated reduced shear, $g(\Theta)$, as

well as the lensed ellipticity, $\epsilon_{\text{lensed}}(\Theta)$, by combining g with the intrinsic ellipticity using Eq. (23).

However, this lensed ellipticity, $\epsilon_{\text{lensed}}(\Theta)$, is not the observed shape measurement in a weak gravitational lensing survey such as KiDS-1000. Additionally, the shape measurement may be distorted by instrumental effects or shape modelling inaccuracies. Some relevant effects are selection biases, noise biases from a low signal-to-noise ratio, biases in the galaxy weights, artefacts, non-linear CCD responses, asymmetries in the point spread function, etc. (see for example Mandelbaum 2018 for a review). To account for this, it is common to parametrically estimate any residual systematics which may be affecting the shape measurements (Hildebrandt et al. 2017, 2021; Zuntz et al. 2018; Giblin et al. 2021). Within the forward simulations, we use the first-order parametric expansion of the observed galaxy shapes given by Heymans et al. (2006) as follows

$$\epsilon_{\text{obs},i \in m}^{(p)} = (1 + M^{(p)})\epsilon_{\text{lensed},i} + \alpha^{(p)}\epsilon_{\text{PSF},m} + \beta^{(p)}\delta\epsilon_{\text{PSF},m} + c^{(p)}, \quad (27)$$

where i is the index for a single galaxy within tomographic bin p and pixel m , $M^{(p)}$ is the multiplicative shear bias as measured for tomographic bin, p , we have $\alpha^{(p)}$ as the fraction of the PSF ellipticity, which remains in the shear estimator, $\epsilon_{\text{PSF}}(m)$ is the local PSF measured within pixel m (see Fig. 6 for a map of ϵ_{PSF} in KiDS-1000), $\beta^{(p)}$ gives the amplitude of the shear bias due to residuals which are not taken into account by the PSF model, $\delta\epsilon_{\text{PSF}}$ represents the residuals in question, and $c^{(p)}$ is the additive shear bias within a given tomographic bin, p . We note that spatially varying multiplicative shear biases, $M^{(p)}$, can also arise, but previous work has found that the mean bias across all positions is sufficient to accurately model the effect (Kitching et al. 2019). In accordance with previous findings (Giblin et al. 2021), we take $\beta^{(p)} = 0 \forall p$. The amplitudes of the remaining systematics as calibrated by Giblin et al. (2021) are shown in Fig. 5. It becomes apparent that these effects are relatively small which explains why only the multiplicative and additive bias is taken into consideration in the modelling on the fiducial KiDS-1000 analysis (Giblin et al. 2021; Joachimi et al. 2021; Asgari et al. 2021; Heymans et al. 2021). Additionally, the measured values of $M^{(p)}$, $\alpha^{(p)}$, and $c^{(p)}$ have associated uncertainties. To take this uncertainty into account and make sure it is represented in the effective likelihood, rather than using a fixed value, we randomly sample different values for $M^{(p)}$, $\alpha^{(p)}$ and $c^{(p)}$ for each realisation of the forward simulations from a normal distribution as shown in Fig. 5. We note that we do not explicitly marginalise over these parameters, but as they are sampled randomly from predetermined probability density distributions, so that their associated uncertainty is propagated into the effective likelihood. This is similar to how the seed for each log-normal random field, and the galaxy positions which populate the associated matter fields are varied from one forward simulation to the next, such that cosmic variance enters the effective likelihood as well.

As a result, for given cosmological parameters, Θ , the KiDS-SBI forward simulations produce a galaxy catalogue containing the galaxies' position on the sky, their spectroscopic redshift, the tomographic bin in which they are detected and their observed ellipticities, while taking into account many relevant systematics which are considered in the modelling of weak gravitational lensing measurements.

3.8. Variable depth

As alluded to in Sect. 3.6 and 3.7, when randomly sampling galaxies and their intrinsic ellipticities, the parameters which

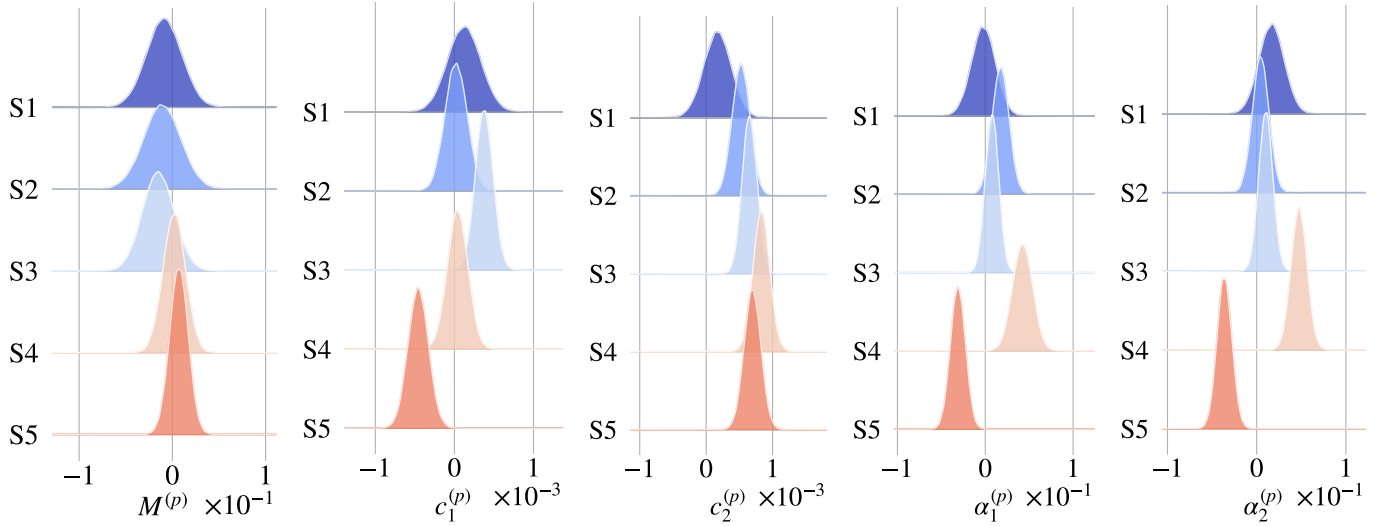


Fig. 5. Distributions for each tomographic bin (S1 to S5) from which the shear bias parameters from KiDS-1000 (Giblin et al. 2021) shown in Eq. (27) are sampled. The first panel from the left shows the multiplicative shear bias, $M^{(p)}$. The second and third panel show the real and imaginary part of the additive shear bias, $c_1^{(p)}$ and $c_2^{(p)}$, respectively. The fourth and fifth panel show the real and imaginary part of the amplitude of the shear bias due to variations in the point-spread function, $\alpha_1^{(p)}$ and $\alpha_2^{(p)}$, respectively.




Fig. 6. Cartesian spatial map ($N_{\text{side}} = 1024$) of the observed magnitude of the point-spread function ellipticities, $|\epsilon_{\text{PSF}}| = \sqrt{\epsilon_{\text{PSF},1}^2 + \epsilon_{\text{PSF},2}^2}$, throughout the KiDS-1000 North field in the upper panel and the KiDS-1000 South field in the lower panel. ϵ_{PSF} is added to the lensed galaxy shapes in the forward simulations within KiDS-SBI in accordance with Eq. (27).

characterise the probability distributions depend on the location on the sky, namely, the pixel m in question. Specifically, the number of galaxies which is sampled (in other words, ‘observed’) within a given pixel depends on the average galaxy density, $n_{\text{gal},m}^{(p)}$, measured within that pixel. The galaxy density varies from pixel to pixel as it depends on the observational depth of the survey in different parts of the sky, and on anisotropies and time variations in the atmospheric seeing (particularly, for surveys such as KiDS which only visit each pointing in the sky once at a given time). Since many parts of the sky are observed at different times, as is the case for ground-based telescopes such as the VLT Survey Telescope and the Visible and the Infrared Survey Telescope for Astronomy (VISTA) used for KiDS-1000, the atmospheric conditions and background light will be different for each pointing for a given exposure time. Thus, the depth of observations may change as seeing, the point spread functions and/or the background flux varies (for example, due to moonlight, zodiacal light, or galactic absorption).

In addition, surveys such as KiDS often have overlapping pointings in order to ensure that the footprint is observed without gaps. However, this implies that galactic sources which happen to be located near the edge of a pointing will be observed more often in different overlapping fields than a source located in the centre of a pointing. Consequently, near the edge of a pointing, there will tend to be a higher S/N which allows for deeper observations. Since these systematic effects in the observational depth can happen at fixed scales and/or have certain periodicities, they can induce significant systematic effects into galaxy clustering and weak gravitational lensing signals. For cosmic shear in KiDS-1000, the bias can be near 1% in the signal and an average of $\sim 10\%$ in the standard deviation of cosmic shear observables (Heydenreich et al. 2020; Joachimi et al. 2021).

As the systematics modify the local observational depth, the galaxies which can be observed within a given pixel are selected according to their magnitude. This will systematically bias the dispersion of intrinsic ellipticities observed, $\sigma_{\epsilon,m}^{(p)}$, and further




Fig. 7. Cartesian spatial map ($N_{\text{side}} = 1024$) of root-mean-square of the observed background noise, σ_{rms} , throughout the KiDS-1000 North field in the upper panel and the KiDS-1000 South field in the lower panel.




Fig. 8. Spatial map ($N_{\text{side}} = 4096$) in a Cartesian projection of root-mean-square of the observed background noise, σ_{rms} , for a $5^\circ \times 5^\circ$ patch of the KiDS-1000 North field at a right ascension (RA) of 180° and a declination (DEC) of 0° .

exacerbate the effect of observational depth variability on the cosmic shear signal.

To take all these effects into consideration within the forward simulations, we calibrate the spatial variability of $n_{\text{eff},m}^{(p)}$ and $\sigma_{\epsilon,m}^{(p)}$ from the KiDS-1000 measurements directly. Rather than defining maps over the entire survey footprint of both quantities for each tomographic bin, we use a map of a direct estimator of observational depth which also correlates with $n_{\text{eff},m}^{(p)}$ and $\sigma_{\epsilon,m}^{(p)}$: the root-mean square of the background noise, σ_{rms} . This quantity gives the detection threshold of counts above the background, and is the dimensionless analogue to the detection threshold typically expressed in mag arcsec^{-2} . Since it directly characterises the detection probability of an object as a function of its angular position, σ_{rms} incorporates variations in the magnitude limit or the seeing (Johnston et al. 2021). We found this quantity to be a good indicator of variable depth as it correlates well with the measured $n_{\text{eff},m}^{(p)}$ and $\sigma_{\epsilon,m}^{(p)}$, while being uncorrelated with the measured photometric redshifts and magnitudes in KiDS-1000.

A similar approach to modelling variable depth has been taken in Joachimi et al. (2021). However, in that analysis, the variable depth along the KiDS-1000 footprint was modelled with a direct estimate of the magnitude limit in the r -band (the band in which galaxy shape measurements are made in KiDS). We have found that the magnitude limit associated with a given galaxy in the KiDS-1000 gold sample is correlated with the galaxy's observed r -band magnitude as well as the photometric redshift estimate for the galaxy. This means that a selection according to the magnitude limit in the r -band can bias cosmological estimates, as the selection is not independent of the cosmic shear signal. In contrast, σ_{rms} is uncorrelated with the r -band magnitude measurements and the photometric redshifts in the KiDS-1000 gold sample, while still being highly correlated with the local magnitude limit in the r -band. Thus, we conclude it is a more direct parametrisation of the variations across the survey footprint of the galaxy selection.

As can be seen in Fig. 7 and 8, a map of σ_{rms} across the KiDS North and South fields clearly shows variations over different pointings as well as with areas of increased overlap. In addition, Fig. 9 shows that $n_{\text{eff}}^{(p)}$ and $\sigma_{\epsilon}^{(p)}$ vary linearly with σ_{rms} for KiDS-1000 data. Thanks to this fact, we do not need to model the variation of $n_{\text{eff},m}^{(p)}$ and $\sigma_{\epsilon,m}^{(p)}$ with the HEALPix pixel m with individual maps for each tomographic bin p . Instead, using the linear relations shown in Fig. 9, we may assign a value for $n_{\text{eff},m}^{(p)}$ and $\sigma_{\epsilon,m}^{(p)}$ for each $\sigma_{\text{rms},m}$ (specifically the mean value of level within each pixel m contained within the footprint observed by KiDS-1000) given that

$$n_{\text{eff},m}^{(p)} = a_{n_{\text{eff}}}^{(p)} \sigma_{\text{rms},m} + b_{n_{\text{eff}}}^{(p)}, \quad (28)$$

$$\sigma_{\epsilon,m}^{(p)} = a_{\sigma_{\epsilon}}^{(p)} \sigma_{\text{rms},m} + b_{\sigma_{\epsilon}}^{(p)}, \quad (29)$$

where $a_{n_{\text{eff}}}^{(p)}$ and $a_{\sigma_{\epsilon}}^{(p)}$ are the slopes of the linear fits shown in Fig. 9 of galaxy density and galaxy shape dispersion for a given tomographic bin, respectively, and $b_{n_{\text{eff}}}^{(p)}$ and $b_{\sigma_{\epsilon}}^{(p)}$ are the associated y-intercepts. The values of these parameters are shown in Table 1. This drastically helps the performance of the forward simulations, as this allows us to just use a single map of σ_{rms} to model the variable depth rather than requiring $2N_{\text{tomo}}$ maps for it. Such a linear relation may not exist in other weak lensing surveys, but it should still be possible to numerically calibrate the relations dependence of $n_{\text{eff}}^{(p)}$ and $\sigma_{\epsilon}^{(p)}$ on σ_{rms} to define an interpolation.

These inhomogeneities in the galaxy selection across the survey footprint also change the redshift distribution of galaxies. As

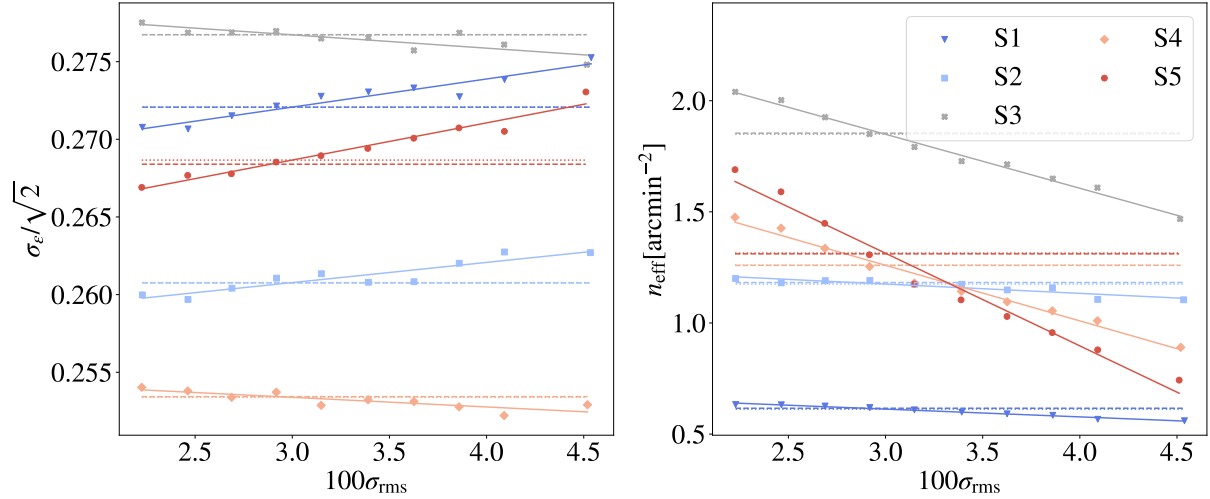


Fig. 9. The dependence of the per-component Gaussian shape dispersion, $\sigma_\epsilon/\sqrt{2}$, (left panel) and the galaxy density, n_{eff} , (right panel) on the root-mean square of the background noise, σ_{rms} in the KiDS-1000 DR4 data. For both panels, the data points represent the mean σ_ϵ or n_{eff} of ten equi-populated bins in σ_{rms} with their boundaries in $\{1.70, 2.33, 2.57, 2.80, 3.04, 3.28, 3.52, 3.76, 4.00, 4.23, 12.96\}$. The solid line shows the linear fit to the aforementioned data points of their respective tomographic bin according to Eqs. (28) and (29). The parameters obtained from this fit are given in Table 1. The dotted horizontal lines show the mean values of σ_ϵ and n_{eff} calculated from the galaxy samples with variable depth per tomographic bin, while the dashed horizontal lines show the values of σ_ϵ and n_{eff} for the respective galaxy samples without any spatial variations in the observational depth. Both of these lines agree exceptionally well by construction, so that for some source bins there is not any observable difference between them.




Fig. 10. In the upper panels, the normalised redshift distributions, $P(z)$, for each tomographic bin (S1 to S5). The redshift distribution from the entire KiDS-1000 DR4 galaxy sample, $P_{\text{total}}(z)$, is shown in black, while the other ten redshift distributions are derived from 10 equi-populated subsamples of DR4 based on their observational depth (i.e. the mean value of the root-mean-square of the background noise, σ_{rms}) which is shown with its respective colour. The lower panels show the associated residual change in the redshift distributions with respect to $P_{\text{total}}(z)$ per unit redshift. It is apparent that variable depth mostly affects the source distributions at high redshifts, while the effect tends to decrease the mean of the redshift distribution with increasing σ_{rms} .

σ_{rms} decreases, the local magnitude limit increases and fainter galaxies can be detected by the survey. As more distant galaxies with high redshifts also tend to be fainter, variable depth can shift the redshift distributions of the local galaxy population as can be seen in Fig. 10. The redshift distributions shown in Fig. 10 have been calculated using the same approach as in KiDS-1000 where the photometric redshifts are calibrated from spectroscopic samples using self-organising maps (Wright et al. 2019). The self-organising map has been applied to ten equi-populated subsamples of each tomographic bin of the KiDS-1000 gold sample which are binned in $100\sigma_{\text{rms}}$ along the following boundaries: $\{1.7, 2.33, 2.57, 2.8, 3.04, 3.28, 3.52, 3.76, 4, 4.23, 12.96\}$. To account for variable depth in our forward model, depending on

the value of σ_{rms} in a given pixel, m , on the survey footprint, we sample galaxy redshifts in that pixel from the associated redshift distribution from Fig. 10. This effect is important to model as an unaccounted shift in the redshift distribution can lead to significant biases in the cosmic shear signal (Heydenreich et al. 2020; Baleato Lizancos & White 2023).

To summarise, in KiDS-SBI, we account for the effects of spatial variations in the observational depth on galaxy number density, galaxy shape dispersion and redshift distributions. Moreover, we also consider anisotropies in the shear bias by including distortions to the lensed galaxy ellipticities calibrated from the spatial variation in the point-spread function ellipticity in the KiDS-1000 gold sample (see Sect. 3.7).

We note that the galaxy selection function considered here does not account for systematic effects which have been previously corrected in KiDS data, such as source blending. Thanks to the selection and calibration of the data discussed in Sect. 2, the KiDS-1000 source catalogue as well as the simulations thereof can be treated as a sample of non-overlapping sources. Any residual effects from source blending which remain after the measurement and calibration stages of the data are included within the shear biases and weights. Consequently, when quantifying the variable depth of the survey, any residual effect on blending is compensated by the inclusion of the galaxy weights and the multiplicative shear bias in the definitions of the effective galaxy density and the intrinsic ellipticity dispersion of galaxies, in Eqs. (18) and (19).

4. Measurements

After sampling the galaxy shapes and shear, then applying all relevant instrumental systematics within KiDS-SBI, we can go on to measure the galaxy shapes, along with the cosmic shear summary statistics in the same manner as the measurements done for the KiDS-1000 gold sample's weak lensing catalogue. The only difference between the measurements of the forward-modelled catalogues and the KiDS-1000 gold sample is the values of the galaxy weights, w_i . In the forward simulations, the galaxy weights are set to unity for all sampled galaxies, $w_i = 1 \forall i$, while the weights used in the measurements from the KiDS-1000 gold sample are given by the `lensfit` weights (see details in Sect. 2). To avoid biasing the statistics of the galaxy populations in the forward simulations with respect to the real data, we sample the galaxy positions and shapes from the galaxy densities, n_{gal} , and intrinsic shape dispersions, σ_ϵ , where the corrections for the `lensfit` weight and the multiplicative shear bias have been applied, as defined in Eqs. (18) and (19).

4.1. Shape measurements

From the KiDS-1000 gold sample data as well as the forward simulated catalogues, we measure the observed galaxy ellipticity, ϵ_{obs} , for each galaxy per tomographic bin. To this end, we apply the same corrections to the shear measurements as described in Giblin et al. (2021). The observed shape measurements are corrected as follows

$$\epsilon_{\text{obs},i}^{\text{corr}(p)}(\Theta) = \frac{1}{1 + \overline{M}^{(p)}} \frac{\sum_{i \in p} w_i (\epsilon_{\text{obs},i}(\Theta) - \langle \epsilon_{\text{obs}}(\Theta) \rangle)}{\sum_{i \in p} w_i}, \quad (30)$$

where w_i is the galaxy `lensfit` weight for galaxy i in the KiDS-1000 gold sample (for the case of the simulations, $w_i = 1 \forall i$), $\langle \epsilon_{\text{obs}}(\Theta) \rangle$ compensates for any additive biases (Asgari et al. 2019), that is the $c^{(p)}$ term and the mean of $\alpha^{(p)}$ in Eq. (27), and it is defined as the mean observed shape of all galaxies within tomographic bin p , and $\overline{M}^{(p)}$ is the mean multiplicative shear bias measured for each tomographic bin p as shown in Fig. 5. We note that $\overline{M}^{(p)}$ is not necessarily the same value as the value of $M^{(p)}$ applied in Eq. (27). For each instance of the simulations, a different $M^{(p)}$ is drawn from a Gaussian probability distribution with mean $\overline{M}^{(p)}$ and the standard deviations shown in Fig. 5. Regardless of the value of $M^{(p)}$ drawn, the shear measurements of each simulation are corrected by the same mean $\overline{M}^{(p)}$. Any discrepancy between these values is going to introduce noise in the simulations which accounts for the uncertainty on the shear bias measurements.

4.2. Observed angular power spectra

With a full simulated cosmic shear catalogue akin to the KiDS-1000 gold sample, we can map any set of cosmological and astrophysical parameters, Θ , to a corresponding set of cosmic shear measurements from the KiDS survey as required for simulation-based inference. Although theoretically possible, it is computationally impractical to conduct cosmological inference at the level of catalogues. For this reason, it is useful to compress the catalogues down to useful statistics which still retain most of the relevant cosmological information about the underlying large-scale structure.

We chose to compress the catalogues down to two-point statistics. The main reason for this is that the forward simulation pipeline is designed to be only percent-level accurate in two-point statistics. Higher-order correlations of the shear field are captured to a certain degree within the simulations, but their accuracy is limited (Hall & Taylor 2022; Piras et al. 2023). In addition, the use of two-point statistics assures that the only distinguishing feature of our analysis when compared to the fiducial cosmic shear KiDS-1000 analyses (Loureiro et al. 2022; Asgari et al. 2021) is the dropping of the assumption of a Gaussian likelihood, so we can carry out a direct comparison between the analyses. This way we can ensure that any non-Gaussianities in our likelihood are either attributed to inherent non-Gaussianities in the likelihood of two-point statistics or non-Gaussianities induced by systematic effects.

Due to its computational efficiency when compared to spatial two-point correlation functions, we chose pseudo-CIs as our two-point statistic. We decompose the observed shear field, $\epsilon_{\text{obs}}^{\text{corr}(p)}(\theta; \Theta)$, into its curl-free E-modes and its divergence-free B-modes as follows

$$\epsilon_{\text{obs}}^{\text{corr}(p)}(\theta; \Theta, \Omega_{\text{survey}}) = \sum_{\ell=0}^{\ell_{\text{max}}} \sum_{m=-\ell}^{\ell} (\tilde{E}_{\ell m}^{(p)}(\Theta) + i \tilde{B}_{\ell m}^{(p)}(\Theta)) {}_2Y_{\ell m}(\theta), \quad (31)$$

where ${}_2Y_{\ell m}(\theta)$ are the spin-2 spherical harmonic functions which define an orthonormal basis for the observed shear field, such that

$$\int d^2\theta {}_2Y_{\ell m}(\theta) {}_2Y_{\ell' m'}^*(\theta) = \delta_{\ell \ell'} \delta_{m m'}, \quad (32)$$

while $\tilde{E}_{\ell m}^{(p)}(\Theta)$ and $\tilde{B}_{\ell m}^{(p)}(\Theta)$ are spherical harmonic coefficients of the curl-free and the divergence-free observed shear fields within each tomographic bin, respectively, and δ^K is the Kronecker delta. These coefficients are defined as follows:

$$\tilde{E}_{\ell m}^{(p)}(\Theta) = \frac{1}{2} \int d^2\theta [\epsilon_{\text{obs}}^{\text{corr}(p)}(\theta; \Theta) {}_2Y_{\ell m}^*(\theta) + \epsilon_{\text{obs}}^{\text{corr}(p)*}(\theta; \Theta) {}_2Y_{\ell m}(\theta)], \quad (33)$$

$$\tilde{B}_{\ell m}^{(p)}(\Theta) = \frac{-i}{2} \int d^2\theta [\epsilon_{\text{obs}}^{\text{corr}(p)}(\theta; \Theta) {}_2Y_{\ell m}(\theta) - \epsilon_{\text{obs}}^{\text{corr}(p)*}(\theta; \Theta) {}_2Y_{\ell m}^*(\theta)]. \quad (34)$$

From these coefficients, we then calculate the pseudo-CIs, $\tilde{C}_{\epsilon\epsilon}^{(pq)}(\ell; \Theta)$, as follows

$$\tilde{C}_{\epsilon\epsilon}^{(pq)}(\ell; \Theta) = \begin{pmatrix} \tilde{C}_{\epsilon\epsilon}^{EE(pq)}(\ell; \Theta) \\ \tilde{C}_{\epsilon\epsilon}^{EB(pq)}(\ell; \Theta) \\ \tilde{C}_{\epsilon\epsilon}^{BB(pq)}(\ell; \Theta) \end{pmatrix} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \begin{pmatrix} \tilde{E}_{\ell m}^{(p)} \tilde{E}_{\ell m}^{(q)*}(\Theta) \\ \tilde{E}_{\ell m}^{(p)} \tilde{B}_{\ell m}^{(q)*}(\Theta) \\ \tilde{B}_{\ell m}^{(p)} \tilde{B}_{\ell m}^{(q)*}(\Theta) \end{pmatrix}, \quad (35)$$

where $\mu \in \{1, 2, 3\}$ and 1 stands for the EE component, 2 for the EB component, and 3 for the BB component. We note that $\tilde{C}_{\epsilon\epsilon}^{EB}(\ell; \Theta) = \tilde{C}_{\epsilon\epsilon}^{BE}(\ell; \Theta)$, since $\tilde{E}_{\ell m} \tilde{B}_{\ell m}^*(\Theta) = \tilde{B}_{\ell m} \tilde{E}_{\ell m}^*(\Theta)$.

Going forward, similar to other previous analyses (Hikage et al. 2019; Loureiro et al. 2022; Tröster et al. 2022), we only take into consideration $\tilde{C}_{\epsilon\epsilon}^{EE(pq)}(\ell; \Theta)$, so any cosmological signal which may have ended up in $\tilde{C}_{\epsilon\epsilon}^{EB(pq)}(\ell; \Theta)$ or $\tilde{C}_{\epsilon\epsilon}^{BB(pq)}(\ell; \Theta)$ because of EE to EB, or EE to BB mode mixing will be lost. $\tilde{C}_{\epsilon\epsilon}^{EE(pq)}(\ell; \Theta)$ is expected to still contain all of the cosmological cosmic shear signal because the shear field resulting from a scalar gravitational field in Λ CDM is predicted to be a curl-free field (Bartelmann & Schneider 2001; Hikage et al. 2011; Kilbinger 2015).

To account for the EE signal from the correlations between the intrinsic galaxy shapes, namely the shape noise bias, we subtract the mean shape noise power spectrum as follows

$$\tilde{C}_{\epsilon\epsilon}^{(pq)}(\ell; \Theta) = \tilde{C}_{\epsilon\epsilon}^{EE(pq)}(\ell; \Theta) - \delta_{pq}^K \langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle, \quad (36)$$

where $\tilde{C}_{\epsilon\epsilon}^{(pq)}(\ell; \Theta)$ is the E-mode pseudo angular power spectrum for cosmic shear with the shape noise bias subtracted, while $\langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle$ is the mean of the curl-free angular power spectrum of the shape noise. The latter is estimated as follows (Becker et al. 2016; Hikage et al. 2019; Nicola et al. 2021; Loureiro et al. 2022),

$$\tilde{C}_{\text{noise}}^{EE(pp)}(\ell) = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \tilde{E}_{\ell m}^{\text{rand}(p)}(\Theta) \tilde{E}_{\ell m}^{\text{rand}(p)*}(\Theta), \quad (37)$$

where $\tilde{E}_{\ell m}^{\text{rand}(p)}$ is the curl-free spherical harmonic coefficient of the randomly rotated shear values, $\epsilon_{\text{obs}}^{\text{rand}(p)}(\theta_m; \Theta)$. In turn, we define this field from the following galaxy shear values as follows:

$$\epsilon_{\text{obs},i}^{\text{rand}}(\Theta) = \epsilon_{\text{obs},i,1}^{\text{rand}}(\Theta) + i \epsilon_{\text{obs},i,2}^{\text{rand}}(\Theta), \quad (38)$$

$$\epsilon_{\text{obs},i,1}^{\text{rand}} = \epsilon_{\text{obs},i,1}^{\text{corr}} \cos(\theta_{\text{rand},i}) - \epsilon_{\text{obs},i,2}^{\text{corr}} \sin(\theta_{\text{rand},i}), \quad (39)$$

$$\epsilon_{\text{obs},i,2}^{\text{rand}} = \epsilon_{\text{obs},i,2}^{\text{corr}} \cos(\theta_{\text{rand},i}) + \epsilon_{\text{obs},i,1}^{\text{corr}} \sin(\theta_{\text{rand},i}), \quad (40)$$

where $\theta_{\text{rand},i}$ is a randomly drawn angle for each galaxy i from a uniform distribution where $\theta_{\text{rand},i} \in [0, 2\pi)$. To calculate the mean of the curl-free angular power spectrum of the shape noise, $\langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle$, we take mean of all modes in $\tilde{C}_{\text{noise}}^{EE(pq)}(\ell)$ as follows

$$\langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle = \frac{\sum_{\ell=\ell_{\min}}^{\ell_{\max}} (2\ell + 1) \tilde{C}_{\text{noise}}^{EE(pq)}(\ell)}{\sum_{\ell=\ell_{\min}}^{\ell_{\max}} (2\ell + 1)}, \quad (41)$$

where the mean is weighted by a factor of $2\ell + 1$ to account for the fact that each angular scale, ℓ , contains $2\ell + 1$ modes, m . We find setting $\ell_{\min} \sim 200$ with $\ell_{\max} = 1500$ yields optimal results, as high angular scales are typically shape noise-dominated, and thus put better constraints on the average shape noise bias. In fact, we find that the measured average shape noise bias, $\langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle$, agrees within $<0.5\%$ with the underlying shape noise signal computed directly from the known intrinsic galaxy shapes. Additionally, we find that $\langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle$ is self-consistent with $\langle \tilde{C}_{\text{noise}}^{BB(pq)}(\ell) \rangle$ within $<0.01\%$. The introduction of a weighted mean given by Eq. (41) allows us to achieve this level of precision with only a single random rotation of the galaxies' shapes. Instead, previous approaches relied on creating many different instances of random rotations and then computed $\langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle$ by taking the average over all the instances while weighting each angular scale, ℓ , equally (see for instance Loureiro et al. 2022). When weighting each ℓ equally, the low- ℓ

scales have an equally strong pull on the estimator as the high- ℓ scales which increases the statistical uncertainty on $\langle \tilde{C}_{\text{noise}}^{EE(pq)}(\ell) \rangle$ as the large scales are mostly dominated by the large-scale structure signal. For this reason, it is common to take the mean over many realisations of the shape noise to reduce the variance of the mean. Although such approaches achieve similar levels of precision when estimating the shape noise bias as the approach shown in this work, computing the angular power spectra for hundreds of realisations of randomly rotated galaxy shape catalogues can become computationally expensive and time-consuming. Hence, it is not feasible for an SBI analysis which requires $\sim 10^4$ realisations and we opt to use the weighted mean of a single realisation of a random rotation to estimate the shape noise bias.

Furthermore, we bin the noise-free observed angular power spectrum, $\tilde{C}_{\epsilon\epsilon}^{(pq)}(\ell; \Theta)$, into 8 log-spaced bins between $\ell = 76$ and $\ell = 1500$ in line with Loureiro et al. (2022). For this, we chose the pseudo-Cl binning scheme described in Brown et al. (2005) given by

$$\tilde{C}_{\epsilon\epsilon,L}^{(pq)}(\Theta) = \frac{1}{2\pi} \sum_{\ell=\ell_L}^{\ell_{L+1}} \frac{\ell(\ell+1)}{(\ell_{L+1} - \ell_L)} \tilde{C}_{\epsilon\epsilon}^{(pq)}(\ell; \Theta), \quad (42)$$

where ℓ_L and ℓ_{L+1} are the lower and upper limits of the L^{th} bin, respectively.

We note that we did not choose to deconvolve the pseudo-Cl's to estimate the full-sky E-mode angular power spectra for cosmic shear. We find that with the sky-coverage of KiDS-1000 and the complexity of the geometry of the KiDS-1000 footprint, the mixing matrix (see Appendix D) is not necessarily invertible, and the deconvolution is not single-valued.

We therefore obtain, for a given forward simulation, a cosmic shear measurement like the one shown in Fig. 1.1 for a single run of the forward simulations. The calculation of the data vector calculated from the theory prediction shown in Fig. 1.1 is described in Appendix D. As can be seen from Fig. 1.1, the measured cosmic shear signal is consistent with theory predictions which we find is the case at all cosmologies throughout the prior volume (see Appendix B). This shows that the log-normal random fields accurately recover the two-point statistics of the galaxy populations throughout parameter space, as expected.

As seen in Fig. 1.1, for the five tomographic bins in KiDS-1000, we obtain $8 \times N_{\text{tomo}}(N_{\text{tomo}} + 1)/2$ data points in $\tilde{C}_{\epsilon\epsilon,L}^{(pq)}(\Theta)$, namely, a 120D data vector. The analysis choices described within Sect. 3 and 4 in conjunction constitute the forward model assumed within this SBI analysis, and it is henceforth labelled as the anisotropic systematics model, in reference to the fact that the model includes novel features such as variable depth, PSF variation, and dropping of the reduced shear approximation.

To summarise, each data vector depends on 12 parameters in Θ : 5 cosmological parameters (σ_8 , the root-mean-square matter fluctuation over 8 Mpc/ h_0 ; ω_b , baryonic matter density; ω_c , cold dark matter density; n_s , the scalar spectral index of the primordial density fluctuation power spectrum, $P_{\mathcal{R}}$; and h_0^{14} , the normalised Hubble constant), and 7 astrophysical and nuisance parameters related to systematics: A_{bary} , the baryonic feedback amplitude within the non-linear 3D matter power spectrum; A_{IA} , the galaxy intrinsic alignment amplitude of the NLA model; and five correlated δ_z parameters which define the shift in the mean of the source redshift distribution, $P(z|z_{\text{ph}})$, of each tomographic bin. All other parameters on which the simulation depends are

¹⁴ Note: throughout we deviate from standard notation and denote the dimensionless Hubble constant as h_0 .

fixed from run to run (such as $\omega_k = 0$, the equation of state of dark energy, $w = -1$, $\sum m_\nu = 0.06$ eV, etc.). The exceptions to this are the seed used to sample the matter overdensities and the galaxies, and the amplitudes of the shear biases as described in Sect. 3.7. These values are varied from run to run in order to simulate cosmic variance, shape noise as well as the uncertainty in the shear bias, respectively.

5. Simulation-based inference (SBI)

The final data vector as defined by the forward simulations is used to determine the effective likelihood of the data using sequential neural likelihood estimation based on the density estimation likelihood-free inference package (pyDELFI¹⁵, Alsing et al. 2019, or DELFI henceforth to denote the method and software package) using the same analysis pipeline outlined in L23.

To implement this, we create a bespoke interface that allows DELFI to be the sampler for CosmoSIS (Zuntz et al. 2015). Initially, the sampler chooses a fixed number of points within the considered parameter space based on the hypercube-generating algorithm from PyDOE¹⁶ with some minor modifications (see L23 for more details). Subsequently, as is outlined in Fig. 11, the measured pseudo-CIs for each evaluation of the parameters on the hypercube are compressed further using score compression (see Sect. 5.3). This reduces the dimensionality of the data vector to the size of the parameter vector (in this case, seven dimensions). The compressed data is then used to train neural density estimators through DELFI (see Sect. 5.4).

As the ensemble of neural density estimators evaluates the sampling distribution throughout the hypercube, through the use of active learning, DELFI finds areas in parameter space where the likelihood is undersampled to select new parameter vectors (see Alsing et al. 2019 for details). As found in L23, this can halve the number of simulations needed to accurately train DELFI by drawing more simulations in the parameter region of interest when compared to a Latin hypercube alone. Once the ensemble of neural density estimators has reached the stopping criterion, controlled for over-fitting by comparison to a hold-out validation data set, we can calculate an effective likelihood from the learnt sampling distribution conditioned on the observed data vector (which can be a mock or measured from real data). We can then obtain a posterior distribution by sampling the product of the likelihood and the prior. This approach of learning the sampling distribution conditioned on the observed compressed data to yield an effective likelihood has the advantage that if the prior or data vector were to be changed, there is nothing new that must be re-trained within DELFI.

We refer to the model choices that we make within KiDS-SBI as the anisotropic systematics model. The model follows the choices outlined in Sects. 3 and 4. It is designed to be consistent with previous analyses of KiDS-1000, specifically, with Loureiro et al. (2022), while adding additional realism. Both Loureiro et al. (2022) and this work assume a flat Λ CDM cosmology to model the cosmic shear signal, both use pseudo-CIs as their data vector of choice, and both consider systematics such as multiplicative shear bias, variable depth in the uncertainty and intrinsic alignments in the signal. The anisotropic systematics model differs from Loureiro et al. (2022) in that it also includes variable depth in the signal modelling (as described in Sect. 3.8), it considers the effect of intrinsic alignments on the likelihood,

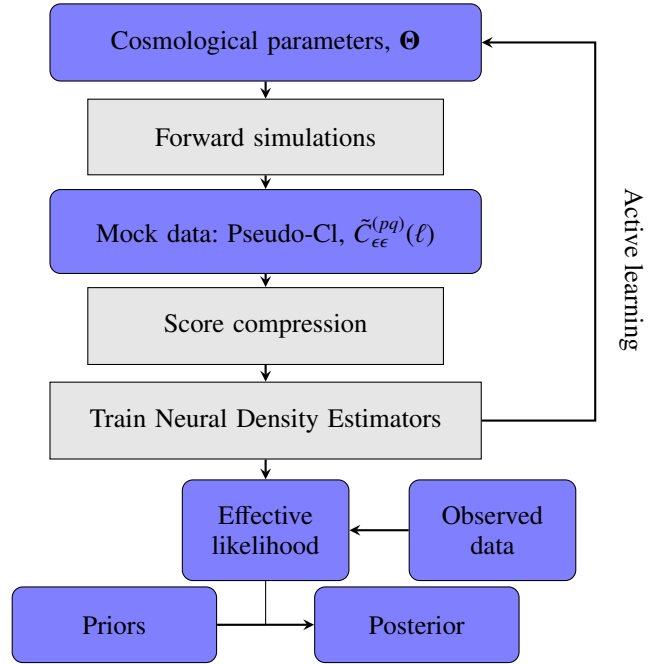


Fig. 11. Flowchart describing the structure of the simulation-based inference pipeline. The dark blue rounded boxes represent the inputs and outputs which are given to the simulation-based inference pipeline. The grey rectangular boxes show steps in the inference pipeline.

and it takes into account the variance in the additive and PSF shear biases (as described in Sect. 3.7). This is a consequence of the fact that in KiDS-SBI the effects which are modelled in the signal are intrinsically considered in the uncertainty model. In Loureiro et al. (2022), these effects are only considered separately as they pertain to the signal or the uncertainty, not both at once.

5.1. Pipeline performance

A single such evaluation of the forward simulations from cosmological parameters to the shear pseudo-CIs as described in Sect. 3 runs within ~ 20 minutes on a single core (with $N_{\text{shells}} = 19$, $N_{\text{tomo}} = 5$ and $N_{\text{side}} = 1024$). We chose this spatial resolution N_{side} , because it implies that the measured two-point statistics should be representative of the input angular power spectra up to an $\ell_{\text{max}} \sim 2 \times N_{\text{side}}$ (Leistedt et al. 2013; Alonso et al. 2019). In addition, for the KiDS-1000's galaxy number density, the pixel size at an $N_{\text{side}} = 1024$ is sufficient for almost all pixels to contain at least one galaxy. This suppresses mode-mixing in the pseudo-CIs due to random masking of empty pixels which were observed, as at this resolution all pixels are on average populated by at least one galaxy (see Appendix D).

In comparison (as can be seen in the bar chart in Fig. 12) log-normal random field simulations similar to the ones presented in Joachimi et al. (2021), take ~ 280 minutes to compute a single forward simulation when run with the same accuracy and precision settings on a single core. In practice, the simulations described in Joachimi et al. (2021) are used exclusively for the characterisation of a numerical covariance matrix. Hence, only a single evaluation of the underlying power spectrum from CAMB is necessary for every simulation. At the same time, the Joachimi et al. (2021) simulations have a lower resolution of shells along the line of sight, so the runtime of FLASK

¹⁵ <https://github.com/justinalsing/pydelfi>

¹⁶ <https://github.com/tisimst/pyDOE>




Fig. 12. Bar chart comparing the run-time of a single evaluation of KiDS-SBI (above) versus a single evaluation of a simulation based on Joachimi et al. (2021) (below), both on a single core ($N_{\text{shells}} = 19$, $N_{\text{tomo}} = 5$ and $N_{\text{side}} = 1024$). Both suites of simulations use CAMB (Lewis et al. 2000; Lewis & Challinor 2002; Howlett et al. 2012) to compute the 3D matter power spectrum. For the reference simulations, we use the non-Limber projection built into CAMB with $limber_phi_lmin = 1200$ rather than Levin with $\ell_{\text{max,nL}} = 1200$. We run FLASK (Xavier et al. 2016) rather than GLASS (Tessore et al. 2023) to compute the underlying matter and convergence fields of each of the 19 shells. Subsequently, we sample galaxies using SALMO in both cases, and then calculate the spatial two-point correlation functions, $\xi_{\pm}(\theta)$, rather than calculating the angular power spectra, $\tilde{C}(\ell)$. To calculate $\xi_{\pm}(\theta)$ in the reference simulations, we use TreeCorr (Jarvis et al. 2004).

(Xavier et al. 2016) in the original setting is lower than the runtime indicated in Fig. 12. We also note that the runtime of TreeCorr (Jarvis et al. 2004) could be drastically reduced from the single-core runtime shown in Fig. 12 by increasing the number of cores and/or by pixelising the galaxy shears prior to calculating the spatial two-point correlation functions, $\xi(\theta)$.

These gains in speed are driven by four main factors: the fast non-Limber integration using the Levin method, the use of recurrence relations to calculate the convergence within GLASS, the choice of using a quick-to-compute summary statistic like angular power spectra, $\tilde{C}(\ell)$, and the efficient subtraction of the shape noise bias from a single random rotation of the galaxy shapes. This is a substantial improvement with respect to other typical codes used for such simulations, while also improving accuracy with respect to previous models by including systematics, improving the resolution along the line of sight, and dropping common approximations like the Limber and reduced shear approximation. Thanks to this, it becomes feasible to increase the precision of our simulations as necessary and add the realism discussed in Sect. 3.6, 3.7, and 3.8, while still being able to compute the $\sim 10^4$ forward simulations needed to adequately characterise the effective likelihood L23.

5.2. Parameters and priors

The parameters which are varied in the simulation-based inference in this work are shown in Table 2. Their priors match those of the previous KiDS-1000 cosmic shear analyses (Asgari et al. 2021; Heymans et al. 2021; Loureiro et al. 2022; van den Busch et al. 2022; Tröster et al. 2022).

To avoid overinformative priors, most are flat top-hat functions. For the same reason, the top-hat priors in S_8 , h_0 , ω_b , A_{IA} , and A_{bary} are chosen to be wide. The prior on n_s spans a smaller range around the theoretical value of unity for scale-invariant primordial fluctuations. This avoids artefacts within the prior volume as this parameter is not well constrained by weak gravitational lensing. The prior on ω_c is defined to be consistent with a range in $\Omega_m \in [0.188, 0.408]$, where the limits are given by

the $\pm 5\sigma$ intervals of the marginal constraints from independent measurements of luminosity distance to Type Ia Supernovae in Scolnic et al. (2018). See Joachimi et al. (2021) for more details on the motivations for the chosen priors. The priors on the nuisance parameters capturing any shifts in the mean of each of the five tomographic bins, δ_z , are given by a multivariate Gaussian, since the δ_z of a given tomographic bin is not independent of the shifts in the other bins. This is quantified by the covariance, C_z , as estimated in Hildebrandt et al. (2021).

5.3. Score compression

To improve the computational efficiency of the analysis and to facilitate the use of DELFI, we compress each of the 120-dimensional data vectors measured from each forward simulation using score compression as described in Alsing et al. (2019).

If the likelihood is known a priori, we can compress a given data vector to a summary of the same dimensionality as the degrees of freedom in the assumed model, such that Fisher information is conserved (Zablocki & Dodelson 2016; Alsing & Wandelt 2018; Alsing et al. 2018, 2019). This allows us to compress the data vector, \mathbf{d} , down to a vector, \mathbf{t} , with the dimensionality of $|\Theta|$ as follows

$$\mathbf{t} = \nabla \mathcal{L}(\mathbf{d}|\Theta_*), \quad (43)$$

where \mathcal{L} is the log-likelihood distribution evaluated at the fiducial cosmology, Θ_* .

In this work, the exact form of the likelihood is not known a priori, as the main motivation is to characterise the form of the effective likelihood. However, it is still possible to perform score compression on the data by assuming an analytical form for the likelihood. To this end, we assume a Gaussian likelihood with a fixed covariance matrix from 1600 forward simulations at the fiducial cosmology given in Table 2 and evaluate its gradient through a five-point stencil near the fiducial cosmology. In this form with a Gaussian likelihood, this compression is equivalent to MOPED (Heavens et al. 2000b) or a linear compression based on Karhunen-Loève eigenvalue decomposition (Tegmark et al. 1997).

This compression is optimal if the true likelihood is Gaussian, which it is asymptotically near the peak of the likelihood and the chosen fiducial set of parameters is equal to the true parameter values. The downside of this is that the compression can lose information if this is not the case. Nevertheless, it was found in L23 that, in KiDS-SBI, the score compression is robust to suboptimal choices of fiducial parameter values and data covariance.

Additionally, we find that the 12 parameters shown in Table 2 are not necessary to characterise the degrees of freedom of the 120-dimensional cosmic shear pseudo-CIs. As all δ_z are broadly consistent with zero, we may compress all data vectors to a 7D summary which is still capable of capturing all the complexity in the data. In any case, the δ_z parameters are still explicitly varied within the simulations to propagate any uncertainties on the mean of the tomographic bins. For simplicity, henceforth all posteriors are implicitly marginalised over the five δ_z parameters.

5.4. Density estimation likelihood-free inference (DELFI)

To estimate the effective likelihood for the anisotropic systematics model, we train neural density estimators (NDEs) using DELFI (Alsing et al. 2019) such that the NDEs learn the sampling distribution between the compressed data and the input

Table 2. Parameters that are varied within the simulation-based inference pipeline and their priors.

Parameter	Symbol	Prior type	Prior range	Initial fiducial	Final fiducial
Density fluctuation amp.	S_8	Flat	[0.1, 1.3]	0.760	0.739
Hubble constant	h_0	Flat	[0.64, 0.82]	0.767	0.653
Cold dark matter density	ω_c	Flat	[0.051, 0.255]	0.118	0.135
Baryonic matter density	ω_b	Flat	[0.019, 0.026]	0.026	0.023
Scalar spectral index	n_s	Flat	[0.84, 1.1]	0.901	0.972
Intrinsic alignment amp.	A_{IA}	Flat	[-6, 6]	0.264	0.603
Baryon feedback amp.	A_{bary}	Flat	[2, 3.13]	3.10	2.66
Redshift displacement	δ_z	Gaussian	$\mathcal{N}(\mathbf{0}, \mathbf{C}_z)$	$\mathbf{0}$	$\mathbf{0}$
Multiplicative shear bias	$M^{(p)}$	Gaussian	$\mathcal{N}(\overline{M}^{(p)}, \sigma_M^{(p)})$	$\overline{M}^{(p)}$	$\overline{M}^{(p)}$
Additive shear bias	$c_{1,2}^{(p)}$	Gaussian	$\mathcal{N}(\overline{c}_{1,2}^{(p)}, \sigma_{c_{1,2}}^{(p)})$	$\overline{c}_{1,2}^{(p)}$	$\overline{c}_{1,2}^{(p)}$
PSF variation shear bias	$\alpha_{1,2}^{(p)}$	Gaussian	$\mathcal{N}(\overline{\alpha}_{1,2}^{(p)}, \sigma_{\alpha_{1,2}}^{(p)})$	$\overline{\alpha}_{1,2}^{(p)}$	$\overline{\alpha}_{1,2}^{(p)}$

Notes. The prior ranges are selected to be exactly in line with previous KiDS-1000 analyses (Asgari et al. 2021; Heymans et al. 2021; Loureiro et al. 2022; van den Busch et al. 2022; Tröster et al. 2022). The upper five rows show the cosmological parameters of interest, while the three rows below show the nuisance parameters which quantify systematic biases. The last three rows show shear bias parameters for each tomographic bin, p , as given in Fig. 5 which are sampled implicitly within each simulation by drawing from the prior distribution, so any posterior is pre-marginalised over these parameters. For flat priors, the lower and upper limits of the normalised rectangular function defines the prior. For the Gaussian prior on δ_z , we use a 5D multivariate Gaussian with its mean at the zero vector and the covariance, \mathbf{C}_z , defined by the one estimated in Hildebrandt et al. (2021). We note that for simplicity, the δ_z are implicitly marginalised throughout this analysis. The fiducial values given here constitute the fiducial parameter choice for the score compression described in Sect. 5.3.

cosmological parameters. To this end, we initialise an ensemble of six independent conditional Masked Autoregressive Flows (MAFs; Papamakarios et al. 2017). Each MAF is made up of between three and eight Masked Autoencoders for Density Estimation (MADEs; Germain et al. 2015) each with two hidden layers of 50 neurons. For a detailed definition and description of these network architectures, see Alsing et al. (2019).

As the first step in the exploration of the prior space, the ensemble of these neural density estimators (NDEs) initially learns the sampling distribution from a set 2000 points on a Latin hypercube. Based on this, we then make use of the active learning feature within DELFI to sample additional points in parameter space which efficiently contribute to learning the effective likelihood within prior space. To ensure that the sampling distribution density, $P(\mathbf{t}|\boldsymbol{\Theta}, \mathbf{w})$ is well learnt, recursive sampling continues until the validation loss stops decreasing for 20 training epochs, where the loss function is defined as

$$-\ln[U(\mathbf{w})] = -\sum_{i=1}^{N_{\text{sims}}} \ln[P(\mathbf{t}_i|\boldsymbol{\Theta}_i, \mathbf{w})], \quad (44)$$

where N_{sims} is the number of forward simulations and the loss function, $-\ln[U(\mathbf{w})]$, is defined such that it is a Monte Carlo estimate of the Kullback-Leibler divergence between the learnt effective likelihood and the true effective likelihood (see Alsing et al. 2019 for a more detailed explanation). This process is repeated over an ensemble of many independent neural density estimator networks, where each network learns the conditional likelihood independently. The final estimate of the effective likelihood is then given by the weighted sum of the learnt distributions from each neural density estimator.

We find that for the anisotropic systematics model, this occurs already after 5000 realisations. We note that the NDEs require fewer realistic KiDS-SBI forward simulations to minimise the loss function as is the case in L23 where the forward simulations were idealised random samples from a Gaussian distribution. In any case, to ensure that the prior space is densely

Table 3. Main inferred cosmological parameters varied within the anisotropic systematics model from the KiDS-1000 gold sample.

Parameter	Marginal $\pm 1\sigma$	MAP \pm PJ-HPD
S_8	0.731 ± 0.033	$0.743^{+0.015}_{-0.051}$
σ_8	$0.73^{+0.10}_{-0.21}$	$0.72^{+0.09}_{-0.20}$
Ω_m	$0.337^{+0.097}_{-0.150}$	$0.323^{+0.220}_{-0.060}$

Notes. The second column shows the marginal as well as the upper and lower 68% confidence intervals, i.e. 1σ , of the marginals. The third column shows the multivariate maximum a posteriori (MAP), and the uncertainties are defined as the upper and lower 68% confidence intervals, i.e. 1σ , given by the projected joint highest posterior density, PJ-HPD (Robert 2007; Joachimi et al. 2021).

sampled, we chose to train the NDEs on 18 000 realisations. This is facilitated by the computational efficiency of the forward simulations making additional realisations relatively inexpensive.

Upon training each of the six MAF NDEs using forward simulations independently, the NDEs are stacked while weighting each according to their relative validation losses. From the stacked NDE, we obtain the final posterior distributions by sampling with nautilus (Lange 2023) and check for consistency with emcee (Foreman-Mackey et al. 2013).

As advised in L23, to minimise the loss of information through from the compression step described in Sect. 5.3, once the simulation-based inference has been conducted the first time on the KiDS-1000 data for a given forward model, we repeat the inference again a second time while taking the fiducial model parameters assumed in the score compression to be given by the maximum a posteriori (MAP) from the initial inference (see Table 3). We note that in the second inference, we do not change the forward model in any way, and keep the exact same realisations of the model as in the first analysis. Only the fiducial data vector used to characterise the likelihood in Eq. (43) are updated

to re-compress all simulated data vectors. Considering that the initial fiducial parameters based on previous KiDS-1000 analyses are mostly consistent with the MAP which we infer in the initial analysis, while score compression is also robust against the biased fiducial parameters (Lin et al. 2023), we find that this step only marginally alters the inferred posterior.

6. Goodness-of-fit in SBI

As the inference is performed using SBI, we cannot perform a goodness of fit test the traditional way when we have access to an analytical likelihood. However, we still wish to keep this metric and so create a novel implementation of the χ^2 goodness-of-fit test. We follow the work of Gelman et al. (1996) and implement a SBI goodness-of-fit test from simulations in the ECP-GF¹⁷ module.

The methodology is as follows, initially, we find the maximum a posteriori (MAP) values from the posterior sampled in the SBI using the Nelder-Mead (Nelder & Mead 1965) optimiser. As the parameters of n_s and A_{bary} are unconstrained by our analysis, and the MAP can therefore vary freely across the prior, instead of making use of the MAP cosmology, we run simulations at the cosmology chosen to perform score compression for those parameters, which is already a pseudo maximum likelihood estimation (see Lin et al. 2023 for more details).

At the MAP cosmological parameter values, we simulate the sampling distribution by running a further 500 forward simulations while randomly varying noise realisations. We can thus define a χ^2 discrepancy measure as follows

$$\chi^2(\mathbf{t}|\Theta) := (\mathbf{t}_i - \mathbb{E}[\mathbf{t}_*|\Theta_*])^T (\text{Cov}(\mathbf{t}_*|\Theta_*))^{-1} (\mathbf{t}_i - \mathbb{E}[\mathbf{t}_*|\Theta_*]), \quad (45)$$

where \mathbf{t}_* is the compressed data vector at the fiducial cosmology, Θ_* the fiducial cosmology and \mathbf{t}_i one of n randomly varied noise realisations of the compressed data vector at the MAP parameter values. A fully Bayesian implementation would simply require forward simulations to be realised from posterior parameter draws as opposed to only at the inferred MAP parameter value. We have stuck to making use of simulations at the MAP to be analogous to the traditional KiDS-1000 analysis. Furthermore, as we are making use of linear score compression, we have linearised our model, and thus our goodness of fit is analogous to the classic χ^2 test. The main difference is that the degrees of freedom of the maximum χ^2 does not trivially reflect the degrees of freedom of our model due to the effects of optimal score compression in performing data dimensionality reduction.

For our setup, we evaluate a mean model data vector and a covariance from a set of 1600 simulations run at the fiducial cosmology. It should be noted that these simulations are the same as the ones required for the score compression step described in Sect. 5.3, so this comes without any additional computational costs. Using this newly obtained model data vector and covariance, and the χ^2 discrepancy measure defined above, we calculate the discrepancy measure for the 100 separate forward simulations from the anisotropic systematics model, and the KiDS-1000 gold sample data, which are depicted in the left panel of Fig. 13. As we obtain a probability to exceed (PTE) of 0.42, we find a good fit to the data. We note that in this test, the degrees of freedom are not recovered directly due to the effects of normalisation and optimal score compression in doing dimensionality reduction.




Fig. 13. Goodness-of-fit test implemented following Gelman et al. (1996). A distribution of χ^2 values for 500 forward simulations at the MAP cosmology is evaluated from a χ^2 function characterised by a Gaussian likelihood defined by a numerical covariance. The vertical dashed line indicates the χ^2 for the MAP when analysing the real KiDS-1000 data vector. The solid blue line shows the histogram upon the application of a smoothing kernel. Since the data's χ^2 falls near the mean and the mode of the χ^2 distribution from random realisations of the anisotropic systematics model, we conclude that the model evaluated at the MAP cosmology provides a good fit to the data, i.e. the data is a feasible realisation of the model.

7. Validation

The SBI pipeline presented in this work has previously been extensively tested in an idealised case. From the findings in L23, we conclude that the SBI pipeline based on DELFI implemented within KiDS-SBI robustly and accurately recovers the posterior distribution of Λ CDM cosmological parameters from KiDS-1000 cosmic shear data, even if the data compression loses information. However, L23 only considered simulated data whose likelihood is by construction Gaussian. For this reason, it becomes necessary to test the robustness of the SBI pipeline again for the anisotropic systematics forward model. We find in Appendix E that accurately and robustly recovers the underlying posterior distribution of Λ CDM parameters.

7.1. Coverage

To determine whether the learnt posterior distribution from the anisotropic systematics model is representative of the true underlying posterior and is unbiased, we conduct the Tests of Accuracy with Random Points (TARP) shown in Fig. 14. TARP (Lemos et al. 2023a) measures the expected coverage probability of random posterior samples within a given credibility level of the learnt posterior empirically. Whilst the original authors have their implementation of the coverage test, we made use of our implementation (ECP-GF¹⁷) to interface with the outputs from our simulations and DELFI in a more straightforward manner.

We did not perform the coverage test over our entire parameter space, as there are many cosmological parameters that are entirely unconstrained by our inference. Instead, the coverage test was only performed on S_8 and Ω_m as those are the parameters of greatest interest for cosmic shear.

To perform this coverage test, we first ran 500 new simulations at the fiducial cosmology for A_{IA} , n_s , h_0 , A_{bary} , freely varying the correlated $\delta_z[5]$ values, but varied $S_8 \in [0.5, 0.9]$,

¹⁷ https://github.com/Kiyam/ecp_gf_tests

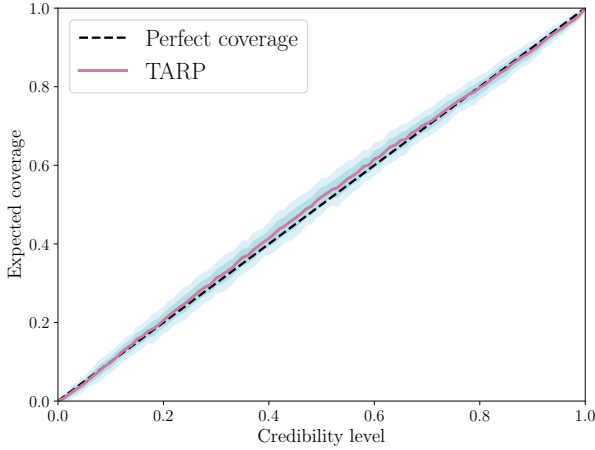


Fig. 14. The expected coverage probability versus the credibility level as defined in the Tests of Accuracy with Random Points (TARP) described in Lemos et al. (2023a) for the posterior shown in Fig. D.2 assuming the anisotropic systematics model. The dark blue region indicates the 1σ fluctuations in the TARP obtained from bootstrapping over 100 realisations, while the light blue region indicates the 2σ region. The dashed line is a reference line for a perfectly linear relation. The credibility level gives the fraction of the total probability density of the learnt posterior being considered, while the expected coverage probability measures the fraction of posterior samples which have a posterior probability smaller than the best estimate at a given credibility level. We note the relation, in this case, indicates slight underconfidence, such that the posteriors in Figs. D.2 and 18 can be considered as conservative. The ideal case is completely within the 1σ confidence interval of the TARP test.

$\omega_b \in [0.019, 0.026]$, $\omega_c \in [0.051, 0.255]$ and $h_0 \in [0.64, 0.82]$, with a latin hypercube. A slightly narrower range was chosen for S_8 than the prior for the full inference so that the testing points were positioned closer to the final parameter posterior range of interest. For each of these simulations, the noise realisation was allowed to freely vary. We also performed a coverage test with A_{IA} also varied, and obtained similar results to what is shown in Fig. 14.

We then used our trained DELFI model to perform inference on the new suite of 500 simulations, each acting as a mock data vector and thus obtaining 1000 posterior points for all 500 simulations. From the posterior samples, the coverage probability was calculated empirically for a given credibility level (see Lemos et al. 2023a for a detailed discussion of the algorithm). Due to the inherently stochastic nature of the coverage probability test, the test was bootstrapped 100 times for each credibility level. Figure 14 depicts both the mean and two standard deviations of the spread of the bootstrapped TARP coverage test.

For a well-estimated posterior that is not biased and representative of the true underlying posterior, we expect the coverage probability to be directly proportional to the credibility level. Any biases in the learnt posterior would lead to deviations from linearity in the TARP.

As can be seen from Fig. 14, the learnt posterior is very close to being ideal. As the mean coverage is seen to be slightly lower than the credibility level but always within one standard deviation of the ideal case. This means that the learnt posterior is slightly underconfident. This is acceptable as it implies that the KiDS-SBI posterior estimates are conservative but not biased. We repeat this test while also varying A_{IA} , and find similar results.

7.2. Impact of variable depth and shear bias

To assess the impact of the additional observational systematics considered in the anisotropic systematics model, we re-train a set of NDE ensembles on another 18 000 simulations of a different model which does not consider variable depth or additional shear biases while following the same procedure as outlined in Sect. 5. We refer to this model as the standard isotropic systematics model as it is more in line with previous ‘standard’ KiDS-1000 analyses such as Asgari et al. (2021), van den Busch et al. (2022), Li et al. (2023a), or Dark Energy Survey and Kilo-Degree Survey Collaboration (2023). This model considers the cosmological signal, intrinsic alignments, and baryonic feedback in the same way as the anisotropic systematics model, but it has two main differences: it does not consider variable depth, and it does not consider the additive or PSF shear biases. Firstly, when sampling galaxies as described in Eq. (20) the galaxy density, n_{gal} , and the photometric redshift distribution, $P(p|i)$, are assumed to be isotropic for a given tomographic bin. Secondly, when biasing the observed galaxy shapes according to the shear biases measured in KiDS-1000, the standard isotropic systematics model only considers the multiplicative shear bias; that is, only $M^{(p)} \neq 0$ while $\alpha^{(p)} = 0$, $c^{(p)} = 0 \forall p$.

This simplification of the forward model impacts the measured cosmic shear signal in a few ways. Firstly, from just comparing the measured two-point statistics, as can be seen in more detail in Appendix C, we find that variable depth can contribute up to 1% to the measured two-point shear signal in KiDS-1000, which is consistent with the predictions from Heydenreich et al. (2020). At the same time, variable depth can significantly alter the uncertainty of the two-point statistic (by $\lesssim 20\%$; Joachimi et al. 2021). As the intrinsic galaxy shape dispersion and the galaxy density vary anisotropically, the shape noise of the sample is no longer determined by the mean value per tomographic bin. Instead, the shape noise becomes an anisotropic distribution too which in the case of KiDS-1000 is skewed below the mean (see Appendix C). In addition, as can be seen by eye from Fig. 7 and 8, variable depth can add angular correlations to the data at the scales of pixels, pointing overlaps, whole pointings or even over the whole footprint, while also changing the signal along the line of sight anisotropically as shown in Fig. 10. These additional correlations can significantly alter the covariance of the two-point statistic or may even lead to non-Gaussian noise.

In fact, these non-cosmological correlations are significant enough for us to find that the overall uncertainty on the two-point statistic of a model including variable depth is larger than the uncertainty in a model which does not consider the effect. Hence, variable depth contributes approximately 1% ($\sim 0.2\sigma$) additional non-cosmological two-point signal.

The other distinguishing systematic between the anisotropic systematics model and the standard isotropic systematics model is the shear bias. When including additive and PSF shear biases in the anisotropic systematics model, which can also add a per-cent level (up to $\sim 0.3\sigma$, but less than 0.1σ at most angular scales) non-cosmological signal to the pseudo-CIs, mostly driven by the PSF variation shear bias (see Appendix C for details). This is consistent with the negligible impact on the measurement found in previous studies (Giblin et al. 2021). At the same time, the inclusion of the $\alpha_1^{(p)}$, $\alpha_2^{(p)}$, $c_1^{(p)}$ and $c_2^{(p)}$ parameters for each tomographic bin and their associated uncertainties adds additional variance to the anisotropic systematics model when we pre-marginalise over these 20 parameters.

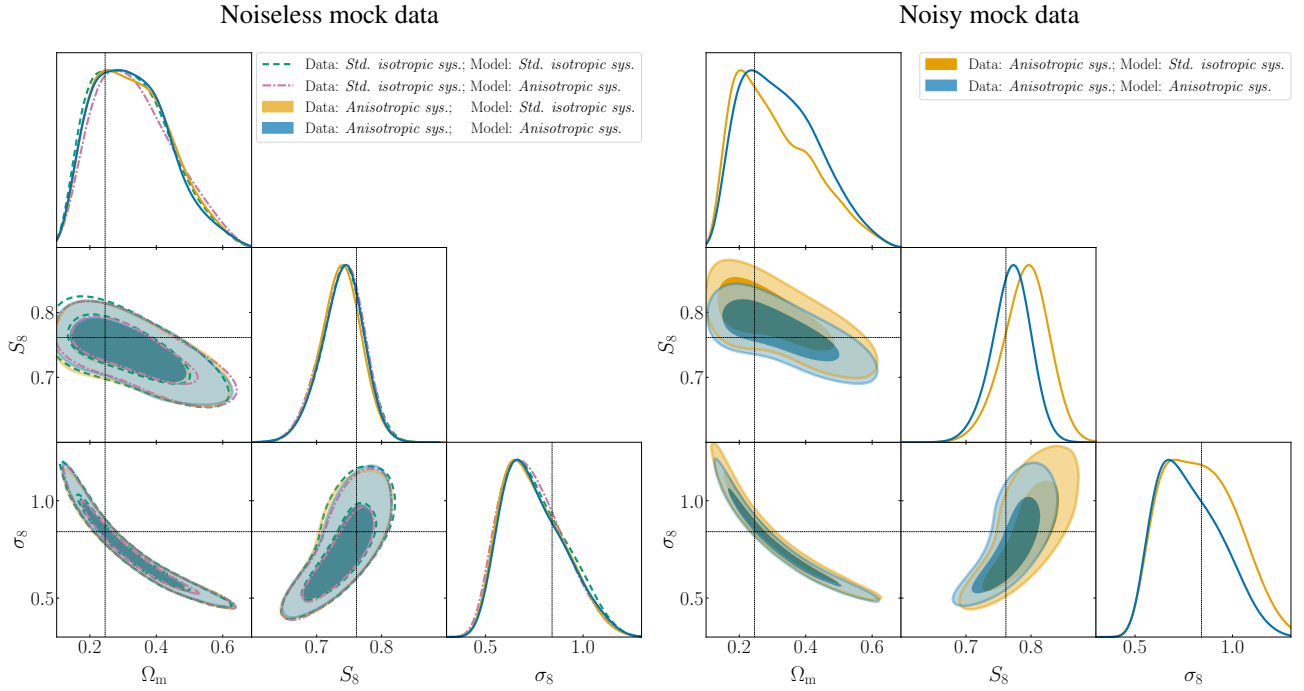


Fig. 15. Contour plots comparing the effects of model misspecification for different model cases and noise levels in the data. *Left panel:* Comparison of marginalised posterior contours from noise-free data assuming the anisotropic systematics model (which includes variable depth and PSF variations) in blue and green, and the standard isotropic systematics model (which is in line with the analysis assumptions in [Asgari et al. 2021](#), i.e. it does not model variable depth or PSF variations) in orange and pink. The orange and the blue posteriors are derived from the same data vector: a noise-free realisation from the anisotropic systematics model at the fiducial cosmology shown in Table 2; while the pink and the green posteriors are derived from a noise-free realisation from the standard isotropic systematics model at the same cosmology. *Right panel:* Comparison of marginalised posterior contours from noisy data simulated with anisotropic systematics model with the same seed as in the noise-free case, and subsequently analysed with the SBI based on the anisotropic systematics model in blue, and standard isotropic systematics model in orange. In both panels, the black solid line indicates the true cosmology assumed for the mock data vector (see Table E.1).

Therefore, the additive and PSF variation shear biases are yet another source of additional non-cosmological signal and variance.

These effects clearly propagate to the posteriors as can be seen when comparing the full simulation-based inference based on a likelihood learnt from 18 000 realisations of the standard isotropic systematics model to a likelihood based on 18 000 realisations of the anisotropic systematics model in Fig. 15.

To assess the effect of model misspecification and whether the inclusion or exclusion of variable depth and shear biases can bias the SBI of cosmological parameters, we repeat the inference of a mock data vector, but this time we generate a random realisation from the standard isotropic systematics model and analyse it using the likelihood learnt from the forward simulation which assumes the anisotropic systematics model, and vice versa. This allows us to determine whether the inclusion of the effects in the forward model would bias our inference if they were not actually present in the data that is analysed. We also show these posteriors in Fig. 15 where we see that when the data model and the likelihood model are mismatched, the constraints are consistent within a given model. This confirms that, given the learned likelihood, the bias in the two-point shear signal from both variable depth and PSF shear bias is negligible in KiDS-1000. These findings are consistent with previous work which found this to be the case for KiDS-1000 for variable depth ([Heydenreich et al. 2020](#); [Joachimi et al. 2021](#)) and for the PSF shear bias ([Giblin et al. 2021](#)).

Nonetheless, systematics such as variable depth and the PSF shear bias do not only affect the measured signal, but also

its uncertainty. As shown in the right panel of Fig. 15, when analysing the same noisy mock data vector simulated with the anisotropic systematics model with the SBI trained on both models (anisotropic systematics and standard isotropic systematics), the marginal in S_8 shifts considerably by about 0.7σ . At the same time, the constraints based on the anisotropic systematics model give $\sim 5\%$ smaller 1σ confidence intervals in the S_8 marginals than standard isotropic systematics model. As each model characterises a different effective likelihood, when analysing the same data vector with both models, each data point is weighted differently depending on the model assumed during the inference. Depending on the noise realisation of the mock data and the effective likelihood of a given model, the posterior may vary considerably. This shows that when assessing the impact of a systematic a measurement, it is not sufficient only checking its contribution to the signal with respect to the noise. Simultaneously, it is important to assess the impact of the systematic on the likelihood itself.

Of course, the noise realisation depicted in the right panel of Fig. 15 is only one possible realisation of many. We find that for other noise realisations from mocks, the difference in the cosmological constraints between anisotropic systematics and standard isotropic systematics models varies as each posterior scatters as shown in Fig. H.1. Consequently, neglecting variable depth and the PSF shear bias in the forward modelling, while it is present in the data that is analysed, can bias the inferred value of S_8 by up to $\sim 1\sigma$. As discussed in Sect. 9, this is consistent with what we find for the KiDS-1000 data.

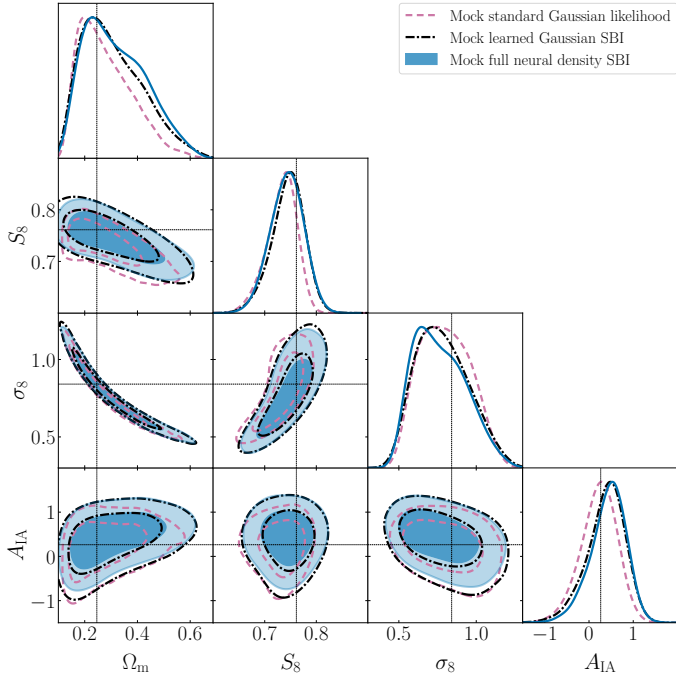


Fig. 16. Comparison of likelihood models with a noise-free mock data vector from the anisotropic systematics forward model. Shown are marginalised posterior contours: in blue, assuming the anisotropic systematics model within the standard setup of KiDS-SBI as described in Sect. 5.4; in black, assuming the same anisotropic systematics forward simulations to train a Mixture Density Network (MDN; Bishop 1994) which is made up of a single multivariate Gaussian; and in pink, applying a standard analysis which assumes a Gaussian likelihood characterised by a numerical sample covariance derived from anisotropic systematics forward simulations and a model given by the analytic pseudo-CI model described in Appendix D. All posteriors are evaluated assuming the same noise realisation and the same model parameters (see Table E.1).

8. Impact of the Gaussian likelihood assumption

One of the main aims of this work is to determine whether the standard assumption of the likelihood of cosmic shear two-point statistics being Gaussian holds under realistic conditions while including systematics. To test this assumption, we performed the same analysis of a mock data vector using SBI (as described in Sect. 5 and Appendix E) again, but with different inference methods.

As shown in Fig. 16, we compare the posterior from the SBI analysis learnt from forward simulations based on the anisotropic systematics model to a standard Gaussian likelihood analysis and an SBI analysis which forces the learnt likelihood to be Gaussian. The standard analysis assumes a Gaussian likelihood characterised by the same numerical covariance from 1600 realisations of the anisotropic systematics model fixed at the fiducial cosmology in Table 2. The model evaluations are calculated from analytical theory for the cosmic shear signal as described in Appendix D, and the sampling is done with the nautilus sampler (Lange 2023) over the same 12D prior space defined in Table 2. In contrast, the Gaussian SBI analysis makes the exact same modelling choices as the main SBI analysis described Sect. 5 and Appendix E with one change: instead of using ensembles of MAFs made up of MADEs as the architecture of the neural density estimators when learning the effective likelihood, we use an ensemble of three single-component

Mixture Density Networks (MDNs; Bishop 1994) to learn the likelihood from the 18 000 realisation of the anisotropic systematics model. As a consequence, the ensemble of MDNs forces the learnt likelihood to be described by a single multi-variate Gaussian distribution. Notably, this is different from a standard Gaussian likelihood analysis with a Gaussian likelihood in that the Gaussian SBI analysis allows the covariance of the likelihood to vary with cosmology.

From Fig. 16, we can conclude that the full non-Gaussian SBI is in exceptionally good agreement with the SBI analysis which forces the likelihood to be Gaussian but allows for cosmology dependence in the uncertainty (particularly, when the mock data vector includes realistic noise). Thus, the uncertainty as modelled by the anisotropic systematics model can be considered to be approximately Gaussian at a given cosmology, since a Gaussian NDE network is fully capable of learning an unbiased likelihood. We conclude that, for a fixed cosmology, the assumption of a Gaussian likelihood for the pseudo-CI cosmic shear signal for $\ell > 76$ holds well. This is in agreement with previous studies that found this to be the case in angular power spectra for $\ell > 50$ (Hamimeche & Lewis 2008; Schneider & Hartlap 2009; Sellentin et al. 2018; Lin et al. 2020).

At the same time, we find that the standard Gaussian likelihood analysis has a noticeably different posterior to both SBI analyses despite the underlying data vector being identical, and the covariance being sampled from the same anisotropic systematics model. In particular, the S_8 marginal is appreciably narrower in the Gaussian likelihood analysis than in the SBI analyses. In the noiseless case shown in Fig. 16, we find that the SBI gives $S_8 = 0.743^{+0.034}_{-0.031}$, while the standard Gaussian likelihood analysis gives $S_8 = 0.733^{+0.034}_{-0.024}$, specifically the 1σ confidence interval in S_8 is $\sim 11\%$ narrower than in the SBI analysis. When adding realistic noise to the data vector from the same seed, we find that the SBI recovers $S_8 = 0.762^{+0.032}_{-0.027}$, while the Gaussian likelihood analysis recovers $S_8 = 0.744^{+0.031}_{-0.022}$, specifically the 1σ confidence interval in S_8 is also $\sim 10\%$ narrower than in the SBI analysis. Additionally, when noise is introduced, the Gaussian likelihood analysis appears more biased with respect to the truth ($S_8 = 0.756$) than the SBI.

The good agreement between the Gaussian likelihood SBI based on MDNs and the non-Gaussian likelihood SBI based on MAFs indicates that both analyses appear to be sensitive to some form of cosmology dependence in the uncertainty model, as that is the only effect which is not accounted for in the standard Gaussian likelihood analysis.

To investigate the origin of this effect further, we sample the learnt likelihood from the full SBI analysis directly at different values of S_8 in Fig. 17. In this plot, we show the 7D score-compressed data space and label each element of the data vector according to the model parameter with which they are most correlated. Then, we vary the value of S_8 while keeping all other parameters fixed at the fiducial cosmology, and plot the likelihood of the data given the chosen model parameters. We find that as S_8 increases the marginal likelihoods of the elements of the score-compressed data vector increase linearly. In particular, the data elements most correlated with ω_c , ω_b , n_s , and h_0 show the strongest linear correlation with S_8 . The implication of this is that the anisotropic systematics model entails that data at higher assumed S_8 values has larger uncertainties than data modelled for low S_8 values, and therefore the learnt likelihood is different from a Gaussian likelihood with the covariance evaluated at a fixed cosmology. This effect appears to drive the increase in the uncertainty on S_8 of the SBI-based posteriors with respect to the Gaussian likelihood analysis in Fig. 16.

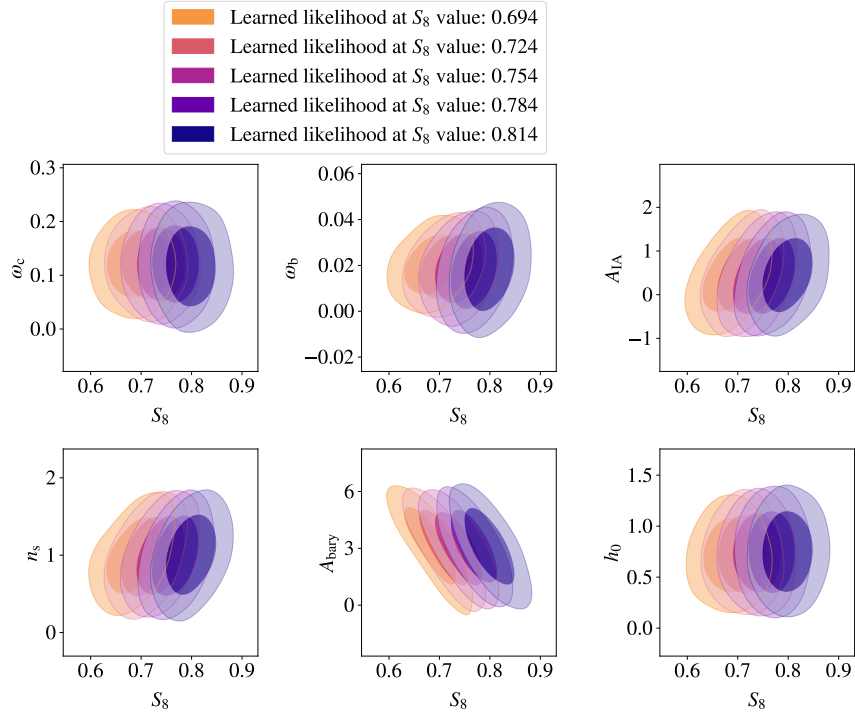


Fig. 17. Likelihood marginals in the compressed data space for five different sets of cosmological parameters given the anisotropic systematics model within KiDS-SBI over the prior space shown in Table 2. The compressed data values are labelled according to the cosmological parameter with which they are most correlated (see Sect. 5.3 for details). For the orange contours, the input data vector is set to $S_8 = 0.694$, $S_8 = 0.724$ for the pink contours, $S_8 = 0.754$ for the purple contours, $S_8 = 0.784$ for the purple contours, and $S_8 = 0.814$ for the blue contours. All other cosmological parameters are taken to be the same as in Table E.1.

In Appendix F, we conduct a detailed investigation into the physical origin of this observed cosmology dependence in the learnt likelihood of cosmic shear two-point statistics. We find that, as has been noted in previous work (Eifler et al. 2009; Reischke et al. 2017), the cosmic variance of the cosmic shear two-point signal (at least, for $76 \leq \ell \leq 1500$) is measurably cosmology-dependent, and mostly S_8 -dependent. We also determine that this effect is expected to be important enough to become detectable in most tomographic bin combinations of the KiDS-1000 gold sample in the regime of $\ell \sim 10^2$. Since the cosmic variance is imprinted into the forward simulations at the step of sampling log-normal random matter fields (see Sect. 3.4), the cosmology dependence of the cosmic variance is present in any forward model presented in this work irrespective of the observational systematic effects considered (namely both the anisotropic systematic and the standard isotropic systematics model are affected by this).

With this being the case, the noise of the measured cosmic shear pseudo-CIs may contain some information which may impact the inferred cosmology from the SBI. Hence, it is important to validate whether the variance realised by log-normal random simulations does not bias the cosmological inference. Firstly, it has been found in previous studies that log-normal simulations can incorporate accurate higher-order statistics into the sampled matter, galaxy and shear fields when compared to N-body dark matter simulations (Hall & Taylor 2022; Piras et al. 2023), so log-normal fields can incorporate three- or four-point statistics to model the variance of two-point statistics. Additionally, we find in our own testing that a sample covariance of cosmic shear two-point statistics in KiDS-1000 as generated from log-normal simulations agrees within 10% with analytical covariance matrices computed from analytical considerations

(Joachimi et al. 2021; Reischke et al. 2024), so the variance of the cosmic shear signal as modelled within KiDS-SBI is consistent with analytical theory.

9. Cosmological inference from KiDS-1000 data

9.1. Blinding

Throughout this analysis, we followed a procedure to avoid any form of unconscious bias when making choices within the KiDS-SBI models. Although the KiDS-1000 two-point shear statistics have been measured before, we fixed the forward model used in the analysis before ever evaluating learnt likelihood function for the KiDS data.

All authors did not conduct any measurements on the KiDS-1000 data nor was it analysed with the SBI pipeline until the forward model as well as all validation tests were finalised and fixed. Finally, upon evaluating the learnt likelihood for the KiDS data, we initially inferred an anomalously low S_8 with the SBI and the standard MCMC caused by an inconsistency between the measurement code for the forward simulations and the real data. Upon re-measuring the KiDS data in a consistent fashion, we obtain the pseudo-CIs from the KiDS-1000 gold sample as shown in Fig. 1.1 and ran it through the entire KiDS-SBI pipeline without making any adjustments to the methodology described in this paper.

9.2. Cosmological parameter constraints

With the learnt likelihood found to be unbiased, and with an understanding of the sensitivity of the likelihood to systematics and cosmology, we proceed to analyse the cosmic

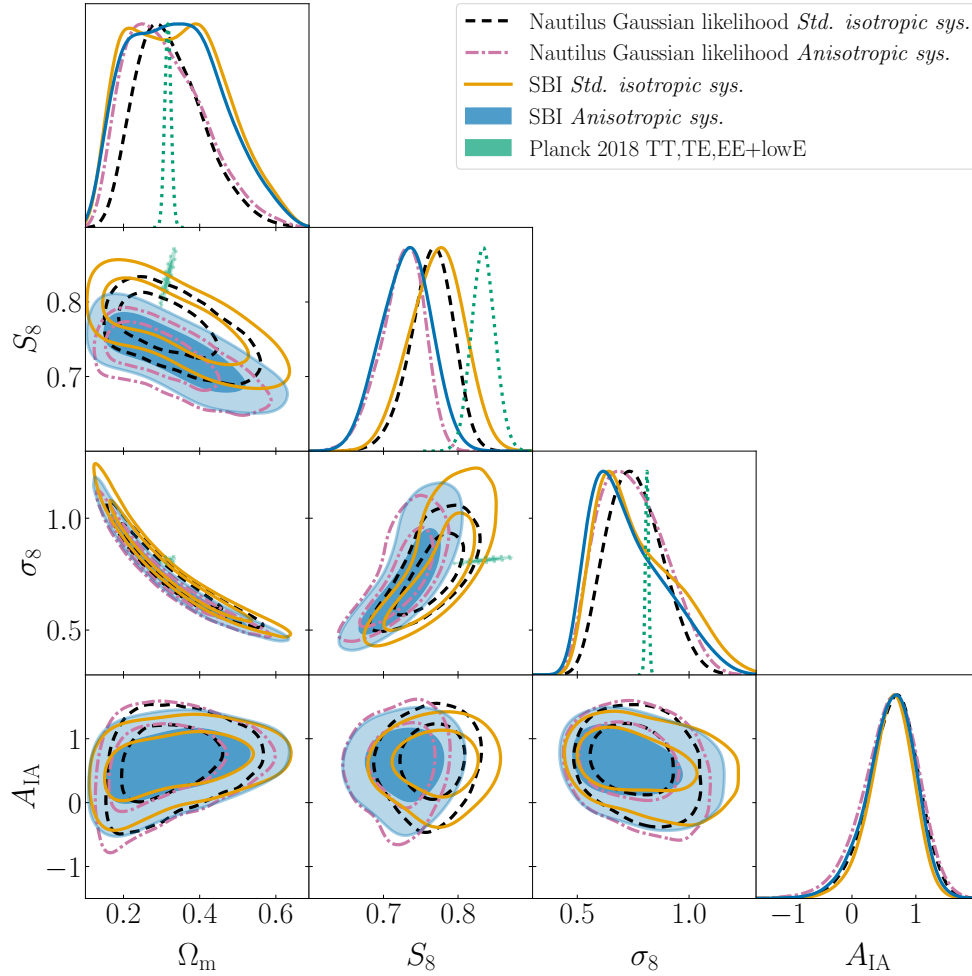


Fig. 18. Posterior contours of the main constrained cosmological parameters from the KiDS-SBI analysis of the KiDS-1000 cosmic shear data assuming the anisotropic systematics model (in blue), which incorporates additional systematics such as variable depth and shear biases, compared against posterior contours from other analyses. In pink, we show the contours for the equivalent analysis assuming a Gaussian likelihood sampled with nautilus (Lange 2023). In orange, we show the posterior from the same data while assuming the standard isotropic systematics model, which considers the systematic effects which are typically modelled in standard cosmic shear analyses. In black, we show the posterior sampled with nautilus based on a Gaussian likelihood defined by the standard isotropic systematics model. The green contour shows the posterior from the cosmic microwave background constraints from the TT,TE,EE+lowE modes (Planck Collaboration VI 2020). We note that the Planck contours do not have any marginals in A_{IA} as the CMB is not sensitive to the intrinsic alignments of galaxies.

shear pseudo-CIs measured from the KiDS-1000 gold sample described in Sect. 2 (see Fig. 1.1 for the full measurement).

Our main results are based on the likelihood learnt from the anisotropic systematics model in KiDS-SBI which incorporates variable depth and additional shear biases when compared to previous KiDS analyses (for example Asgari et al. 2021). The marginalised posterior and the constraints on the main parameters are shown in Fig. 18 and Table 3, respectively (see Appendix I for the full posteriors and all parameter best estimates). We report the maximum a posteriori (MAP) \pm the projected joint highest posterior density (PJ-HPD; Robert 2007; Joachimi et al. 2021) of $0.743^{+0.015}_{-0.051}$, and a mean marginal with 68% confidence intervals of 0.731 ± 0.033 . We note that Fig. 18 also shows that the SBI significantly broadens the posterior in Ω_m with respect to a Gaussian likelihood analysis, while the mean value of Ω_m from SBI is highly consistent with the value measured from the early Universe with Planck 2018 TT,TE,EE+lowE (Planck Collaboration VI 2020).

As shown in Fig. 13, we find that the MAP best-fit cosmology given the anisotropic systematics model provides a good fit to the data with a probability-to-exceed (PTE) of 0.42. Addition-

ally, we plot the line of best fit at the MAP cosmology in Fig. 1.1 and find that it is in good agreement with the measured data vector and the analytical theory at the MAP cosmology.

In addition to our main results based on the anisotropic systematics model, we re-analyse the KiDS-1000 gold sample data with the SBI pipeline based on the standard isotropic systematics model which makes similar assumptions to previous KiDS-1000 analyses and does not consider the effects of variable depth and some shear biases (see Sect. 7.2 for details). When considering this model instead, we find a MAP \pm PJ-HPD of $S_8 = 0.780^{+0.020}_{-0.048}$, and a mean marginal with 68% confidence intervals of $0.772^{+0.038}_{-0.032}$, while the fit to the data at the MAP is similarly good with a χ^2 that gives a PTE of 0.52. With the standard isotropic systematics model, the marginal in S_8 is shifted by 0.9σ upwards. This is broadly consistent with the expectation from the validation discussed in Sect. 7.2, specifically that the extra systematics in the anisotropic systematics model can bias the inference when the data is noisy despite the bias in the signal itself being negligible at the precision of KiDS-1000. This is driven by the uncertainty modelling being considerably

different in both models (see Sect. 7.2): the shape noise model is different, variable depth changes the correlation functions due to the anisotropic selection, and the inclusion of additional shear biases adds statistical noise from 20 more parameters. Consequently, both models respond differently to the noise measured in the KiDS-1000 data.

We also conduct a Bayesian model comparison between the two models by evaluating the Bayesian evidence, Z , for each posterior using *nautilus* (Lange 2023). We find that for the anisotropic systematics model, $\ln(Z) = -14.18$, while for the standard isotropic systematics model, $\ln(Z) = -15.07$. Hence, assuming equal prior probability of both models, the log of the Bayes factor, $|\ln(B)| = 1.11$ which implies no significant preference for either model when compared to the other.

As also shown in Fig. 18, we find that when re-analysing the KiDS-1000 data assuming a Gaussian likelihood instead, we find results consistent with the testing conducted on mocks (see Fig. 16): the Gaussian likelihood assumption causes S_8 to be overconstrained by $\sim 10\%$ with respect to the SBI irrespective of how the observational systematics are modelled as it neglects to consider the cosmology dependence of cosmic variance. Additionally, the posterior mean marginal in S_8 assuming the Gaussian likelihood shifts from $0.764^{+0.031}_{-0.025}$ when deriving the sample covariance from the standard isotropic systematics model to a value of $0.725^{+0.034}_{-0.023}$ when assuming the anisotropic systematics model in the covariance instead. We note that both assume the same modelling in the signal described in Appendix D.

We highlight that the Gaussian likelihood analysis based on the standard isotropic systematics model agrees well with the constraints from Asgari et al. (2021) where $S_8 = 0.758^{+0.017}_{-0.026}$. We note that the constraints in our work are broader despite making consistent modelling assumptions. This is caused by the use of a different MCMC sampler. Here we have employed the *nautilus* sampler (Lange 2023), while in Asgari et al. (2021) the *multinest* sampler (Feroz et al. 2009) was used. It has been found that *multinest* can be overconfident by approximately $\sim 10\%$ (Lemos et al. 2023b) which we also find in our analysis and show in Appendix G. When accounting for the difference in sampler used, we find that the Gaussian likelihood constraints from the standard isotropic systematics model agree exceptionally well with the results from Asgari et al. (2021).

Likewise the inferred S_8 value when assuming a Gaussian likelihood derived from the anisotropic systematics model is consistent with the result from Loureiro et al. (2022) where the mean marginal is measured to be $0.742^{+0.034}_{-0.023}$. Despite both analyses being based on pseudo-CIs, the one presented in this work finds a slightly lower value of S_8 ($\sim 0.27\sigma$) which can be driven by the additional signal and uncertainty induced by the PSF variations into the anisotropic systematics model (see Appendix C) which was not considered in the modelling of Loureiro et al. (2022).

When comparing the constraints from the anisotropic systematics model SBI to previous cosmic shear analyses in Fig. 19, we find that our results are in agreement within 1σ of S_8 best-fits from all previous KiDS-1000 analyses (Asgari et al. 2021; Loureiro et al. 2022; van den Busch et al. 2022; Li et al. 2023a), the HSC-Y3 pseudo-CI analysis (Dalal et al. 2023), and the DES-Y3 two-point analysis (Amon et al. 2022; Secco et al. 2022). The one exception is the anisotropic systematics DES-Y3 analysis (Dark Energy Survey and Kilo-Degree Survey Collaboration 2023) whose marginal from the main hybrid analysis is different by 1.65σ from our marginal estimate of S_8 based on the anisotropic systematics model. However, in the same work, all re-analyses of the KiDS-1000 data alone (with




Fig. 19. Comparison between different constraints of the growth of structure parameter, S_8 , from cosmic shear measurements and from the cosmic microwave background. The circular points indicate MAP+PJ-HPD constraints and the triangular points reflect estimates from the marginal posterior mean and 68% confidence intervals. The blue dashed regions mark the 1σ region from the KiDS-SBI constraint based on the anisotropic systematics model shown in this work.

KiDS-like, DES-like, and hybrid analysis) are consistent within less than 1σ with the best-fit for S_8 presented in this work. These discrepancies may be partially attributed to cosmic variance, different modelling choices for certain systematics, and the impact of the cosmology dependence in the learnt likelihood which is included in our SBI analysis, but it is not considered in any other previous analyses.

In fact, as discussed in detail in Sect. 8, since the uncertainty is cosmology-dependent and scales with S_8 in KiDS-SBI, the SBI infers a larger uncertainty in S_8 on top of the added uncertainty from variable depth and shear biases. Additionally, variable depth and the PSF shear bias can impact the likelihood to significantly re-weight the data vector despite their impact on the signal being negligible. This may explain why the best estimate for S_8 from the anisotropic systematics model is about 0.5σ lower than previous standard analyses of KiDS-1000.

When comparing the results from the SBI with the anisotropic systematics model to early-Universe probes of large-scale structure, such as the results from the Planck collaboration shown in Fig. 18, we find that the main KiDS-SBI S_8 marginal is in 2.9σ tension with the best estimate of $S_8 = 0.834 \pm 0.016$ from Planck 2018 TT,TE,EE+lowE (Planck Collaboration VI 2020).

10. Conclusions

We performed a novel simulation-based inference (SBI) analysis, *KiDS-SBI*, of the cosmic shear two-point statistics

as measured by the Kilo-Degree Survey’s fourth data release (KiDS-1000; Kuijken et al. 2019).

We present a new suite of stochastic forward simulations of cosmic shear observables within KiDS-SBI, which take into consideration all the systematic effects which are typically modelled in a weak lensing analysis plus additional systematics that have the potential to significantly bias the signal. At the same time, KiDS-SBI introduces considerable improvements in accuracy and performance with respect to similar simulation environments in previous works (Joachimi et al. 2021): non-Limber projections of angular power spectra (with the new Levin module), higher resolution along the line of sight, inclusion of field-level intrinsic alignments, consideration of variable depth and anisotropic shear biases, and significant computational performance improvements.

On the basis of these forward simulations, we implemented the first simulation-based inference analysis of large-scale structure observables as measured by KiDS-1000 with the same level of complexity as other standard analyses in the field. This has enabled us to fully propagate all uncertainties from the data vector all the way to the inferred cosmological parameters in a Bayesian way. The analysis is based on neural density estimators which learn the effective likelihood through the mapping between the model parameters and the forward-simulated data vectors. We conducted an inference of the cosmological parameters in flat Λ CDM from cosmic shear two-point statistics, specifically, pseudo-CIs. We assumed a forward model consistent with previous KiDS-1000 cosmic shear analyses (Asgari et al. 2021; Loureiro et al. 2022), which additionally includes variable depth and anisotropic shear biases: the anisotropic systematics model. Variable depth accounts for the anisotropy in the galaxy selection function in KiDS-1000 and how it impacts the measured galaxy densities, shapes, and redshifts. The inclusion of anisotropic shear biases models the angular variation of the residual bias in the galaxy shapes due to variations in the point-spread function of the telescope from the mean model used in the shape calibration.

To assess the quality of any best fit obtained from KiDS-SBI, we also introduced a Bayesian goodness-of-fit measurement based on Gelman et al. (1996). From a small number of realisations of the forward model, we characterised a distribution of χ^2 values for the given model. By determining where the χ^2 of the best fit on the KiDS-1000 data lies, we evaluated the quality of the fit and whether the data stand as a feasible realisation of the forward model. On this basis, we were able to conclude that the KiDS-1000 data is fitted well by the anisotropic systematics model at the maximum a posteriori cosmology (with a probability-to-exceed of 0.42).

Through the testing conducted in this work and in a previous paper (Lin et al. 2023), we determine that KiDS-SBI recovers accurate and robust posteriors of the 12D space of cosmological and astrophysical parameters. At the same time, we find that the likelihood learnt from KiDS-SBI forward simulations is well approximated by a Gaussian distribution at a fixed cosmology. However, we determined that the uncertainty on the cosmic shear signal measurably varies with the assumed cosmological parameters, which are driven by the cosmic variance scaling with the underlying value of S_8 , as noted in previous works (Eifler et al. 2009). Consequently, we find that the uncertainty of the two-point statistics in KiDS-1000 is measurably cosmology-dependent which increases the nominal uncertainty of S_8 by $\sim 10\%$. Hence, the standard assumption of the covariance of a Gaussian likelihood being fixed for all parameter points in the prior space does not necessarily hold. Therefore, such

an assumption may not be valid for current Stage-III galaxy surveys and this issue may be also be relevant for upcoming stage-IV galaxy surveys, such as *Euclid* (Laureijs et al. 2011) or Rubin (*LSST Science Collaboration* 2009). Addressing this assumption further will improve the precision of cosmic shear measurements.

Given the anisotropic systematics model in KiDS-SBI, we obtain a mean marginal with 68% confidence intervals of 0.731 ± 0.033 . These constraints are in agreement within 1σ with previous cosmic shear analyses of KiDS-1000 (Asgari et al. 2021; Loureiro et al. 2022; van den Busch et al. 2022; Li et al. 2023a), HSC-Y3 (Li et al. 2023b; Dalal et al. 2023), and DES-Y3 (Amon et al. 2022; Secco et al. 2022). The tension in S_8 from the KiDS-1000 analysis, in this work and measurements from the cosmic microwave background from *Planck* 2018 TT,TE,EE+lowE (Planck Collaboration VI 2020), is at 2.9σ , which is similar to the 3.0σ tension found in an equivalent Gaussian likelihood analysis of KiDS-1000 (Asgari et al. 2021). Alternatively, this is also close to the 2.8σ tension found the analysis of KiDS-1000, which incorporated variable depth in the Gaussian likelihood (Loureiro et al. 2022), and the 3.1σ tension found in the KiDS-1000 beyond Λ CDM (Tröster et al. 2021). This establishes that the σ_8 -tension between late-Universe and early-Universe probes of large-scale structure still persists even when we drops the assumption of a Gaussian likelihood for the cosmic shear signal and considers additional contaminant systematics.

To evaluate and isolate the impact of the inclusion of variable depth and shear biases in the anisotropic systematics model, we reproduce the KiDS-SBI analysis with a second model which makes the same assumptions, but does not consider these specific systematic effects: the standard isotropic systematics model (which is more in line with previous KiDS-1000 analyses, such as Asgari et al. 2021). We find that the inclusion of both systematics in the forward model increases the S_8 marginal by 0.9σ or 5% and decreases the marginal 1σ confidence intervals by 6%. We find that this is not driven by these systematics’ impact on the signal but, rather, on their impact on the uncertainty model. The variable depth alone appears to impact S_8 by approx. 0.38σ when comparing the Gaussian likelihood results from Loureiro et al. (2022) with the constraints from the standard isotropic systematics model. The variations in the PSF modelling residuals appear to account for approx. 0.36σ in the shift in S_8 when comparing the Gaussian likelihood results from the anisotropic systematics model with the constraints from Loureiro et al. (2022). However, we note that the PSF variations and the variable depth are known to be correlated as both are influenced by some of the same observational systematics, such as atmospheric seeing or thermal effects in the telescope optics; thus, it may not be physically sound to consider these effects separately.

In any case, this reflects the fact that observational anisotropies, such as variable depth and PSF shear variation bias, can significantly alter the uncertainty modelling despite having a negligible impact on the shear signal at the precision levels of KiDS-1000. The KiDS-1000 data appears to be consistent with the presence of these effects since the anisotropic systematics model provides a good fit to the data, but when examining the Bayesian evidence, we do not find a clear preference between the two models by the data. Nevertheless, both systematic effects have been independently measured in the KiDS-1000 weak lensing data (Joachimi et al. 2021; Giblin et al. 2021).

This work stresses the importance of testing the common assumption of Gaussian likelihoods in realistic conditions. At the

same time, we also highlight the potential of simulation-based inference to rigorously perform Bayesian uncertainty propagation even when the models of the uncertainty of the signal or the systematics are analytically intractable. The approach using SBI with likelihood-based neural density estimation comes with the added advantage that it does not require the expensive re-evaluation of the posterior distribution if the input measurements or priors are updated, as would be the case for a standard MCMC analysis. Since the likelihood for a given forward model is learnt independently of the input data and the priors, SBI can significantly improve the efficiency of inference analysis where the forward model is not changed between measurements. All these advantages showcase the potential of simulation-based inference becoming the method of choice in future parameter inference and model testing analyses in observational cosmology.

We conclude that SBI is a powerful tool that may be able to tackle many of the physical and statistical challenges faced by future galaxy surveys, as their observations become more precise, less limited by statistical noise, and more limited by the accurate modelling of systematics. To address this need, future works should expand upon the forward model in KiDS-SBI to also include other probes of large-scale structure such as galaxy clustering and galaxy-galaxy lensing, while incorporating the relevant modelling for the relevant systematic effects, such as magnification bias, non-linear galaxy bias, and field-level baryonic feedback.

Acknowledgements. We thank Stephen Feeney, Justin Alsing, Hendrik Hildebrandt, Marika Asgari and Alan Heavens for helpful discussions. We are grateful to the referee for the constructive comments on the manuscript. MvWK thanks the Science and Technology Facilities Council for support. KL thanks the UCL Cosmoparticle Initiative. NT and BJ are supported by the ERC-selected UKRI Frontier Research Grant EP/Y03015X/1 and the UK Space Agency grants ST/W002574/1 and ST/X00208X/1. BJ is also supported by STFC Consolidated Grant ST/V000780/1. AL acknowledges support from the research project grant ‘Understanding the Dynamic Universe’ funded by the Knut and Alice Wallenberg Foundation under Dnr KAW 2018.0067. RR is supported by the European Research Council (Grant No. 770935). AHW is supported by the Deutsches Zentrum für Luft- und Raumfahrt (DLR), made possible by the Bundesministerium für Wirtschaft und Klimaschutz, and acknowledges funding from the German Science Foundation DFG, via the Collaborative Research Center SFB1491 ‘Cosmic Interacting Matters – From Source to Signal’. This work used facilities provided by the UCL Cosmoparticle Initiative. We thank Dr Edd Edmonson for technical support at the UCL HPC facilities. The results in this paper are based on observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs: 088.D-4013, 092.A-0176, 092.D-0370, 094.D-0417, 177.A-3016, 177.A-3017, 177.A-3018 and 179.A-2004, and on data products produced by the KiDS consortium. The KiDS production team acknowledges support from: Deutsche Forschungsgemeinschaft, ERC, NOVA and NWO-M grants; Target; the University of Padova, and the University Federico II (Naples). Contributions to the data processing for VIKING were made by the VISTA Data Flow System at CASU, Cambridge and WFAU, Edinburgh. The work presented in this paper made use of the following software packages: Numpy (Harris et al. 2020), SciPy (Virtanen et al. 2020), AstroPy (Astropy Collaboration 2018), HealPy (Zonca et al. 2019), GetDist (Lewis 2019), Jupyter Lab (Kluyver et al. 2016), numba (Lam et al. 2015), and Matplotlib (Hunter 2007).

References

Abbott, T. M. C., Aguena, M., Alarcon, A., et al. 2022, *Phys. Rev. D*, **105**, 023520

Abramo, L. R., Secco, L. F., & Loureiro, A. 2016, *MNRAS*, **455**, 3871

Abramo, L. R., Dinarte Ferri, J. V., Tashiro, I. L., & Loureiro, A. 2022, *JCAP*, **2022**, 073

Aihara, H., AlSayyad, Y., Ando, M., et al. 2022, *PASJ*, **74**, 247

Alsing, J., & Wandelt, B. 2018, *MNRAS*, **476**, L60

Alsing, J., Wandelt, B., & Feeney, S. 2018, *MNRAS*, **477**, 2874

Alsing, J., Charnock, T., Feeney, S., & Wandelt, B. 2019, *MNRAS*, **488**, 4440

Amon, A., & Efstathiou, G. 2022, *MNRAS*, **516**, 5355

Amon, A., Gruen, D., Troxel, M. A., et al. 2022, *Phys. Rev. D*, **105**, 023514

Asgari, M., Heymans, C., Hildebrandt, H., et al. 2019, *A&A*, **624**, A134

Asgari, M., Lin, C.-A., Joachimi, B., et al. 2021, *A&A*, **645**, A104

Astropy Collaboration (Price-Whelan, A. M., et al.) 2018, *AJ*, **156**, 123

Baleato Lizancos, A., & White, M. 2023, *JCAP*, **2023**, 044

Bartelmann, M., & Schneider, P. 2001, *Phys. Rep.*, **340**, 291

Beaumont, M. A. 2019, *Ann. Rev. Stat. Appl.*, **6**, 379

Becker, M. R., Troxel, M. A., MacCrann, N., et al. 2016, *Phys. Rev. D*, **94**, 022002

Benítez, N. 2000, *ApJ*, **536**, 571

Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003, *ApJS*, **148**, 1

Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393

Bishop, C. M. 1994, *Mixture density networks* (Aston University)

Blazek, J., McQuinn, M., & Seljak, U. 2011, *JCAP*, **2011**, 010

Blazek, J., Vlah, Z., & Seljak, U. 2015, *JCAP*, **2015**, 015

Bocquet, S., Grandis, S., Bleem, L. E., et al. 2024, *Phys. Rev. D*, **110**, 083510

Böhm, V., Hilbert, S., Greiner, M., & Enßlin, T. A. 2017, *Phys. Rev. D*, **96**, 123510

Bridle, S., & King, L. 2007, *New J. Phys.*, **9**, 444

Brown, M. L., Taylor, A. N., Hambly, N. C., & Dye, S. 2002, *MNRAS*, **333**, 501

Brown, M. L., Castro, P. G., & Taylor, A. N. 2005, *MNRAS*, **360**, 1262

Castro, P. G., Heavens, A. F., & Kitching, T. D. 2005, *Phys. Rev. D*, **72**, 023516

Catelan, P., Kamionkowski, M., & Blandford, R. D. 2001, *MNRAS*, **320**, L7

Chen, A., Harness, A., & Melchior, P. 2023, *J. Astron. Telescopes Instrum. Syst.*, **9**, 025002

Chiu, I. N., Klein, M., Mohr, J., & Bocquet, S. 2023, *MNRAS*, **522**, 1601

Clerkin, L., Kirk, D., Manera, M., et al. 2017, *MNRAS*, **466**, 1444

Coles, P., & Jones, B. 1991, *MNRAS*, **248**, 1

Dalal, R., Li, X., Nicola, A., et al. 2023, *Phys. Rev. D*, **108**, 123519

Dark Energy Survey and Kilo-Degree Survey Collaboration (Abbott, T. M. C., et al.) 2023, *Open J. Astrophys.*, **6**, 36

de Jong, J. T. A., Verdoes Kleijn, G. A., Erben, T., et al. 2017, *A&A*, **604**, A134

de la Bella, L. F., Tessore, N., & Bridle, S. 2021, *JCAP*, **2021**, 001

Edge, A., Sutherland, W., Kuijken, K., et al. 2013, *The Messenger*, **154**, 32

Eifler, T., Schneider, P., & Hartlap, J. 2009, *A&A*, **502**, 721

Erben, T., Schirmer, M., Dietrich, J. P., et al. 2005, *Astron. Nachr.*, **326**, 432

Euclid Collaboration (Ajani, V., et al.) 2023, *A&A*, **675**, A120

Fenech Conti, I., Herbonnet, R., Hoekstra, H., et al. 2017, *MNRAS*, **467**, 1627

Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601

Fluri, J., Kacprzak, T., Lucchi, A., et al. 2022, *Phys. Rev. D*, **105**, 083518

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306

Fortuna, M. C., Hoekstra, H., Johnston, H., et al. 2021, *A&A*, **654**, A76

Franco-Abellán, G., Cañas-Herrera, G., Martinelli, M., et al. 2024, *JCAP*, **2024**, 057

Garrel, C., Pierre, M., Valageas, P., et al. 2022, *A&A*, **663**, A3

Gatti, M., Chang, C., Friedrich, O., et al. 2020, *MNRAS*, **498**, 4060

Gatti, M., Sheldon, E., Amon, A., et al. 2021, *MNRAS*, **504**, 4312

Gatti, M., Jeffrey, N., Whiteway, L., et al. 2024, *Phys. Rev. D*, **109**, 063534

Gelman, A., Meng, X. L., & Stern, H. 1996, *Stat. Sin.*, **6**, 733

Gerardi, F., Feeney, S. M., & Alsing, J. 2021, *Phys. Rev. D*, **104**, 083531

Germain, M., Gregor, K., Murray, I., & Larochelle, H. 2015, in *Proceedings of the 32nd International Conference on Machine Learning*, eds. F. Bach, & D. Blei (Lille, France: PMLR), *Proc. Mach. Learn. Res.*, **37**, 881

Ghirardini, V., Bulbul, E., Artis, E., et al. 2024, *A&A*, **689**, A298

Giblin, B., Heymans, C., Asgari, M., et al. 2021, *A&A*, **645**, A105

Giacoli, C., Di Meo, S., Meneghetti, M., et al. 2017, *MNRAS*, **470**, 3574

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759

Guzik, J., & Bernstein, G. 2005, *Phys. Rev. D*, **72**, 043503

Hall, A., & Taylor, A. 2022, *Phys. Rev. D*, **105**, 123527

Hamimeche, S., & Lewis, A. 2008, *Phys. Rev. D*, **77**, 103013

Handley, W., & Lemos, P. 2021, *Phys. Rev. D*, **103**, 063529

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357

Heavens, A., Refregier, A., & Heymans, C. 2000a, *MNRAS*, **319**, 649

Heavens, A. F., Jimenez, R., & Lahav, O. 2000b, *MNRAS*, **317**, 965

Heydenreich, S., Schneider, P., Hildebrandt, H., et al. 2020, *A&A*, **634**, A104

Heymans, C., & Heavens, A. 2003, *MNRAS*, **339**, 711

Heymans, C., Van Waerbeke, L., Bacon, D., et al. 2006, *MNRAS*, **368**, 1323

Heymans, C., Grocutt, E., Heavens, A., et al. 2013, *MNRAS*, **432**, 2433

Heymans, C., Tröster, T., Asgari, M., et al. 2021, *A&A*, **646**, A140

Hikage, C., Takada, M., Hamana, T., & Spergel, D. 2011, *MNRAS*, **412**, 65

Hikage, C., Oguri, M., Hamana, T., et al. 2019, *PASJ*, **71**, 43

Hilbert, S., Hartlap, J., & Schneider, P. 2011, *A&A*, **536**, A85

Hilbert, S., Xu, D., Schneider, P., et al. 2017, *MNRAS*, **468**, 790

Hildebrandt, H., Viola, M., Heymans, C., et al. 2017, *MNRAS*, **465**, 1454

Hildebrandt, H., van den Busch, J. L., Wright, A. H., et al. 2021, *A&A*, **647**, A124

- Hirata, C. M., Mandelbaum, R., Seljak, U., et al. 2004, *MNRAS*, **353**, 529
- Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, **567**, 2
- Howlett, C., Lewis, A., Hall, A., & Challinor, A. 2012, *JCAP*, **2012**, 027
- Hu, W. 2000, *Phys. Rev. D*, **62**, 043007
- Hu, T., Khaire, V., Hennawi, J. F., et al. 2022, *MNRAS*, **515**, 2188
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Jarvis, M., Bernstein, G., & Jain, B. 2004, *MNRAS*, **352**, 338
- Jeffrey, N., Alsing, J., & Lanusse, F. 2021, *MNRAS*, **501**, 954
- Jeffrey, N., Whiteway, L., Gatti, M., et al. 2025, *MNRAS*, **536**, 1303
- Joachimi, B., Mandelbaum, R., Abdalla, F. B., & Bridle, S. L. 2011, *A&A*, **527**, A26
- Joachimi, B., Lin, C. A., Asgari, M., et al. 2021, *A&A*, **646**, A129
- Johnston, H., Wright, A. H., Joachimi, B., et al. 2021, *A&A*, **648**, A98
- Joudaki, S., Blake, C., Johnson, A., et al. 2018, *MNRAS*, **474**, 4894
- Kaiser, N. 1992, *ApJ*, **388**, 272
- Kaiser, N. 1998, *ApJ*, **498**, 26
- Kannawadi, A., Hoekstra, H., Miller, L., et al. 2019, *A&A*, **624**, A92
- Kayo, I., Taruya, A., & Suto, Y. 2001, *ApJ*, **561**, 22
- Kilbinger, M. 2015, *Rep. Prog. Phys.*, **78**, 086901
- Kilbinger, M., Heymans, C., Asgari, M., et al. 2017, *MNRAS*, **472**, 2126
- King, L., & Schneider, P. 2002, *A&A*, **396**, 411
- Kitching, T. D., & Heavens, A. F. 2017, *Phys. Rev. D*, **95**, 063522
- Kitching, T. D., Paykari, P., Hoekstra, H., & Cropper, M. 2019, *Open J. Astrophys.*, **2**, 5
- Cluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, *Jupyter Notebooks—a publishing format for reproducible computational workflows* (IOS Press), 87
- Krause, E., Eifler, T., & Blazek, J. 2016, *MNRAS*, **456**, 207
- Kuijken, K. 2011, *The Messenger*, **146**, 8
- Kuijken, K., Heymans, C., Hildebrandt, H., et al. 2015, *MNRAS*, **454**, 3500
- Kuijken, K., Heymans, C., Dvornik, A., et al. 2019, *A&A*, **625**, A2
- Lam, S. K., Pitrou, A., & Seibert, S. 2015, *Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC*, 1
- Lange, J. U. 2023, *MNRAS*, **525**, 3181
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, *ArXiv e-prints* [arXiv:1110.3193]
- Leclercq, F. 2018, *Phys. Rev. D*, **98**, 063511
- Legin, R., Hezaveh, Y., Perreault Levasseur, L., & Wandelt, B. 2021, *ArXiv e-prints* [arXiv:2112.05278]
- Leistedt, B., Peiris, H. V., Mortlock, D. J., Benoit-Lévy, A., & Pontzen, A. 2013, *MNRAS*, **435**, 1857
- Lemos, P., Challinor, A., & Efstathiou, G. 2017, *JCAP*, **2017**, 014
- Lemos, P., Jeffrey, N., Whiteway, L., et al. 2021, *Phys. Rev. D*, **103**, 023009
- Lemos, P., Coogan, A., Hezaveh, Y., & Perreault-Levasseur, L. 2023a, in *40th International Conference on Machine Learning*, 202, 19256
- Lemos, P., Cranmer, M., Abidi, M., et al. 2023b, *Mach. Learn. Sci. Technol.*, **4**, 01LT01
- Lemos, P., Parker, L. H., Hahn, C., et al. 2023c, in *Mach. Learn. Astrophys.*, **18**, 18
- Leonard, C. D., Ferreira, T., Fang, X., et al. 2023, *Open J. Astrophys.*, **6**, 8
- Levin, D. 1996, *J. Comput. Appl. Math.*, **67**, 95
- Lewis, A. 2019, *ArXiv e-prints* [arXiv:1910.13970]
- Lewis, A., & Challinor, A. 2002, *Phys. Rev. D*, **66**, 023531
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, **538**, 473
- Li, S.-S., Hoekstra, H., Kuijken, K., et al. 2023a, *A&A*, **679**, A133
- Li, X., Zhang, T., Sugiyama, S., et al. 2023b, *Phys. Rev. D*, **108**, 123518
- Limber, D. N. 1953, *ApJ*, **117**, 134
- Lin, C.-A., & Kilbinger, M. 2015, *A&A*, **583**, A70
- Lin, C.-A., Kilbinger, M., & Pires, S. 2016, *A&A*, **593**, A88
- Lin, C.-H., Harnois-Déraps, J., Eifler, T., et al. 2020, *MNRAS*, **499**, 2977
- Lin, K., von Wietersheim-Kramsta, M., Joachimi, B., & Feeney, S. 2023, *MNRAS*, **524**, 6167
- Loureiro, A., Whittaker, L., Spurio Mancini, A., et al. 2022, *A&A*, **665**, A56
- Loureiro, A., Whiteway, L., Sellentin, E., et al. 2023, *Open J. Astrophys.*, **6**, 6
- LoVerde, M., & Afshordi, N. 2008, *Phys. Rev. D*, **78**, 123506
- Alonso, D., Sanchez, J., Slosar, A., & LSST Dark Energy Science Collaboration 2019, *MNRAS*, **484**, 4127
- LSST Science Collaboration (Abell, P. A., et al.) 2009, *ArXiv e-prints* [arXiv:0912.0201]
- Lu, T., Haiman, Z., & Li, X. 2023, *MNRAS*, **521**, 2050
- Mandelbaum, R. 2018, *ARA&A*, **56**, 393
- Mantz, A. B., von der Linden, A., Allen, S. W., et al. 2015, *MNRAS*, **446**, 2205
- McFarland, J. P., Verdoes-Kleijn, G., Sikkema, G., et al. 2013, *Exp. Astron.*, **35**, 45
- Mead, A. J., Peacock, J. A., Heymans, C., Joudaki, S., & Heavens, A. F. 2015, *MNRAS*, **454**, 1958
- Mead, A. J., Heymans, C., Lombriser, L., et al. 2016, *MNRAS*, **459**, 1468
- Mead, A. J., Brieden, S., Tröster, T., & Heymans, C. 2021, *MNRAS*, **502**, 1401
- Miller, L., Kitching, T. D., Heymans, C., Heavens, A. F., & van Waerbeke, L. 2007, *MNRAS*, **382**, 315
- Miller, L., Heymans, C., Kitching, T. D., et al. 2013, *MNRAS*, **429**, 2858
- Mishra-Sharma, S., & Cranmer, K. 2022, *Phys. Rev. D*, **105**, 063017
- Miyatake, H., Sugiyama, S., Takada, M., et al. 2023, *Phys. Rev. D*, **108**, 123517
- Moser, B., Kacprzak, T., Fischbacher, S., et al. 2024, *JCAP*, **2024**, 049
- Nelder, J. A., & Mead, R. 1965, *Comput. J.*, **7**, 308
- Nicola, A., García-García, C., Alonso, D., et al. 2021, *JCAP*, **2021**, 067
- Papamakarios, G., & Murray, I. 2016, *ArXiv e-prints* [arXiv:1605.06376]
- Papamakarios, G., Pavlakou, T., & Murray, I. 2017, in *Advances in Neural Information Processing Systems*, eds. I. Guyon, U. V. Luxburg, S. Bengio, et al. (Curran Associates Inc.), 30
- Peebles, P. J. E. 1973, *ApJ*, **185**, 413
- Piras, D., Joachimi, B., & Villaescusa-Navarro, F. 2023, *MNRAS*, **520**, 668
- Planck Collaboration VI. 2020, *A&A*, **641**, A6
- Porqueres, N., Heavens, A., Mortlock, D., & Lavaux, G. 2022, *MNRAS*, **509**, 3194
- Porqueres, N., Heavens, A., Mortlock, D., Lavaux, G., & Makinen, T. L. 2023, *ArXiv e-prints* [arXiv:2304.04785]
- Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A., & Feldman, M. W. 1999, *Molecular biology and evolution*, **16**, 1791
- Reinecke, M. 2011, *A&A*, **526**, A108
- Reischke, R., Kiessling, A., & Schäfer, B. M. 2017, *MNRAS*, **465**, 4016
- Reischke, R., Unruh, S., Asgari, M., et al. 2024, *A&A*, submitted
- Robert, C. P. 2007, *The Bayesian choice: from decision-theoretic foundations to computational implementation* (Springer), 2
- Rubin, D. B. 1984, *Ann. Stat.*, **12**, 1151
- Schneider, P. 2016, *A&A*, **592**, L6
- Schneider, P., & Hartlap, J. 2009, *A&A*, **504**, 705
- Schneider, A., & Teyssier, R. 2015, *JCAP*, **2015**, 049
- Schneider, A., Teyssier, R., Stadel, J., et al. 2019, *JCAP*, **2019**, 020
- Schneider, A., Giri, S. K., Amodeo, S., & Refregier, A. 2022, *MNRAS*, **514**, 3802
- Scolnic, D. M., Jones, D. O., Rest, A., et al. 2018, *ApJ*, **859**, 101
- Secco, L. F., Samuroff, S., Krause, E., et al. 2022, *Phys. Rev. D*, **105**, 023515
- Seitz, S., & Schneider, P. 1996, *A&A*, **305**, 383
- Sellentin, E., & Heavens, A. F. 2018, *MNRAS*, **473**, 2355
- Sellentin, E., Heymans, C., & Harnois-Déraps, J. 2018, *MNRAS*, **477**, 4879
- Shirasaki, M., Hamana, T., Takada, M., Takahashi, R., & Miyatake, H. 2019, *MNRAS*, **486**, 52
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, *ArXiv e-prints* [arXiv:1503.03757]
- Spurio Mancini, A., Reischke, R., Pettorino, V., Schäfer, B. M., & Zumalacárregui, M. 2018a, *MNRAS*, **480**, 3725
- Spurio Mancini, A., Taylor, P. L., Reischke, R., et al. 2018b, *Phys. Rev. D*, **98**, 103507
- Spurio Mancini, A., Docherty, M. M., Price, M. A., & McEwen, J. D. 2023, *RAS Techn. Instrum.*, **2**, 710
- Sugiyama, S., Takada, M., Miyatake, H., et al. 2022, *Phys. Rev. D*, **105**, 123537
- Taylor, P. L., Kitching, T. D., Alsing, J., et al. 2019, *Phys. Rev. D*, **100**, 023519
- Tegmark, M., Taylor, A. N., & Heavens, A. F. 1997, *ApJ*, **480**, 22
- Tessore, N., Loureiro, A., Joachimi, B., von Wietersheim-Kramsta, M., & Jeffrey, N. 2023, *Open J. Astrophys.*, **6**, 11
- Tröster, T., Asgari, M., Blake, C., et al. 2021, *A&A*, **649**, A88
- Tröster, T., Mead, A. J., Heymans, C., et al. 2022, *A&A*, **660**, A27
- Upham, R. E., Brown, M. L., & Whittaker, L. 2021, *MNRAS*, **503**, 1999
- van den Busch, J. L., Wright, A. H., Hildebrandt, H., et al. 2022, *A&A*, **664**, A170
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Methods*, **17**, 261
- Wright, A. H., Hildebrandt, H., Kuijken, K., et al. 2019, *A&A*, **632**, A34
- Wright, A. H., Hildebrandt, H., van den Busch, J. L., & Heymans, C. 2020, *A&A*, **637**, A100
- Xavier, H. S., Abdalla, F. B., & Joachimi, B. 2016, *MNRAS*, **459**, 3693
- Zablocki, A., & Dodelson, S. 2016, *Phys. Rev. D*, **93**, 083525
- Zaldarriaga, M., & Seljak, U. 1997, *Phys. Rev. D*, **55**, 1830
- Zieser, B., & Merkel, P. M. 2016, *MNRAS*, **459**, 1586
- Zonca, A., Singer, L., Lenz, D., et al. 2019, *J. Open Source Software*, **4**, 1298
- Zuntz, J., Paterno, M., Jennings, E., et al. 2015, *Astron. Comput.*, **12**, 45
- Zuntz, J., Sheldon, E., Samuroff, S., et al. 2018, *MNRAS*, **481**, 1149

Appendix A: Numerical non-Limber integration

Levin's method (Levin 1996) casts a quadrature problem of an oscillatory integral into the solution of a system of ordinary and linear differential equations. Assuming a set basis for the solution the solution to the set of differential equations can be cast into a simple linear algebra problem. It has been used in a variety of cosmic shear applications (Zieser & Merkel 2016; Spurio Mancini et al. 2018a,b; Baleato Lizancos & White 2023). The method relies on integrals of the type:

$$I[h(x, k)] = \int_{x_1}^{x_2} dx \langle \mathbf{f}(x, k), \mathbf{w}(k, x) \rangle, \quad (\text{A.1})$$

where $\langle \cdot, \cdot \rangle$ denotes a scalar product, which, in this case, is taken between the vector of non-oscillatory functions \mathbf{f} and the oscillatory part \mathbf{w} , which has to satisfy:

$$\frac{d\mathbf{w}(x)}{dx} = \mathbf{A}\mathbf{w}(x), \quad (\text{A.2})$$

with a matrix \mathbf{A} . Therefore any oscillatory function with recursion relations is very well suited for this formalism.

The integral in Eq. A.1 is approximated by finding a vector \mathbf{p} such that:

$$\langle \mathbf{p}, \mathbf{w} \rangle' = \langle \mathbf{p}' + \mathbf{A}^T \mathbf{p}, \mathbf{w} \rangle \approx \langle \mathbf{F}, \mathbf{w} \rangle. \quad (\text{A.3})$$

Thus $\langle \mathbf{p}, \mathbf{w} \rangle' = \langle \mathbf{F}, \mathbf{w} \rangle$, at n collocation points $x_j, j = 1, 2, \dots, n$, is subject to the following set of equations:

$$\langle \mathbf{p}' + \mathbf{A}^T \mathbf{p} - \mathbf{F}, \mathbf{w} \rangle(x_j) = 0, \quad j = 1, \dots, n. \quad (\text{A.4})$$

Here the trivial solution is the null vector:

$$\mathbf{p}'(x_j) + \mathbf{A}^T(x_j)\mathbf{p}(x_j) = \mathbf{F}(x_j). \quad (\text{A.5})$$

In a last step a solution is constructed for \mathbf{p} using n differentiable basis functions $u_m(x)$ that \mathbf{p} can be expanded in:

$$p_i(x) = \sum_m c_i^{(m)} u_m(x), \quad i = 1, \dots, d; \quad m = 1, \dots, n. \quad (\text{A.6})$$

Implying:

$$\sum_m c_i^{(m)} u_m'(x_j) + \sum_{m,q} A_{qi} c_q^{(m)} u_m(x_j) = F_i(x_j), \quad (\text{A.7})$$

where $i, q = 1, \dots, d; j, m = 1, \dots, n$. Any basis is suitable, here we chose equidistant collocation points x_j and the n lowest-order polynomials as basis functions. Finally, the vector \mathbf{w} depends on the specific type of integral considered. Equation 8 allows for two possibilities: (i) first integrate over the co-moving distance χ ; (ii) first integrate over the wave-vector k . Both approaches have advantages and disadvantages: due to the separation of the χ and χ' dependence, integrating over the co-moving distance first will be quicker. However, the residual integral over k can still be very oscillatory and must be sampled accordingly, especially with narrow window functions $W(\chi)$. The second option is to integrate over k first, this smooths out the oscillations from the Bessel function. However, there are two remaining integrals, thus slowing down the computation, especially for many radial bins. If, however, the radial bins do not overlap, the resulting spectrum for bins further apart than a single bin index will be

very small and can be neglected. A more detailed discussion on this is provided in Leonard et al. (2023). For integrals of the type

$$I_{\ell,k} = \int dx f_k(x) j_\ell(kx), \\ I_{\ell,k_1,k_2} = \int dx f_{k_1,k_2}(x) j_\ell(k_1 x) j_\ell(k_2 x), \quad (\text{A.8})$$

we chose

$$\mathbf{w}(x) = \begin{pmatrix} j_\ell(xk) \\ j_{\ell+1}(xk) \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} \frac{\ell}{x} & -k \\ -\frac{\ell+2}{x} & k \end{pmatrix}; \\ \mathbf{w}(x) = \begin{pmatrix} j_\ell(xk_1)j_\ell(xk_2) \\ j_{\ell+1}(xk_1)j_\ell(xk_2) \\ j_\ell(xk_1)j_{\ell+1}(xk_2) \\ j_{\ell+1}(xk_1)j_{\ell+1}(xk_2) \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} \frac{2\ell}{x} & -k_1 & -k_2 & 0 \\ k_1 & -\frac{2}{x} & 0 & -k_2 \\ k_2 & 0 & -\frac{2}{x} & -k_1 \\ 0 & k_2 & k_1 & -\frac{2\ell+2}{x} \end{pmatrix}, \quad (\text{A.9})$$

respectively.

Appendix B: Accuracy of the cosmic shear signal

To test the robustness of the cosmic shear signal as predicted by the log-normal simulations generated in KiDS-SBI in addition to the self-consistency tests already conducted within Sect. 3, we perform a comparison of the noise-free pseudo-CLs as measured in KiDS-SBI assuming the standard isotropic systematics model (see Sects. 4.2 and 7.2) to the pseudo-CLs predicted from analytical theory (see Appendix D). This comparison is shown for the KiDS-SBI MAP cosmology in Fig. B.1. When assuming 19 shells and a comoving-volume weighted kernel as discussed in Sect. 3.2, we find per-cent or sub-per-cent level agreement for all tomographic bin combinations and scales (within 0.5σ or less). The agreement between the models is also shown to be good in Fig. 16 where the posterior from the SBI analysis based on KiDS-SBI is unbiased with respect to the posterior sampled with an MCMC while assuming analytical theory. Overall, we find that for small ℓ the forward simulations predict additional power with respect to the pseudo-CLs predicted from theory. This can be explained by discretisation effects along the line of sight due to the division of the observable Universe into concentric shells weighted by a window function (see Sect. 3.2), and by the linear interpolation of convergence and shear across HEALPix pixels which are built into SALMO. Both effects have been shown to produce a similar minor excess in angular power in testing conducted in Tessore et al. (2023).

When increasing the resolution along the line of sight by increasing the number of shells, and when changing the kernel to a uniform window function as suggested in Tessore et al. (2023), we find that for 26 shells the agreement with the signal predicted by analytical theory is marginally better, but not enough to impact the SBI. When increasing the number of shells further to 31 shells, such that each shell has a width of $\sim 100 \text{ Mpc } h^{-1}$, the agreement with analytical theory is poorer for large ℓ and for low ℓ in the autocorrelations, as discretisation effects become important. Consequently, we conclude that the choice of 19 shells allows us to carry out an accurate forward modelling of the cosmic shear signal in KiDS-SBI, while still allowing for computational efficiency (see Sect. 5.1).

These effects are not significant enough to bias the cosmological inference at the precision of KiDS-1000, particularly, since they are more prevalent in tomographic bin

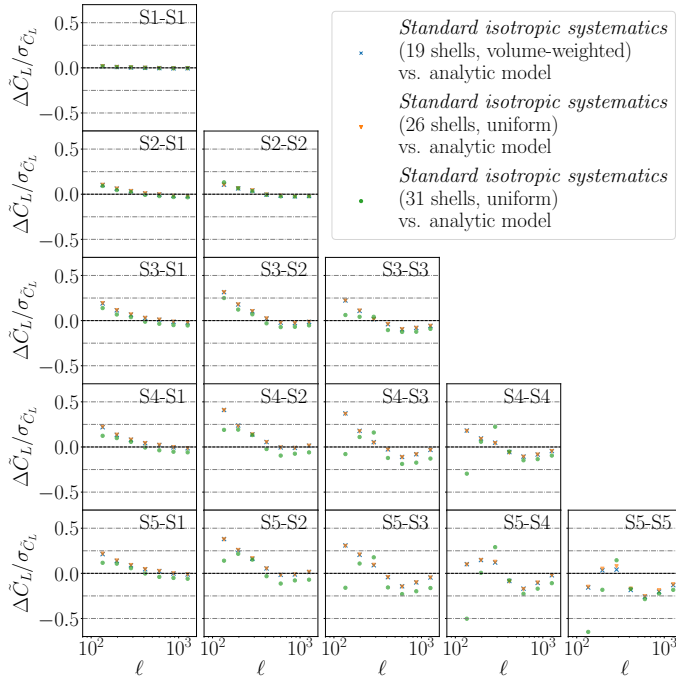


Fig. B.1. Relative difference between the mean pseudo-CIs measured from the standard isotropic systematics model in KiDS-SBI with different line-of-sight resolutions and the expectation from analytical theory as described in Appendix D. The blue crosses show the standard isotropic systematics model with the resolution assumed in the analysis discussed in this work: 19 shells with a comoving-volume weighted kernel. The orange triangles and the green circles show difference in the signal of standard isotropic systematics model with a uniform kernel and at higher resolutions: 26 shells and 31 shells, respectively. All are evaluated at the MAP cosmology inferred from the KiDS-1000 data using the standard isotropic systematics in KiDS-SBI (see Table F.1).

combinations with a lower signal-to-noise ratio. However, future simulation-based inference analyses of galaxy surveys with smaller uncertainties (for instance stage-IV surveys) based on similar simulations should aim to account for such effects.

Moreover, Fig. 16 shows that in the bin combinations at higher redshifts (with the highest signal-to-noise ratio) the angular power in the high ℓ modes tends to be slightly underestimated by the standard isotropic systematics model when compared to the analytical theory. We find that this is driven by more cosmic shear signal leaking from the EE modes to the BB modes than predicted by the mixing matrix discussed in Appendix D. This should not bias the SBI analysis, since the KiDS-1000 data and the forward simulations share the same measurement procedure, so mode mixing effects should be intrinsically ingrained in both the real and the forward-modelled data vectors.

Appendix C: Effects of variable depth and shear biases

To quantify the physical effect of the inclusion of variable depth and shear biases into the anisotropic systematics model as discussed in Sect. 3.8 and 3.7, we isolate the bias each systematic causes in the observed two-point cosmic shear signal. First, we forward model variable depth in KiDS-1000 as described in Sect. 3.8, and compare it to a set of forward simulations where the galaxy density, intrinsic galaxy shape dispersion and the photometric redshift distributions are isotropic for each tomographic bin, as is the case in the standard isotropic systemat-




Fig. C.1. The impact of variable depth in KiDS-1000 on the cosmic shear real-space two-point correlation functions, $\xi(\theta)$ in the autocorrelation of the S5 bin. The blue line shows the difference in the mean signal from 5,000 realisations of the anisotropic systematics model with variable depth versus the mean signal from the same number of realisations when assuming that the observational depth is isotropic. The orange line shows the same difference in signal as predicted for KiDS-1000 by the analytical model in Heydenreich et al. (2020). We note that variable depth constitutes approximately 1% of the measured cosmic shear signal in KiDS-1000.

ics model. When comparing the cosmic shear signal through the spatial two-point correlations functions (2PCF) measured from TreeCorr (Jarvis et al. 2004) in Fig. C.1, we find that variable depth adds additional two-point correlation signal at all scales and particularly at scales where $\theta < 100$ arcmin. In this case, we measure the cosmic shear signal with 2PCF instead of pseudo-CIs as it simplifies the comparison of the results with previous work. This additional signal is equivalent to about 1% ($\sim 0.2\sigma$) additional non-cosmological shear signal which is in agreement with predictions made for KiDS-1000 from Heydenreich et al. (2020). The effect is similar, but slightly smaller in the lower redshift tomographic bin combinations of KiDS-1000 where the shear signal is weaker.

Additionally, when modelling variable depth, the shape noise becomes anisotropic across the sky as a function of the variations in the effective galaxy density and the intrinsic galaxy shape dispersion. When considering a Gaussian likelihood, this implies that the shape noise scales with $\sigma_\epsilon^2/n_{\text{eff}}$ (Kaiser 1992, 1998). As discussed in Sect. 3.8 both quantities scale linearly with the variable depth measure, σ_{rms} , which implies that the shape noise in a given tomographic bin either scales with σ_{rms}^3 in S1, S2 and S5, or with σ_{rms}^{-1} in S3 and S5. Since in KiDS-1000 the distribution of σ_{rms} is skewed below the mean (see Fig. C.2), variable depth in fact reduces the overall contribution of the shape noise to the uncertainty of the cosmic shear signal. Nevertheless, as shown in Figs. 7 and 8, variable depth also causes additional non-cosmological angular correlations in the data at the scales of pixels, pointing overlaps, whole pointings or even over the whole footprint, while also changing the signal along the line of sight anisotropically as shown in Fig. 10. These additional correlations can significantly alter the covariance of the two-point statistic or may even lead to non-Gaussian noise. Such correlations can exacerbate the uncertainty contributions from cosmic variance and super-sample covariance, which causes the noise predicted in KiDS-1000 by a model including variable depth to




Fig. C.2. Histogram showing the distribution of the pixelised root-mean-square of the background noise, σ_{rms} , within the KiDS-1000 footprint shown in Fig. 7. The solid black line shows the mean value of σ_{rms} , while the dashed black line indicates the median of the distribution. This shows that the majority of the pixels in the KiDS-1000 footprint have lower σ_{rms} than the mean which is realised in the case of isotropic depth, i.e. in the standard isotropic systematics model.

ultimately exceed the noise predicted by a model which does not consider the systematic (Reischke et al. 2024).

When isolating the effect of additive and PSF variation shear bias on the observed cosmic shear signal as described in Sect. 3.7 in Fig. C.3, we find that, while the additive shear bias has a negligible effect on the cosmic shear signal, the shear bias originating from PSF variations can bias the pseudo-Cl's by up to 0.3σ particularly at large scales. Hence, the inclusion of the PSF shear bias can contribute a percentage level non-cosmological signal on certain scales.

Appendix D: Theoretical signal modelling

To provide consistency checks, we also theoretically model the expected signal for a given mock. We model the shape-shape angular power spectrum, $C_{\epsilon\epsilon}^{(pq)}(\ell)$, as follows

$$C_{\epsilon\epsilon}^{(pq)}(\ell) = C_{\text{gg}}^{(pq)}(\ell) + C_{\text{gl}}^{(pq)}(\ell) + C_{\text{lg}}^{(pq)}(\ell) + C_{\text{II}}^{(pq)}(\ell), \quad (\text{D.1})$$

where ϵ indicates correlations with a galaxy shape field, g indicates correlations with a cosmic shear field, and I indicates correlations between galaxy shapes due to intrinsic alignments. To model each of these terms, we calculate the matter power spectrum using CAMB and proceed to make a Limber projection to obtain the following expression

$$C_{ab}^{(pq)}(\ell) = \int_0^\infty \frac{d\chi}{f_k^2(\chi)} W_a^{(p)}(\chi) W_b^{(q)}(\chi) P_\delta\left(\frac{\ell + 1/2}{f_k(\chi)}, \chi\right), \quad (\text{D.2})$$

where W_a is the kernel for a given field, such that $a, b \in \{\text{g}, \text{I}\}$. For cosmic shear, the kernel is given by

$$W_{\text{g}}^{(p)}(\chi) = \frac{3H_0^2\Omega_m}{2} \frac{f_k(\chi)}{a(\chi)} \int_\chi^{\chi_{\text{hor}}} d\chi' n_s^{(p)}(\chi') \frac{f_k(\chi' - \chi)}{f_k(\chi')}, \quad (\text{D.3})$$

where $n_s^{(p)}$ is the redshift distribution of the objects in p^{th} tomographic bin as shown in Fig. 2. For consistency with the forward modelling described in Sect. 3.5, we chose an intrinsic alignment kernel in accordance with the non-linear alignment model

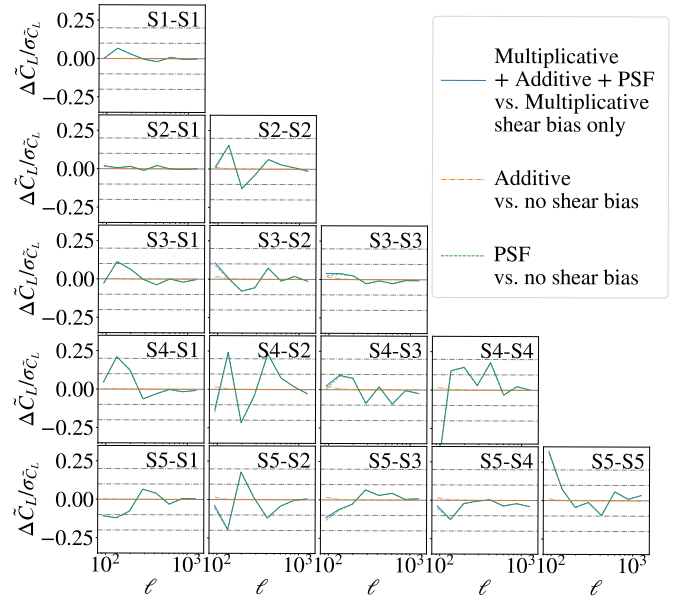


Fig. C.3. Relative difference between the mean pseudo-Cl's measured in KiDS-SBI for different combinations of shear biases in KiDS-1000 as described in Sect. 3.7. The blue solid line compares the pseudo-Cl's of a model that adds multiplicative, additive and PSF shear biases (as is the case in the anisotropic systematics model) to the signal measured from a model which only considers multiplicative shear bias (i.e. the standard isotropic systematics model). The orange dot-dashed line shows the effect of the additive shear bias alone, while the green dashed line shows the effect of the PSF shear bias alone. We note that, in particular, the PSF shear bias causes an up to 0.3σ bias in the pseudo-Cl's, especially, at large scales.

(NLA; Catelan et al. 2001; Hirata et al. 2004; Bridle & King 2007) given by

$$W_1^{(p)}(\chi) = -A_{\text{IA}} \frac{C_1 \bar{\rho}_{\text{cr}}(\chi) \Omega_m}{D(\chi)} n_s^{(p)}(\chi). \quad (\text{D.4})$$

Knowing the full-sky angular power spectra, $C_{\epsilon\epsilon, \nu}^{(pq)}(\ell'; \Theta)$, we can define the pseudo-Cl, $\tilde{C}_{\epsilon\epsilon, \mu}^{(pq)}(\ell; \Theta)$, as follows

$$\tilde{C}_{\epsilon\epsilon, \mu}^{(pq)}(\ell; \Theta) = \sum_{\ell'=0}^{\ell'_{\text{max}}} \sum_{\ell''=0}^{\ell''_{\text{max}}} \sum_{\nu=1}^3 \sum_{\nu'=1}^3 M_{\mu\nu', \ell\ell''} M_{\nu'\nu, \ell'\ell''}^{(pq)} C_{\epsilon\epsilon, \nu}^{(pq)}(\ell'; \Theta), \quad (\text{D.5})$$

where $\mu, \nu, \nu' \in \{1, 2, 3\}$ such that 1 stands for the EE component, 2 for the BB component and 3 for the EB component, $M_{\ell\ell''}$ is the mixing matrix of the survey mask, $W(\theta)$, which does not vary between tomographic bins, and $M_{\ell'\ell''}^{(pq)}$ is the mixing matrix of the effective mask imposed by the average random variations in the observed galaxies in a given tomographic bin over many realisations.

The first mixing matrix from the left in Eq. (D.5), $M_{\ell\ell''}$, is the standard block mixing matrix to account for the partial sky coverage in a survey (Hivon et al. 2002; Bennett et al. 2003; Brown et al. 2005). The second mixing matrix, $M_{\ell'\ell''}^{(pq)}$, accounts for an additional effect: the mode mixing induced by a mask of uncorrelated noise originating from the random sampling described in Sect. 3.6. In other words, when randomly Poisson sampling galaxies within a discrete pixel k given a probability of $P^{(m)}(N)$, there will be a non-zero probability of observing zero counts in that given pixel. In essence, all the pixels which are sampled to have zero galaxy counts impose an additional

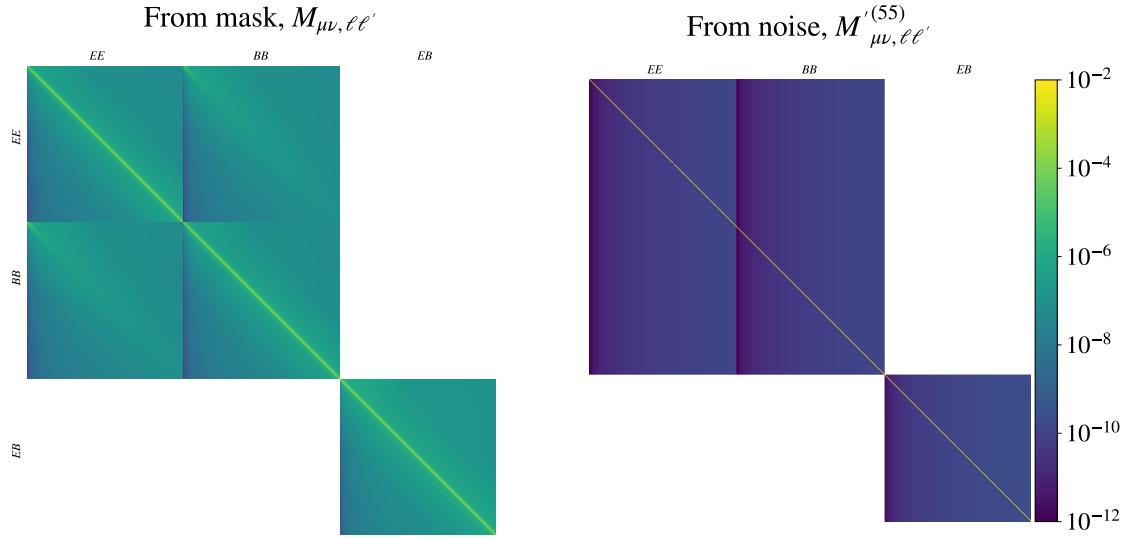


Fig. D.1. Bitmap of the mixing matrices to model the theoretical signal for the pseudo-Cl's as seen by KiDS-1000. On the right panel is the mixing matrix derived from the KiDS-1000 mask, $M_{\mu\nu, \ell\ell'}$, (the input mask is at a resolution of $N_{\text{side}} = 1024$). On the left panel is the mixing matrix caused by the selection due to randomly sampling galaxies for the auto-correlation of the fifth tomographic bin, $M_{\mu\nu, \ell\ell'}^{(55)}$. Both matrices are decomposed into block matrices separating the $EE \rightarrow EE$, $EE \rightarrow BB$ and $EB \rightarrow EB$ mixing. Each block matrix has dimensions of 8193×8193 with $\ell \in \{\ell \in \mathbb{Z}^{0+} | \ell \leq 8192\}$.

masking on our sample. This can become significant at the resolutions and galaxy densities which we probe in this work. For example, for an HealPix $N_{\text{side}} = 2048$ and a galaxy density of $n = 1 \text{ arcmin}^{-2}$ for a given tomographic bin k (similar to KiDS-1000), $P^{(m)}(N = 0) \approx 0.07$. This means that nearly a tenth of all pixels could be 'masked' in a given realisation of simulations at a higher resolution. To allow for these cases, we decide to also model the mode mixing induced by the sampling variance. However, for the settings discussed in this paper for the main pipeline of forwards simulations where $N_{\text{side}} = 1024$, this effect is not so relevant. Even for the sparsest source bin, S1, with an $n_{\text{gal}}^{(1)} = 0.62 \text{ arcmin}^{-2}$, $P^{(m)}(N = 0) \approx 0.0005$, so only a negligible fraction of pixels are masked due to random sample variance.

To provide a general overview, the distinguishing factor between the two mixing matrices is how we define the mask $W(\theta)$ which characterises them. For $M_{\mu\nu, \ell\ell'}$, $W(\theta)$ is defined as the mask of the survey footprint, Ω_{survey} , such that

$$W(\theta) = \begin{cases} 1, & \theta \in \Omega_{\text{survey}} \\ 0, & \theta \notin \Omega_{\text{survey}} \end{cases}. \quad (\text{D.6})$$

For $M_{\mu\nu, \ell\ell'}^{(pq)}$, rather than defining it using an explicit spatial mask from the pixels masked by each random realisation, $W^{(p)}(\theta)$, we can define the expectation value of the mixing matrix due to the random sampling of galaxies directly from the angular power spectrum of $W^{(p)}(\theta)$ as

$$\begin{aligned} \frac{\mathcal{W}^{(pq)}(\ell)}{2\ell + 1} &= \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} w_{\ell m}^{(p)} w_{\ell m}^{(q)*} \\ &= \frac{4\pi\sigma^{(pq)}}{N_{\text{pix}} - 1} + 4\pi \left[\mu^{(pq)} - \frac{\sigma^{(pq)}}{N_{\text{pix}} - 1} \right] \delta_{\ell 0}^{\text{K}}, \end{aligned} \quad (\text{D.7})$$

where $\mu^{(pq)} = P^{(p)}(N > 0)P^{(q)}(N > 0)$, namely the product of the mean probabilities of success within each tomographic bin (the probability that a given pixel will be populated by at least one galaxy), while $2\sigma^{(pq)2} = P^{(p)2}(N > 0)(1 - P^{(p)}(N > 0))^2 +$

$P^{(q)2}(N > 0)(1 - P^{(q)}(N > 0))^2$, expressly the associated standard deviation of the mean, and N_{pix} is the number of pixels.

As described in Sect. 3.6, the galaxies are Poisson sampled according to the probability of success is given in Eq. (20). As we are interested in the mean over many iterations, we can assume that the matter overdensities are small, $1 + b^{(p)}\delta^{(p)} \approx 1$, so we can rewrite the probability as

$$\begin{aligned} P^{(p)}(N > 0) &= 1 - P^{(p)}(N = 0) \\ &= 1 - e^{-\langle N_m^{(p)} \rangle(\Theta)} \\ &\approx 1 - e^{-n_{\text{gal}}^{(p)} A_{\text{pix}}}, \end{aligned} \quad (\text{D.8})$$

where $n_{\text{gal}}^{(p)}$ is the mean galaxy density per tomographic bin p and A_{pix} is the mean pixel size for a given HEALPix resolution. From this, we can calculate both mixing matrices as shown in Fig. D.1. In practice, we can see that from Fig. D.1 that the mode mixing $M_{\mu\nu, \ell\ell'}^{(pq)}$ is negligible when compared to $M_{\mu\nu, \ell\ell'}$ at an $N_{\text{side}} = 1024$. However, we still include the term in the signal modelling to allow the flexibility to choose higher resolutions within KiDS-SBI.

With Eqs. (D.6) and (D.7) defining the selections which define $M_{\mu\nu, \ell\ell'}$ and $M_{\mu\nu, \ell\ell'}^{(pq)}$, respectively, we can compute the matrices as follows (Brown et al. 2005; Hikage et al. 2011):

$$M_{\mu\nu, \ell\ell'} = \begin{pmatrix} W_{\ell\ell'}^{++} & (W_{\ell\ell'}^{+-} + W_{\ell\ell'}^{-+}) & W_{\ell\ell'}^{--} \\ -W_{\ell\ell'}^{+-} & (W_{\ell\ell'}^{++} - W_{\ell\ell'}^{--}) & W_{\ell\ell'}^{+-} \\ W_{\ell\ell'}^{--} & -(W_{\ell\ell'}^{+-} + W_{\ell\ell'}^{-+}) & W_{\ell\ell'}^{++} \end{pmatrix}, \quad (\text{D.9})$$

$$M_{\mu\nu, \ell\ell'}^{(pq)} = \begin{pmatrix} W_{\ell\ell'}'^{++} & (W_{\ell\ell'}'^{+-} + W_{\ell\ell'}'^{-+}) & W_{\ell\ell'}'^{-+} \\ -W_{\ell\ell'}'^{+-} & (W_{\ell\ell'}'^{++} - W_{\ell\ell'}'^{-+}) & W_{\ell\ell'}'^{+-} \\ W_{\ell\ell'}'^{-+} & -(W_{\ell\ell'}'^{+-} + W_{\ell\ell'}'^{-+}) & W_{\ell\ell'}'^{++} \end{pmatrix}^{(pq)}, \quad (\text{D.10})$$

where each of the Wigner symbols is given by

$$W_{\ell\ell'}^{\pm\pm} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \sum_{m'=-\ell'}^{\ell'} W_{\ell\ell' m m'}^{\pm\pm} (W_{\ell'\ell m' m}^{\pm\pm})^*. \quad (\text{D.11})$$

Since the shear field is a spin 2 field

$$\begin{aligned} W_{\ell\ell'mm'}^+ &= \frac{1}{2}(2W_{\ell\ell'}^{mm'} + {}_{-2}W_{\ell\ell'}^{mm'}), \\ W_{\ell\ell'mm'}^- &= \frac{i}{2}(2W_{\ell\ell'}^{mm'} - {}_{-2}W_{\ell\ell'}^{mm'}), \end{aligned} \quad (\text{D.12})$$

where

$${}_sW_{\ell\ell'mm'}^+ = \int d\theta {}_sY_{\ell'm'}(\theta) W(\theta) {}_sY_{\ell m}^*(\theta). \quad (\text{D.13})$$

The latter is difficult to compute, so it is more convenient to compute in phase space such that

$$\begin{aligned} \sum_{m=-\ell}^{\ell} \sum_{m'=-\ell'}^{\ell'} {}_sW_{\ell\ell'}^{mm'} ({}_sW_{\ell'\ell}^{m'm})^* \\ = \frac{(2\ell+1)(2\ell'+1)}{4\pi} \sum_{\ell''=0}^{\ell''_{\max}} \mathcal{W}_{\ell''} \begin{pmatrix} \ell & \ell' & \ell'' \\ -s & s & 0 \end{pmatrix} \begin{pmatrix} \ell & \ell' & \ell'' \\ -s' & s' & 0 \end{pmatrix}, \end{aligned} \quad (\text{D.14})$$

where $s \in \{-2, 2\}$ and

$$\mathcal{W}_{\ell} = \sum_{m=-\ell}^{\ell} w_{\ell m} w_{\ell m}^*, \quad (\text{D.15})$$

where $w_{\ell m}$ are the coefficients of the spherical harmonic transform of the mask, $W(\theta)$, such that

$$w_{\ell m} = \int d\theta W(\theta) Y_{\ell m}^*(\theta). \quad (\text{D.16})$$

Appendix E: Parameter recovery

Figure D.2 shows the posterior distribution of the seven main parameters varied in the analysis as learnt from 18,000 realisations of the anisotropic systematics model. The data vector used for this posterior is a noisy mock data vector generated with the anisotropic systematics model while assuming the input parameters given in Table 2. We find that the posteriors from KiDS-SBI accurately recover the true input parameters (see Table E.1 for the inferred parameter values). As shown in Fig. D.2, the anisotropic systematics posterior mostly only constrains S_8 and A_{IA} . At the same time, it becomes apparent that the posterior distributions over ω_b , n_s , A_{bary} , and h_0 are mainly prior driven, namely, the posterior distribution is mostly flat throughout prior space. As the cosmic shear signal is largely degenerate to small changes in these parameters, this is expected.

In summary, we find that the SBI inference pipeline within KiDS-SBI accurately and robustly recovers the underlying posterior distribution of Λ CDM parameters from two-point statistic measurements of the cosmic shear signal in KiDS-1000 when considering a forward model-based log-normal random fields and biased by relevant systematic effects.

Appendix F: Cosmology dependence of the cosmic shear two-point statistic covariance

To determine the origin of the observed cosmology dependence in the likelihood in Fig. 17, we investigate whether there is a theoretical justification for the uncertainty in the cosmic shear signal in KiDS-1000 to vary significantly with S_8 . First, we

Table E.1. Main inferred cosmological and astrophysical parameters from mocks generated with the anisotropic systematics model.

Parameter	Mock truth	Marginal $\pm 1\sigma$	MAP \pm PJ-HPD
S_8	0.756	$0.752^{+0.034}_{-0.029}$	$0.752^{+0.029}_{-0.031}$
σ_8	0.706	$0.720^{+0.093}_{-0.19}$	$0.613^{+0.096}_{-0.190}$
Ω_m	0.344	$0.36^{+0.10}_{-0.14}$	$0.451^{+0.168}_{-0.140}$
h_0	0.657	< 0.747	$0.672^{+0.023}_{-0.078}$
ω_c	0.292	$0.162^{+0.075}_{-0.048}$	$0.182^{+0.060}_{-0.057}$
ω_b	0.022	—	$0.022^{+0.002}_{-0.002}$
n_s	1.0	< 0.981	$0.911^{+0.071}_{-0.067}$
A_{IA}	0.396	$0.11^{+0.48}_{-0.31}$	$0.266^{+0.068}_{-0.322}$
A_{bary}	3.133	—	$2.351^{+0.114}_{-0.250}$

Notes. For each parameter, we show the underlying true value which was input into the mock data vector, and the inferred value recovered by KiDS-SBI. The third column shows the marginal means as well as the upper and lower 68% confidence intervals, i.e. 1σ , of the marginals. h_0 , n_s , ω_b , and A_{bary} are not well enough constrained in order to calculate a meaningful marginal parameter estimate. The fourth column shows the maximum a posteriori (MAP), and the uncertainties are defined as the upper and lower 68% confidence intervals, i.e. 1σ , given by the projected joint highest posterior density, PJ-HPD (Robert 2007; Joachimi et al. 2021). Note: cosmic shear is only expected to significantly constrain the value of S_8 and A_{IA} .

assume that the likelihood of a full-sky shear angular power spectrum, $C(\ell)$, at a given cosmology, is Gaussian, specifically $\mathcal{L} \propto \exp(-\frac{1}{2}\chi^2)$, where χ is the goodness of fit given by

$$\chi^{(pq)} = [\hat{C}^{(pq)}(\ell) - C_{\text{fid}}^{(pq)}(\ell)] / \sigma_{\text{fid}}^{(pq)}(\ell), \quad (\text{F.1})$$

where $\hat{C}^{(pq)}$ is the observed data vector, $C_{\text{fid}}^{(pq)}$ is the modelled data vector at the fiducial cosmology, and $\sigma_{\text{fid}}^{(pq)}$ is the standard deviation of the data at the fiducial cosmology. Assuming that the underlying shear field is described by a Gaussian random field, $\sigma_{\text{fid}}^{(pq)}$ is given by

$$(\sigma_{\text{fid}}^{(pq)})^2 = \frac{[(C_{\text{fid}}^{(pq)}(\ell)]^2 + C_{\text{fid}}^{(pp)}(\ell) C_{\text{fid}}^{(qq)}(\ell)}{2\ell + 1} + [N^{(pq)}(\ell)]^2, \quad (\text{F.2})$$

where the first term on the right-hand side defined the contribution due to cosmic variance, while $N^{(pq)}$ is the shape noise contribution for a given tomographic bin combination (Joachimi et al. 2011). To assess the impact of varying S_8 on such a likelihood distribution, we introduce a toy parameter, a , which linearly scales the model's fiducial data vector similar to S_8 as follows

$$C_{\text{fid},a}^{(pq)}(\ell) = a C_{\text{fid}}^{(pq)}(\ell). \quad (\text{F.3})$$

Hence, we can define $\chi_a^{(pq)}$ as the goodness-of-fit after rescaling the model by combining Eqs. (F.1), (F.2) and (F.3) as follows

$$\chi_a^{(pq)} = \frac{1 - a \pm \chi^{(pq)} \sqrt{\frac{1 + [R^{(pq)}(\ell)]^{-2}}{2\ell + 1}} + \left[\frac{N^{(pq)}(\ell)}{C_{\text{fid}}^{(pq)}(\ell)} \right]^2}{\sqrt{a^2 \frac{1 + [R^{(pq)}(\ell)]^{-2}}{2\ell + 1}} + \left[\frac{N^{(pq)}(\ell)}{C_{\text{fid}}^{(pq)}(\ell)} \right]^2}}, \quad (\text{F.4})$$




Fig. D.2. Posterior contours of the seven cosmological and astrophysical parameters which are varied given the anisotropic systematics model within KiDS-SBI over the prior space shown in Table 2 when analysing a mock data vector from the same model. The black solid lines indicate the true cosmology of the input mock data vector generated from the anisotropic systematics model while adding noise. All the aforementioned values are shown in Table E.1. These posteriors are obtained from training neural density estimators in DELFI (Alsing et al. 2019) on 18,000 realisations of the forward simulations assuming the anisotropic systematics model, in line with the choices made in L23. The posterior is obtained from the combined posteriors of six independent conditional Masked Autoregressive Flows (MAFs) each is made up of three to eight Masked Autoencoder for Density Estimations (MADEs) each with two hidden layers of 50 neurons.

where $R^{(pq)}(\ell) \equiv C_{\text{fid}}^{(pq)}(\ell) / \sqrt{C_{\text{fid}}^{(pp)}(\ell) C_{\text{fid}}^{(qq)}(\ell)}$. By calculating the distribution of $\exp(-\frac{1}{2}\chi_a^2)$ over χ for different values of ℓ and a , we can assess the impact of rescaling the shear angular power spectrum directly on the Gaussian likelihood across different angular scales, ℓ . To facilitate this, we can analytically compute the standard deviation of this distribution, σ_{χ_a} , as follows

$$\sigma_{\chi_a}^2 = 1 + (a^2 - 1) \left(1 + (2\ell + 1) \frac{[N^{(pq)}(\ell)/C_{\text{fid}}^{(pq)}(\ell)]^2}{1 + [R^{(pq)}(\ell)]^{-2}} \right)^{-1}. \quad (\text{F.5})$$

Applying Eq. (F.5) to the five KiDS-1000 tomographic bins, while assuming a shape noise contribution consistent with the values given in Table 1, we obtain Fig. F.1. We find that, as expected, at large angular scales ($\ell = 500$), the likelihood is dominated by the shape noise contribution, so its width does not scale with a under any circumstance. However, at $\ell = 70$

and $\ell = 100$, in some tomographic bin combinations where the signal-to-noise ratio is large enough (i.e. the higher redshift bins), the width of the likelihood distribution scales linearly with the scaling applied to the shear angular power spectrum, a . This implies that the uncertainty of the KiDS-1000 cosmic shear measurements is dependent on the amplitude of the measurement itself, and therefore the underlying cosmology. In fact, as shown in the upper right panel of Fig. F.1, the broadening of the analytical likelihood with a appears to be generally consistent with the widening observed in the likelihood marginals from the anisotropic systematics model in KiDS-SBI shown in Fig. 17.

Appendix G: Impact of the sampler on the Gaussian likelihood posteriors

To show that the Gaussian likelihood constraints shown in Fig. 18 and Fig. I.2 obtained from the nautilus sampler

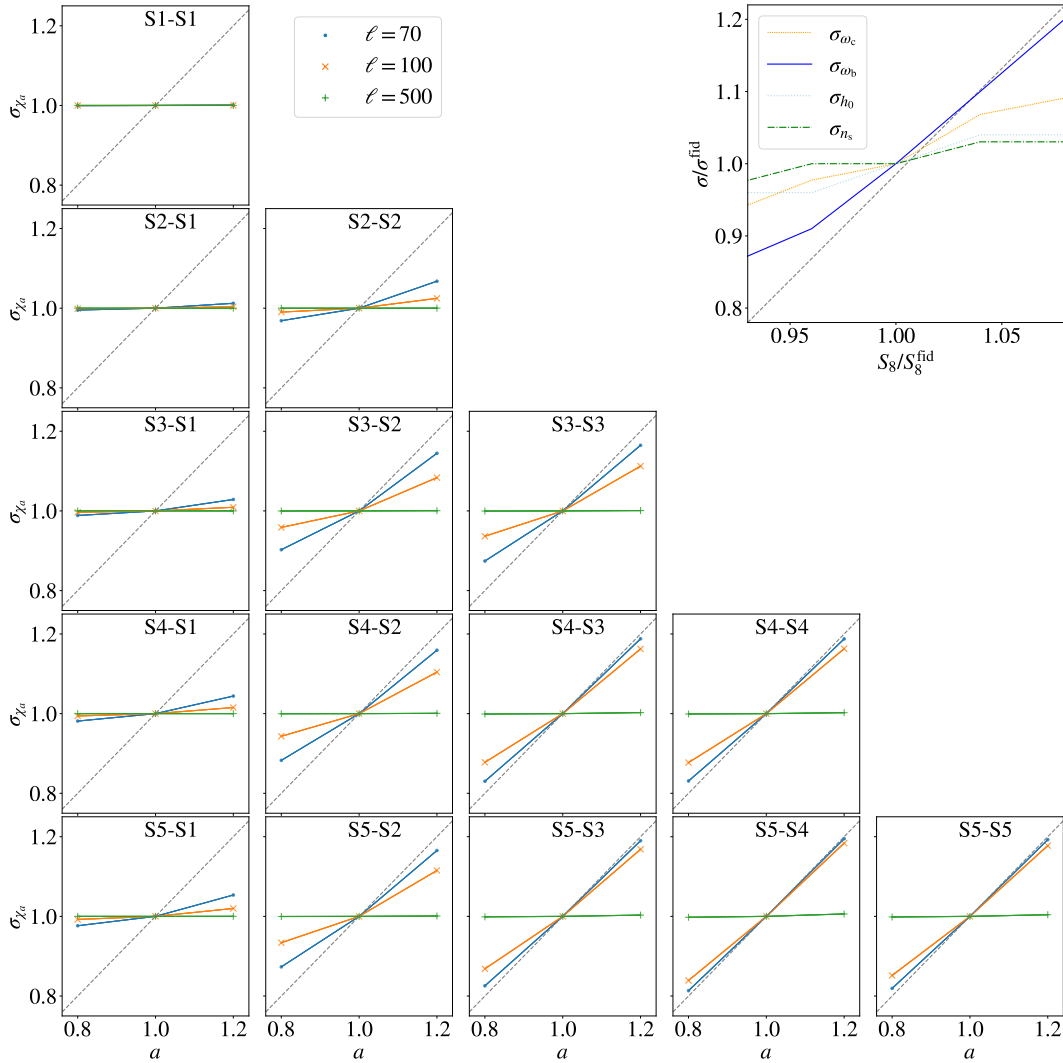


Fig. F.1. The standard deviation of the analytical likelihood distribution of χ_a (which is given by the goodness-of-fit of the cosmic shear two-point statistic signal scaled by a factor of a) as a function of the factor a . Each of the panels shows the effect on the shear signal of a different combination of the five KiDS-1000 tomographic bins (S1 to S5). The blue dots represent the σ_{χ_a} values at $\ell = 70$ which is just below the scale cuts applied in the anisotropic systematics model. The orange crosses are evaluated at $\ell = 100$, while the green plus signs assume $\ell = 500$, where the uncertainty is dominated by the shape noise. The panel in the upper right corner shows how the 1σ intervals of the likelihood marginals from KiDS-SBI vary with respect to the change in S_8 relative to $S_8^{\text{fid}} = 0.754$ as shown in Fig. 17. The grey dashed line in each panel shows a direct proportionality for reference.

(Lange 2023) are consistent with previous KiDS-1000 analyses (e.g. Asgari et al. 2021; Loureiro et al. 2022) which used the multinest (Feroz et al. 2009), we compare the posteriors obtained from both samplers for the same data vector and covariance matrix in Fig. G.1. As has been noted in previous work (Lemos et al. 2023b), we find that multinest consistently gives $\sim 15\%$ narrower S_8 posterior marginals than the ones sampled nautilus irrespective of the model assumed to obtain the sample covariance.

Appendix H: Posterior variations with noise realisations

As discussed in Sect. 6, different uncertainty realisations of the anisotropic systematics model in KiDS-SBI at the same cosmology produces a significant random scatter in the measured final pseudo-CIs. Figure H.1 shows how the choice of noise realisation impacts the inferred posteriors at a given cosmology. We find that overall the variations in the posteriors between uncertainty reali-

sations are small with the mean marginal of S_8 fluctuating by up to $\pm 0.5\sigma$ and the A_{IA} marginal can fluctuate up to $\pm 0.75\sigma$.

Appendix I: KiDS-SBI posteriors and parameter constraints

To expand on the results of the KiDS-SBI analysis of the KiDS-1000 gold sample presented in Fig. 18 and Table 3, we present the full posterior learnt in KiDS-SBI to fit the data in Fig. 1.1 in Fig. 1.2 for both the anisotropic systematics and standard isotropic systematics models. Table F.1 also provides the parameter best estimates for all cosmological and astrophysical parameters varied in the forward models. As shown in Fig. 1.2, both models leave A_{bary} , n_s and h_0 unconstrained, while constraining A_{IA} . The posteriors of both models strongly prefer values of A_{IA} greater than zero and less than 1.5. Despite the wide range of the A_{IA} priors, the posterior stays within a physical regime for the parameter which indicates a low chance of residual systematics strongly biasing the parameter constraints.

Table F.1. All inferred cosmological and astrophysical parameters varied within the anisotropic systematics and the standard isotropic systematics models from the KiDS-1000 gold sample.

	SBI: anisotropic sys.		SBI: std. isotropic sys.		MCMC: anisotropic sys.		MCMC: std. isotropic sys.	
	Marginal $\pm 1\sigma$	MAP \pm PJ-HPD	Marginal $\pm 1\sigma$	MAP \pm PJ-HPD	Marginal $\pm 1\sigma$	MAP \pm PJ-HPD	Marginal $\pm 1\sigma$	MAP \pm PJ-HPD
S_8	0.731 ± 0.033	$0.743^{+0.015}_{-0.051}$	$0.772^{+0.038}_{-0.032}$	$0.780^{+0.020}_{-0.048}$	$0.725^{+0.034}_{-0.023}$	$0.731^{+0.023}_{-0.033}$	$0.764^{+0.031}_{-0.025}$	0.770 ± 0.029
σ_8	$0.73^{+0.10}_{-0.21}$	$0.72^{+0.09}_{-0.20}$	$0.78^{+0.12}_{-0.23}$	$0.73^{+0.12}_{-0.20}$	$0.75^{+0.12}_{-0.18}$	$0.72^{+0.19}_{-0.11}$	$0.76^{+0.11}_{-0.14}$	$0.74^{+0.14}_{-0.11}$
Ω_m	$0.337^{+0.097}_{-0.150}$	$0.323^{+0.220}_{-0.060}$	$0.333^{+0.093}_{-0.160}$	$0.345^{+0.216}_{-0.083}$	$0.311^{+0.072}_{-0.130}$	$0.310^{+0.078}_{-0.108}$	$0.324^{+0.068}_{-0.110}$	$0.323^{+0.110}_{-0.086}$
h_0	$0.736^{+0.083}_{-0.044}$	$0.669^{+0.085}_{-0.026}$	$0.726^{+0.043}_{-0.084}$	$0.640^{+0.11}_{-0.001}$	—	$0.674^{+0.066}_{-0.005}$	—	$0.673^{+0.071}_{-0.022}$
ω_c	$0.158^{+0.084}_{-0.045}$	$0.125^{+0.095}_{-0.023}$	$0.151^{+0.061}_{-0.089}$	$0.122^{+0.091}_{-0.027}$	$0.143^{+0.040}_{-0.073}$	$0.118^{+0.069}_{-0.034}$	$0.147^{+0.039}_{-0.057}$	$0.123^{+0.077}_{-0.021}$
ω_b	$0.0224^{+0.0022}_{-0.0033}$	$0.0190^{+0.004}_{-0.0001}$	0.0225 ± 0.0020	$0.0190^{+0.0046}_{-0.0001}$	—	$0.0222^{+0.0016}_{-0.0012}$	—	0.0225 ± 0.0019
n_s	$0.987^{+0.110}_{-0.039}$	$0.972^{+0.128}_{-0.132}$	$0.961^{+0.054}_{-0.120}$	$0.949^{+0.151}_{-0.109}$	—	$0.988^{+0.045}_{-0.105}$	$0.958^{+0.042}_{-0.110}$	$0.956^{+0.046}_{-0.108}$
A_{IA}	$0.60^{+0.42}_{-0.32}$	$0.63^{+0.36}_{-0.40}$	$0.61^{+0.38}_{-0.30}$	$0.57^{+0.39}_{-0.29}$	$0.59^{+0.48}_{-0.36}$	$0.73^{+0.35}_{-0.45}$	$0.64^{+0.42}_{-0.34}$	$0.73^{+0.32}_{-0.43}$
A_{bary}	$2.60^{+0.52}_{-0.20}$	$2.66^{+0.47}_{-0.66}$	$2.50^{+0.23}_{-0.49}$	$2.00^{+1.13}_{-0.001}$	> 2.45	$3.02^{+0.09}_{-0.70}$	> 2.48	$3.13^{+0.001}_{-0.66}$

Notes. The columns marked with SBI show the constraint obtained from the KiDS-SBI analysis, while the columns marked with MCMC represent the parameter constraints obtained from a standard MCMC analysis which assumes a Gaussian likelihood with a sample covariance matrix derived from the respective forward model. The second, fourth, sixth and seventh columns show the mean marginals as well as the upper and lower 68% confidence intervals, i.e. 1σ , of the marginals. The third and fifth columns show the maximum a posteriori (MAP), and the uncertainties are defined as the upper and lower 68% confidence intervals, i.e. 1σ , given by the projected joint highest posterior density, PJ-HPD (Robert 2007; Joachimi et al. 2021).

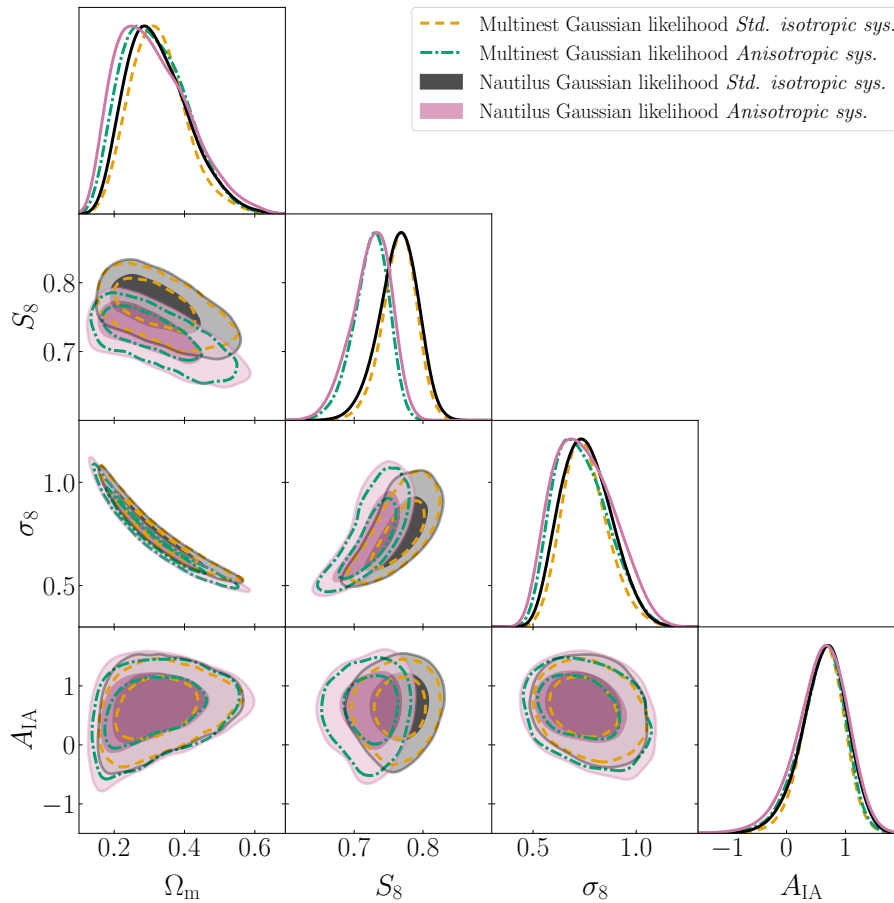


Fig. G.1. Comparison of posteriors obtained from a standard Gaussian likelihood analysis when using different MCMC samplers for the same noiseless mock data vector. In pink and in grey, we show the posteriors sampled with *nautilus* (Lange 2023) from a likelihood defined by a sample covariance based on mocks from the anisotropic systematics model and the standard isotropic systematics, respectively. In green and orange, we show posteriors sampled from the same two likelihoods as before, but sampling with the *multinest* sampler (Feroz et al. 2009) instead. All inference shown is done using the real KiDS data vector.




Fig. H.1. Posterior contours from five different mock data vectors given the anisotropic systematics model within KiDS-SBI over the prior space shown in Table 2. Each data realisation assumes the same input parameters (as given by the black line and shown in Table E.1) while varying the random seeds of the underlying matter fields, galaxy positions, intrinsic galaxy shapes, and shear biases. The posteriors shown in Fig. D.2, 14, 15, and 16 are based on the 3817 noise realisation.

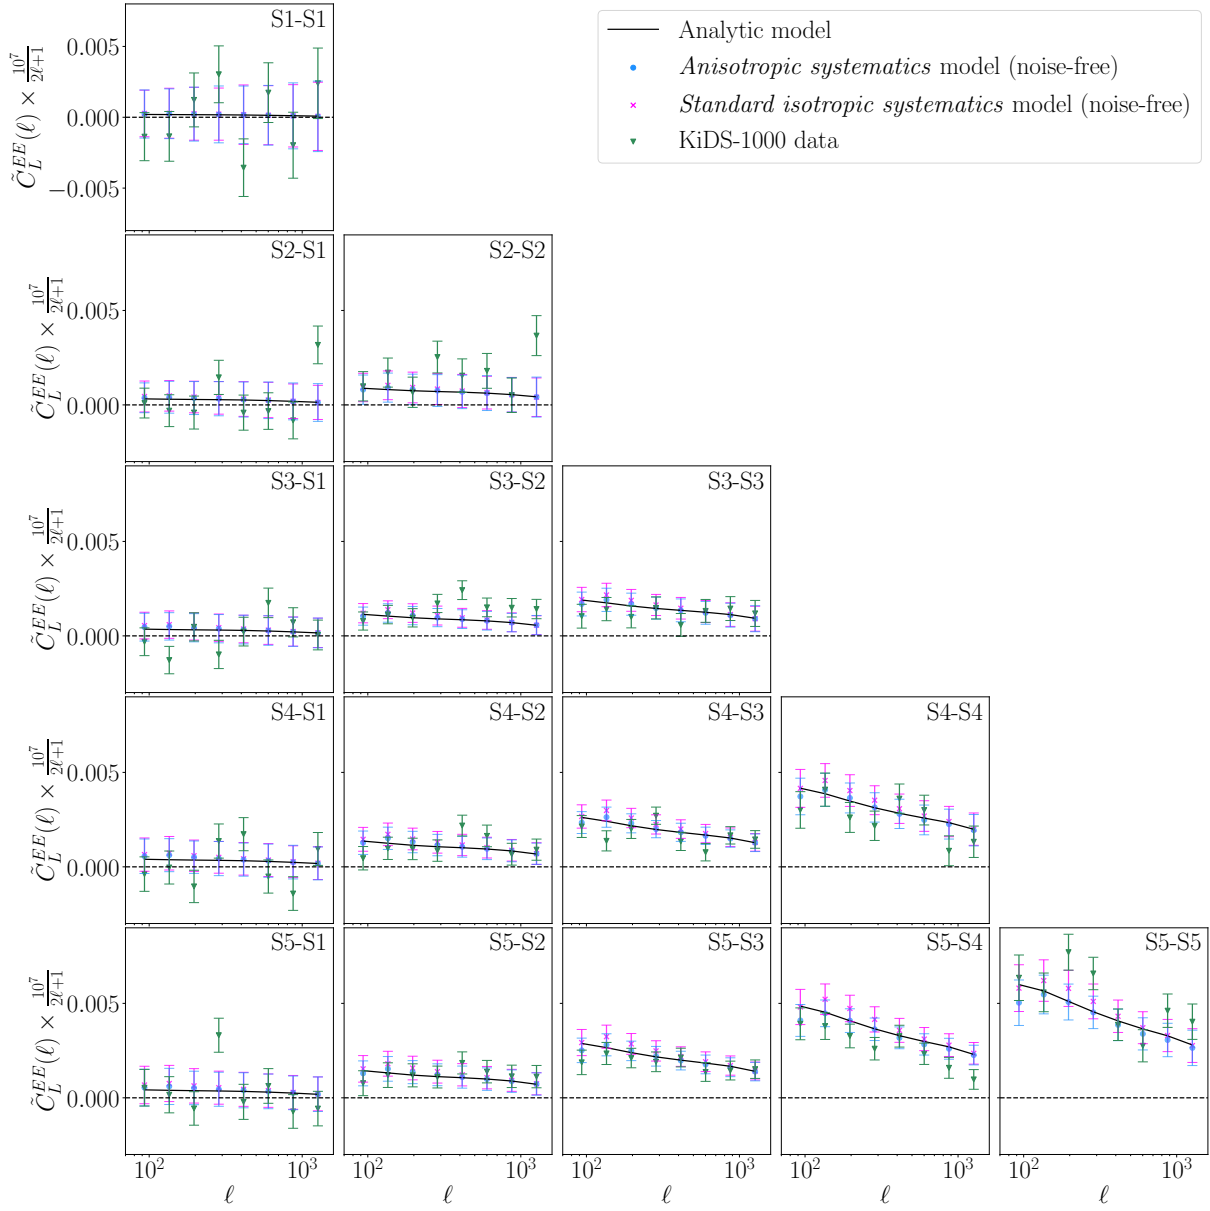


Fig. I.1. Measured pseudo-Cl's, $\tilde{C}_{\ell\ell,L}^{(pq)}(\Theta)$, for all combinations for five tomographic bins (S1 to S5). The green triangles show the pseudo-Cl's from the KiDS-1000 gold sample data. The blue points show the same for a noise-free realisation single realisation of the anisotropic systematics model in KiDS-SBI (see Sect. 5) at the maximum-a-posteriori, MAP (see Table F.1), while the magenta crosses show the measurements for noise-free realisation single realisation of the standard isotropic systematics model in KiDS-SBI (see Sect. 7.2) at the MAP for this model (see Table F.1). The uncertainties on the measurements are derived from the covariance matrix described in Sect. 5.3. The solid black line shows the pseudo-Cl's, $\tilde{C}_{\ell\ell}$, as derived from theory (see Appendix D for details on this) at the MAP from KiDS-SBI based on the anisotropic systematics model. The horizontal dashed line marks the line of zero signal.

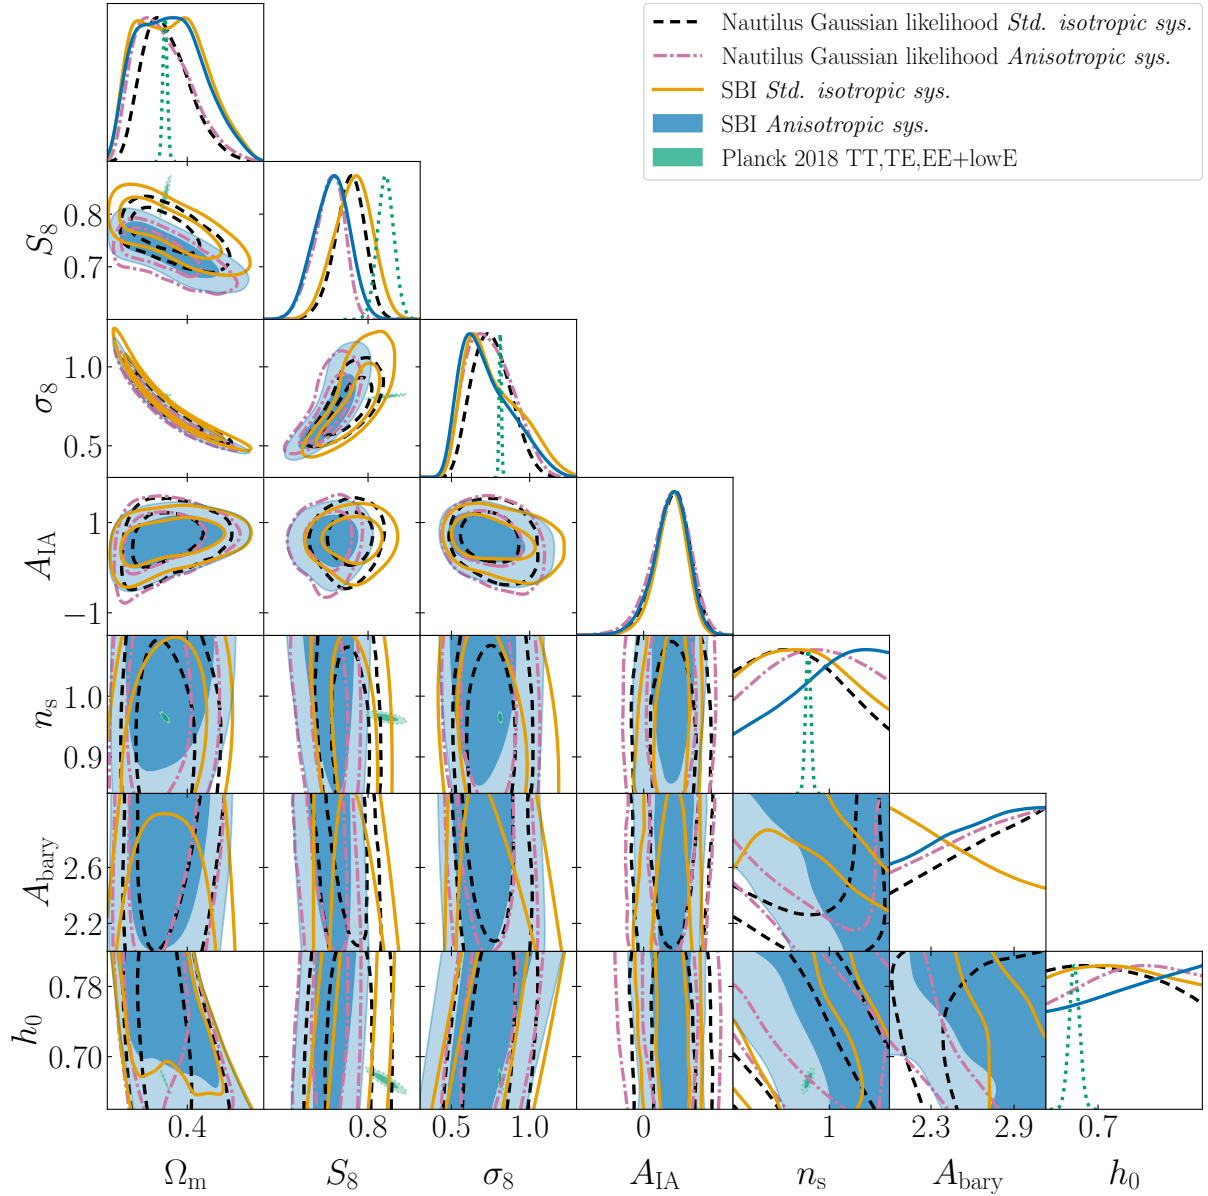


Fig. I.2. Posterior contours of the main constrained cosmological parameters from the KiDS-SBI analysis of the KiDS-1000 cosmic shear data assuming the anisotropic systematics model (in blue) compared against posterior contours from other analyses. In orange, we show the posterior from the same data while assuming the standard isotropic systematics model. In black and pink, we show the posteriors from the equivalent analyses done assuming a Gaussian likelihood for a given model. For clarity purposes, the δ_z parameters are marginalised out, while ω_b and ω_c are folded into Ω_m . The green contour shows the posterior from the cosmic microwave background constraints from the TT,TE,EE+lowE modes (Planck Collaboration VI 2020). We note that the Planck TT,TE,EE+lowE contours do not have any marginals in A_{IA} or A_{bary} as the CMB is not sensitive to the IAs of galaxies or baryonic feedback.