



Universiteit
Leiden

The Netherlands

Cosmic tomography with weak gravitational lensing

Li, S.S.

Citation

Li, S. S. (2023, October 25). *Cosmic tomography with weak gravitational lensing*. Retrieved from <https://hdl.handle.net/1887/3645974>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3645974>

Note: To cite this publication please use the final published version (if applicable).

CHAPTER 4

KiDS-1000: Cosmology with improved cosmic shear measurements

ABSTRACT

We present refined cosmological parameter constraints derived from a cosmic shear analysis of the fourth data release from the Kilo-Degree Survey (KiDS-1000). Our refinements are driven by enhanced galaxy shape measurements using an updated version of the *lensfit* code, and improved shear calibration achieved with a newly developed suite of multi-band image simulations. Additionally, we incorporate recent advancements in cosmological inference from the joint Dark Energy Survey Year 3 and KiDS-1000 cosmic shear analysis. Assuming a spatially flat standard cosmological model, we constrain $S_8 \equiv \sigma_8 (\Omega_m/0.3)^{0.5} = 0.776^{+0.029+0.002}_{-0.027-0.003}$, where the second set of uncertainties accounts for the systematic uncertainties within the shear calibration. These systematic uncertainties stem from minor deviations from realism in the image simulations and the sensitivity of the shear measurement algorithm to the morphology of the galaxy sample. Despite these changes, our results align with previous KiDS studies and other weak lensing surveys, and find a $\sim 2.3\sigma$ level of tension with the *Planck* cosmic microwave background constraints on S_8 .

S.-S. Li, H. Hoekstra, K. Kuijken et al.
Astronomy & Astrophysics, submitted

4.1 Introduction

Weak gravitational lensing by large-scale structure, also known as cosmic shear, is a powerful technique for studying the matter distribution in the Universe without assuming a specific correlation between dark and baryonic matter (e.g. Blandford et al. 1991; Miralda-Escude 1991; Kaiser 1992)¹. Owing to its remarkable potential in exploring the cosmic matter distribution, cosmic shear analysis gained popularity since its first detection over twenty years ago (Bacon et al. 2000; Kaiser et al. 2000; Van Waerbeke et al. 2000; Wittman et al. 2000). When distance information for source galaxies is also known, we can differentiate them along the line of sight and perform a tomographic analysis, which entails reconstructing the three-dimensional matter distribution from multiple two-dimensional projections. This tomographic cosmic shear analysis is especially effective for constraining dark energy properties, as it sheds light on the evolution of cosmic structures (e.g. Hu 1999; Huterer 2002).

Recent surveys, such as the Kilo-Degree Survey (KiDS, de Jong et al. 2013), the Dark Energy Survey (DES, Dark Energy Survey Collaboration et al. 2016), and the Hyper Suprime-Cam (HSC) survey (Aihara et al. 2018), primarily focus on constraining the amplitude of matter density fluctuations. Conventionally, this quantity is characterised by the parameter $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5}$, where Ω_m is the matter density parameter and σ_8 is the standard deviation of matter density fluctuations in spheres of radius $8h^{-1}$ Mpc, computed using linear theory, where the Hubble constant $H_0 = 100h$ km s⁻¹ Mpc⁻¹. Interestingly, the S_8 values derived from these weak lensing surveys are consistently lower than those predicted by cosmic microwave background (CMB) observations from the *Planck* satellite.

Specifically, the latest cosmic shear analyses from KiDS ($0.759^{+0.024}_{-0.021}$, Asgari et al. 2021, A21 hereafter), DES ($0.759^{+0.025}_{-0.023}$, Amon et al. 2022; Secco et al. 2022), and HSC ($0.769^{+0.031}_{-0.034}$, Li et al. 2023c; $0.776^{+0.032}_{-0.033}$, Dalal et al. 2023) provide S_8 values that are roughly 2σ lower than the *Planck* predictions (0.832 ± 0.013 , Planck Collaboration et al. 2020) based on the standard spatially flat Λ cold dark matter (Λ CDM) cosmological model. Most recently, a joint cosmic shear analysis of the DES Y3 and KiDS-1000 by the two survey teams (DES and KiDS Collaboration et al. 2023, DK23 hereafter) yields an S_8 constraint of $0.790^{+0.018}_{-0.014}$, which is closer to the *Planck* results, but still shows a level of 1.7σ difference. This mild difference in the S_8 constraints between the weak lensing surveys and CMB observations triggered extensive discussions from various perspectives, encompassing potential systematic errors in the data (e.g. Efstathiou & Lemos 2018; Köhlinger et al. 2019), the influence of the baryonic physics (e.g. Schneider et al. 2002; Amon & Efstathiou 2022; Preston et al. 2023), and a potential deviation from the standard Λ CDM model (see Perivolaropoulos & Skara 2022 for a recent review).

Here, we focus on the control of systematics in the cosmic shear analysis, particularly

¹However, with increasing precision in weak lensing observations, the impact of baryonic processes, such as radiative cooling and feedback from star formation and active galactic nuclei, on the observed matter distribution can no longer be ignored for small-scale structures (e.g. van Daalen et al. 2011; Semboloni et al. 2011).

those arising during the KiDS shear measurement process. Measuring lensing-induced shear from noisy pixelised galaxy images is a challenging task, complicated further by distortions caused by the point spread function (PSF) resulting from instrumental and observational conditions, as well as blending effects that arise when two or more objects are close on the sky (see Mandelbaum 2018 for a review). These factors can introduce significant measurement biases (e.g. Paulin-Henriksson et al. 2008; Melchior & Viola 2012; Refregier et al. 2012; Massey et al. 2013; Dawson et al. 2016; Euclid Collaboration et al. 2019) and alter the selection function of the source sample, leading to selection bias (e.g. Hartlap et al. 2011; Chang et al. 2013; Hoekstra et al. 2021). Therefore, obtaining unbiased shear measurements relies on careful calibration, which can be performed using either pixel-level image simulations (e.g. Miller et al. 2013; Hoekstra et al. 2015; Fenech Conti et al. 2017, FC17 hereafter; Samuroff et al. 2018; Mandelbaum et al. 2018) or the data themselves (e.g. Huff & Mandelbaum 2017; Sheldon & Huff 2017; Sheldon et al. 2020).

Additionally, in the case of large-area imaging surveys, determining the distance information for individual source galaxies depends on redshifts derived from broad-band photometric observations. These photometric redshift estimates, which are subject to significant uncertainty, require careful calibration using spectroscopic reference samples (e.g. Hoyle et al. 2018; Tanaka et al. 2018; Hildebrandt et al. 2021). Furthermore, recent studies showed that the blending of source images results in the coupling of shear and redshift biases (e.g. MacCrann et al. 2022; Li et al. 2023b, L23 hereafter). Consequently, a joint calibration of these two estimates becomes essential, necessitating the use of multi-band image simulations for future cosmic shear analyses.

In light of all these concerns, we implemented several improvements to the cosmic shear measurements in KiDS, as detailed in L23. We enhanced the accuracy of the galaxy shape measurements by using an upgraded version of the *lensfit* code (Miller et al. 2007, 2013; Kitching et al. 2008), complemented by an empirical correction scheme that reduces PSF contamination. More notably, in L23 we introduced SKiLLS (SURFS-based KiDS-Legacy-Like Simulations), a suite of multi-band image simulations that enables joint calibration of shear and redshift estimates. This is an important element for the forthcoming weak lensing analysis of the complete KiDS survey, known as the KiDS-Legacy analysis (Wright et al. in prep.).

In this chapter, we take an intermediate step towards the forthcoming KiDS-Legacy analysis by applying the improvements from L23 to a cosmic shear analysis based on the fourth data release of KiDS (KiDS-1000, A21). In contrast to previous KiDS cosmic shear analyses, which used shear calibration methods developed in FC17 and Kannawadi et al. (2019, K19 hereafter) based on single-band image simulations, the current analysis adopted SKiLLS, marking the first instance of multi-band image simulations being used for KiDS cosmic shear analysis². We also incorporated recent advancements in cosmological inference and updated the current cosmological parameter constraints from KiDS. In particular, we updated the code for the non-linear evolution of the matter

²K19 did attempt to assign photo- z estimates from data to simulations, but the actual photo- z measurements were not simulated.

power spectrum calculation from `HMCODE` to the latest `HMCODE-2020` version (Mead et al. 2021). We also investigated the impact of the intrinsic alignment model by incorporating amplitude priors inspired by Fortuna et al. (2021a).

The remainder of this chapter is structured as follows. In Sect. 4.2, we introduce and validate the updated KiDS shear catalogue, followed by the shear and redshift calibration in Sect. 4.3. We describe our cosmological inference method in Sect. 4.4 and present the results in Sect. 4.5. Finally, we summarise the results in Sect. 4.6.

4.2 Updated weak lensing shear catalogue

Our shear catalogue is based on the KiDS-ESO-DR4 data release (Kuijken et al. 2019), which combines optical observations in the *ugri* bands from KiDS using the ESO VLT Survey Telescope (de Jong et al. 2013) and near-infrared observations in the *ZYJHK_s* bands from the ESO VISTA Kilo-degree INfrared Galaxy (VIKING) survey using the VISTA telescope (Edge et al. 2013). The data set covers 1006 deg² survey tiles and includes nine-band photometry measured using the GAAP technique (Kuijken et al. 2015). The photometric redshifts (photo-*z*s) for individual source galaxies were estimated using the `BPZ` code (Benítez 2000). After masking, the effective area of the data set in the CCD pixel frame is 777.4 deg² (Giblin et al. 2021). To perform the cosmic shear analysis, we divided the source sample into five tomographic bins based on the `BPZ` estimates (z_B). The first four bins have a spacing of $\Delta z_B = 0.2$ in the range $0.1 < z_B \leq 0.9$, while the fifth bin covers the range $0.9 < z_B \leq 1.2$, following the previous KiDS cosmic shear analyses.

4.2.1 Galaxy shapes measured with the updated *lensfit*

When preparing the shear measurements for the upcoming data release of KiDS, we upgraded the *lensfit* code (Miller et al. 2007, 2013; Kitching et al. 2008) from version 309c to version 321 (see L23 for details). The latest version includes a correction to an anisotropic error in the original likelihood sampler, which previously caused a small yet noticeable residual bias that was not related to the PSF or underlying shear (Miller et al. 2013; Hildebrandt et al. 2016; Giblin et al. 2021). We used the new code to re-measure the galaxy shapes in the KiDS-ESO-DR4 data set, resulting in a new shear catalogue. Throughout the paper, we refer to the new shear catalogue as KiDS-1000-v2 to distinguish it from the previous KiDS-1000(-v1) shear catalogue.

The raw measurements from the *lensfit* code suffer from biases primarily due to the PSF anisotropy, but also because of the object selection and weighting scheme. To address these biases, FC17 introduced an empirical correction scheme to isotropise the original measurement weights, which was used in previous KiDS studies (see also K19). However, this approach is insufficient for the current version of the *lensfit* code. Furthermore, L23 found that the method was susceptible to variations in the sample size, posing challenges for consistent application to both data and simulations.

Therefore, a new correction scheme was introduced by L23 that modifies both the

measured ellipticities and weights to ensure the average PSF leakage, defined as the fraction of the PSF ellipticity leaking into the shear estimator, is negligible in each tomographic bin. For further details, we direct readers to L23. In summary, the new correction scheme first isotropises the measurement weights, then adjusts the measured ellipticities to eliminate any remaining noise bias and selection effects. We note that this correction scheme is not designed to refine the shape measurements of individual galaxies; rather, it aims to ensure that the collectively weighted shear signal is robust against PSF leakage. In this paper, we applied this newly developed empirical correction to the KiDS-1000-v2 shear catalogue.

4.2.2 Validation of the shear estimates

In order to use the weak lensing shear catalogue for cosmological inference, it is crucial to first verify the accuracy of the shear estimation and ensure that the residual contamination from systematic effects is within the acceptable level for scientific analysis. To achieve this, Giblin et al. (2021) proposed a series of null-tests to assess the robustness of the KiDS-1000-v1 shear catalogue. With the updated galaxy shape measurements in the KiDS-1000-v2 catalogue, it is necessary to repeat some of these tests to confirm the reliability of the new catalogue.

As the KiDS-1000-v2 catalogue updates only the galaxy shape measurements while maintaining the established photometry and PSF models, we did not repeat tests related to photometry and PSF modelling. We started by examining the PSF leakage in the weighted *lensfit* shear estimator, using the first-order systematics model proposed by Heymans et al. (2006). This model takes the form (Giblin et al. 2021)

$$\epsilon_k^{\text{obs}} = (1 + m_k)(\epsilon_k^{\text{int}} + \gamma_k) + \alpha_k \epsilon_k^{\text{PSF}} + c_k, \quad [k = 1, 2], \quad (4.1)$$

where ϵ^{obs} denotes the measured galaxy ellipticity, m is the multiplicative shear bias³, ϵ^{int} refers to the intrinsic galaxy ellipticity, γ stands for the cosmic shear signal (which is the parameter of interest), α is the PSF leakage factor, and c is an additive term comprising residual biases unrelated to the PSF or underlying shear. The subscript $k = 1, 2$ denotes the two ellipticity components. We note that we did not include PSF modelling errors in Eq. (4.1), as we used the same PSF model as Giblin et al. (2021), who had already confirmed its accuracy. Assuming that $(\epsilon_k^{\text{int}} + \gamma_k)$ averages to zero for a large galaxy sample (a property validated with the KiDS data; see, for example, Sect. 3 in Giblin et al. 2021), we can determine the α and c parameters from the data using a simple linear regression method.

Figure 4.1 presents the measured PSF leakage α and the additive term c for the KiDS-1000-v2 catalogue, alongside the measurements from the KiDS-1000-v1 catalogues for comparison. As expected, the KiDS-1000-v2 catalogue exhibits a mean α -term consistent with zero for all redshift bins, owing to the empirical correction scheme

³Throughout this paper, we interchangeably use ‘multiplicative bias’ and ‘shear bias’, as our simulation-based shear calibration only addresses this parameter. Conversely, PSF leakage and the additive term are empirically corrected.




Figure 4.1: PSF contamination α (top panels) and additive term c (bottom panels) as a function of tomographic bin labelled by the central z_B value. The measurements are obtained from a weighted linear fitting using Eq. (4.1). The red points represent measurements from the KiDS-1000-v2 catalogue, while the grey points show the measurements from the KiDS-1000-v1 catalogue. The red/grey coloured bars correspond to results from the entire sample without tomographic binning.

4

outlined in Sect. 4.2.1 (see also Sect. 4 in L23). The upgraded *lensfit* code has reduced the overall c_2 -term by half, reaching a level of $c_2 \sim (3 \pm 1) \times 10^{-4}$ for the entire sample. However, despite this improvement, the c term has not been eliminated, particularly in distant tomographic bins where a small but noticeable c term still persists, which was not seen in the simulations.

To correct for these residual small additive c -terms, we used the same empirical correction method as in previous KiDS analyses. Specifically, we subtracted the weighted average ellipticity from the observed ellipticity for each redshift bin as $\epsilon_{\text{corr}}^{\text{obs}} = \epsilon^{\text{obs}} - \overline{\epsilon^{\text{obs}}}$. Nevertheless, we caution that subtracting the mean c -term does not guarantee the removal of all additive biases, especially when detector-level effects, such as ‘charge transfer inefficiency’ (e.g. Rhodes et al. 2007; Massey 2010) and ‘pixel bounce’ (e.g. Toyozumi & Ashley 2005), can introduce position-dependent bias patterns. Although we have detected such effects in KiDS data (Hildebrandt et al. 2020; Giblin et al. 2021), their level does not affect the current cosmic shear analysis. More specifically, Asgari et al. (2019) demonstrated that even if current detector-level effects were increased by a factor of ten, they would not cause significant bias for KiDS-like analyses.

The cosmic shear signal is conventionally measured using the two-point shear




Figure 4.2: Ratio of the PSF contamination ξ_+^{sys} , computed using Eq. (4.3), to the predicted amplitude of the cosmic shear signal $\xi_+^{\Lambda\text{CDM}}$ across all 15 tomographic bin combinations. The red lines represent results derived from the KiDS-1000-v2 catalogue, while the grey lines correspond to those obtained from the KiDS-1000-v1 catalogue. The blue shaded regions denote $\pm 10\%$ of the standard deviation of the measured cosmic shear signal, determined from the covariance matrix based on the statistics from the KiDS-1000-v2 catalogue.

correlation function, defined as⁴

$$\xi_{\pm}^{ij}(\theta) = \frac{\sum_{ab} w_a w_b \left[\epsilon_t^i(\mathbf{x}_a) \epsilon_t^j(\mathbf{y}_b) \pm \epsilon_{\times}^i(\mathbf{x}_a) \epsilon_{\times}^j(\mathbf{y}_b) \right]}{\sum_{ab} w_a w_b}, \quad (4.2)$$

where θ represents the separation angle between a pair of galaxies (a, b), the tangential and cross ellipticities $\epsilon_{t,\times}$ are computed with respect to the vector $\mathbf{x}_a - \mathbf{y}_b$ that connects the galaxy pair, and the associated measurement weight is denoted by w . Therefore, it is crucial to examine the systematics in the two-point statistics. Following the method of Bacon et al. (2003), we estimate the PSF leakage into the two-point correlation function measurement using

$$\xi_{\pm}^{\text{sys}} = \frac{\langle \epsilon^{\text{obs}} \epsilon^{\text{PSF}} \rangle^2}{\langle \epsilon^{\text{PSF}} \epsilon^{\text{PSF}} \rangle}, \quad (4.3)$$

where the $\langle \cdot \rangle$ represents the correlation function.

In Fig. 4.2, we present the ratio of the measured ξ_{+}^{sys} to the theoretical predictions of the cosmic shear signal. The blue shaded region denotes $\pm 10\%$ of the standard deviation of the cosmic shear signal, extracted from the analytical covariance. This covariance is calculated using an independent implementation of the methodology of Joachimi et al. (2021), and it incorporates the sample statistics of the updated catalogue. We compare the results from the KiDS-1000-v2 catalogue with those from the KiDS-1000-v1 catalogue. We observe general improvements, particularly in the high-redshift bins, where the PSF contamination is now negligible. The only exceptions are found in some large-scale bins ($\theta > 60$ arcmin), where the expected fiducial cosmic shear signal is relatively small and overwhelmed by high statistical noise.

To leading order, the weak lensing effect introduces only curl-free gradient distortions (E -mode signal), which makes the curl distortions (B -mode signal) a useful null-test for residual systematics in the shear measurement⁵. Following the convention of KiDS (Hildebrandt et al. 2017; Giblin et al. 2021), we use the complete orthogonal sets of E/B-integrals (COSEBIs, Schneider et al. 2010) to measure the B -mode signal. The COSEBIs provide an optimal E/B separation by combining different angular scales from the $\xi_{\pm}$ measurements.

Figure 4.3 presents the measured B -mode signals for all combinations of tomographic bins in our analysis, alongside the B -mode measurements from the KiDS-1000-v1 catalogue for comparison. To enable direct comparison, we used the same scale range of $(0.5, 300')$ as in Giblin et al. (2021) for calculating the COSEBIs B -mode⁶. Assuming

⁴In this study, all measurements of the two-point shear correlation function are conducted using the TREECORR code (Jarvis et al. 2004; Jarvis 2015).

⁵Some higher-order effects from lensing, such as source redshift clustering (e.g. Schneider et al. 2002), and intrinsic alignment of nearby galaxies (e.g. Troxel & Ishak 2015; Joachimi et al. 2015) can also introduce B -mode signals. However, these contributions are expected to be negligible for current weak lensing surveys (e.g. Hilbert et al. 2009).

⁶We also evaluated an alternate scale range of $(2', 300')$, consistent with our fiducial cosmic shear analysis. As anticipated, the B -mode signal was more negligible in this scenario due to reduced small-scale contamination.




Figure 4.3: Measurements of the B -mode signals using COSEBIs for the KiDS-1000-v2 catalogue (red points) compared to the KiDS-1000-v1 catalogue (grey points). The error bars originate from the diagonal of an analytical covariance matrix, accounting solely for measurement noise. For the KiDS-1000-v2 catalogue, we re-calculated the covariance using the method introduced by Joachimi et al. (2021), incorporating the updated statistics. The p -values for the KiDS-1000-v2 catalogue, shown in the top right corner of each panel, were calculated with 20 degrees of freedom, which corresponds to the number of modes used in each correlation.

a null signal, we computed the p -value for each B -mode measurement, setting the degrees of freedom equal to the number of modes in each measurement ($n = 20$). The covariance matrix, accounting only for shot noise, was estimated using an analytical model from Joachimi et al. (2021) applied to the updated catalogue. It is noteworthy that our covariance matrix differs from the one used in Giblin et al. (2021). This is due to the changes in sample statistics resulting from the updated shape measurement code and redshift calibration relative to the KiDS-1000-v1 catalogue used in Giblin et al. (2021). Most diagonal entries in our matrix show reduced uncertainties, ranging from a level of per cent to ten per cent. Therefore, if the absolute systematic levels are comparable between the two catalogues, our test would likely show a slight increase in the final p -values compared to those in Giblin et al. (2021). As indicated in the top right corner of each panel, the estimated p -values suggest that the measured B -mode signals align with a null signal across all bin combinations. The lowest p -value, $p = 0.02$, was found in the cross-correlation between the first and third tomographic bins.

After conducting all these tests, we can conclude that the KiDS-1000-v2 catalogue has reduced systematics when compared to the results from the KiDS-1000-v1 catalogue. These improvements are largely attributed to the updated version of the *lensfit* code, as well as the implementation of a new empirical correction scheme for PSF contamination. These results give us the confidence to use the updated catalogue for cosmological inference.

4.3 Shear and redshift calibration

The main improvement in our calibration comes from the use of SKiLLS multi-band image simulations, as developed in L23. These simulations fuse cosmological simulations with high-quality observational data to create mock galaxies with photometric and morphological properties closely resembling real-world galaxies. The observational data used by SKiLLS, drawn from the catalogue of Griffith et al. (2012), is identical to that used in K19. In L23, we developed a vine-copula-based algorithm that learns the measured morphological parameters from this catalogue and assigns them to the SURFS-Shark mock galaxies (Elahi et al. 2018; Lagos et al. 2018). We verified that the learning procedure maintains the observed multi-dimensional correlations between morphological parameters, magnitude, and redshifts. Nevertheless, both the observed catalogue from Griffith et al. (2012) and the learning algorithm possess inherent limitations, resulting in unavoidable uncertainties in our simulation input catalogue. These uncertainties are addressed in our shear calibration in Sect. 4.3.2.

To create KiDS+VIKING-like nine-band images, SKiLLS replicated the instrumental and observational conditions of 108 representative tiles selected from six sky pointings evenly distributed across the KiDS-DR4 footprint. The star catalogue was generated for each sky pointing using the TRILEGAL population synthesis code (Girardi et al. 2005) to account for the variation in stellar densities across the footprint. For the primary *r*-band images, on which the galaxy shapes were measured, SKiLLS included the correlated pixel noise introduced by the stacking process and the PSF variation between CCD images.

On the data processing side, SKiLLS followed the entire KiDS procedure, including object detection, PSF homogenisation, forced multi-band photometry, photo-*z* estimation, and shape measurements. The end result is a self-consistent joint shear-redshift mock catalogue that matches KiDS observations in both shear and redshift estimates. By taking this end-to-end approach, we accounted for photo-*z*-related selection effects in our shear bias estimation and enabled redshift calibration using the same mock catalogue. While our current analysis focuses on the improvement in shear calibration, it represents an intermediate step towards the KiDS-Legacy analysis, which will implement joint shear and redshift calibrations facilitated by the SKiLLS mock catalogue.

4.3.1 Calibration

To correct for shear bias in our measurements, we followed the method used in previous KiDS studies (FC17, K19). We applied an average shear bias correction factor, denoted as $\overline{m}^i$, to each tomographic bin i . This factor was calculated by averaging the individual m values of all sources within the bin, with each individual m value obtained using Eq. (4.1). In order to better align the simulations with the target data, we adhered to KiDS conventions by re-weighting the simulation estimates using the *lensfit* reported model signal-to-noise ratio and resolution, which is defined as the ratio of the PSF size to the measured galaxy size. More information on the re-weighting procedure can be found in Sec. 5.1 of L23.

Although the averaging method addresses the noise in individual source’s m estimation, it does not account for correlations involving shear bias. Thus, we have $\langle [1 + m^i(\theta')][1 + m^j(\theta' + \theta)] \rangle = (1 + \overline{m}^i)(1 + \overline{m}^j)$, with θ and θ' representing different separation angles between galaxy pairs. To test this assumption, we directly measured $\langle [1 + m^i(\theta')][1 + m^j(\theta' + \theta)] \rangle$ from image simulations and compared it to $(1 + \overline{m}^i)(1 + \overline{m}^j)$. Further details on this test can be found in Appendix 4.A. In summary, we found a negligible difference between the two estimators, a result that falls well within the current KiDS requirements. This validates the assumption for the KiDS analysis.

Given that the updated galaxy shape measurements also lead to changes in the sample selection function, it is necessary to repeat the redshift calibration for the KiDS-1000-v2 catalogue, even though our primary focus is to improve shear calibration. To quantify the changes in galaxy samples introduced by the modifications in shape measurements from the KiDS-1000-v1 to KiDS-1000-v2 catalogues, we compared their effective number densities before applying any redshift calibration. The observed percentage differences in each tomographic bin, from low to high redshift bins, are -1.8% , -0.4% , 0.2% , 1.3% , and 3.2% . Here, negative values indicate a decrease in density from the v1 to the v2 catalogue, while positive values signify an increase. These differences are largely attributed to changes in the weighting scheme from *lensfit* version 309c to version 321, as well as the implementation of the new empirical correction scheme for PSF leakage, as discussed in Sect. 4.3 and in L23. For this, we employed a methodology identical to the one used by Wright et al. (2020b), Hildebrandt et al. (2021) and van den Busch et al. (2022) (vdB22 hereafter). It is based on a direct calibration method (Lima et al. 2008) implemented with a self-organising map (SOM, Kohonen 1982; Masters et al. 2015). More information on our implementation is provided in Appendix 4.B, while Wright et al. (2020b), Hildebrandt et al. (2021) and vdB22 offer more comprehensive discussions.

The SOM-based redshift calibration method uses a ‘gold selection’ criterion to filter out sources that are not represented in the spectroscopic reference sample (see Appendix 4.B). However, this process influences shear biases as it alters the selection function of the final sample. To ensure a consistent estimation of shear biases, we created the SKiLLS-gold catalogue by mimicking this quality control on the SKiLLS mock catalogue, using the same SOM trained by the spectroscopic reference sample as the real

Table 4.1: Data properties for the KiDS-1000-v2 catalogue.

Bin	$n_{\text{eff}} [\text{arcmin}^{-2}]$	$\sigma_{\epsilon,i}$	$\delta z = z_{\text{est}} - z_{\text{true}}$	m_{raw}	m_{final}	σ_m
1	0.68	0.27	0.000 ± 0.0096	-0.023	-0.021	0.019
2	1.30	0.26	0.002 ± 0.0114	-0.025	-0.023	0.008
3	1.97	0.28	0.013 ± 0.0116	-0.013	-0.015	0.007
4	1.39	0.27	0.011 ± 0.0084	0.018	0.015	0.006
5	1.35	0.29	-0.006 ± 0.0097	0.032	0.031	0.006

Comparable summary statistics for the KiDS-1000-v1 catalogue can be found in Table 1 of A21. We note that the differences in summary statistics between our work and A21 stem from both the updated *lensfit* code and the enhanced redshift calibration outlined in vdB22. The effective number density n_{eff} and the ellipticity dispersion per ellipticity component $\sigma_{\epsilon,i}$ are calculated using the formulae provided in Appendix C of Joachimi et al. (2021). The n_{eff} values in this table are derived from an effective area of 777.4 square degrees in the CCD pixel frame, making them directly comparable to the values in Table 1 of A21. The correlated Gaussian redshift priors are based on the differences between the estimated and true redshifts, $\delta z = z_{\text{est}} - z_{\text{true}}$, as reported in vdB22. The priors are denoted as $\mu_i \pm \sigma_i$, where μ_i represents the mean shift and σ_i corresponds to the square root of the covariance matrix's diagonal elements. The m_{raw} results are derived from idealised constant shear simulations, while the m_{final} results, our fiducial outcomes, include corrections for the shear-interplay effect and PSF modelling bias. Statistical uncertainties, determined by the simulation volume, are directly computed from the fiducial simulations and denoted as σ_m .

data. We derived the appropriate shear bias correction factors from this SKiLLS-gold catalogue for individual tomographic bins, and present these values in Table 4.1. It is worth noting that the shear bias estimates presented in this work differ slightly from those in L23, which did not include the gold selection procedure. Despite this, the differences in the estimated shear biases are relatively minor across all tomographic bins, with the first tomographic bin showing the most noticeable change of 0.008.

Our fiducial results, m_{final} , account for the impact of PSF modelling uncertainties and the ‘shear interplay’ effect, which occurs when galaxies from different redshifts are blended together. For more details on these effects, we refer to L23 and MacCrann et al. (2022). Additionally, we provide the idealised m_{raw} results, which do not consider these higher-order effects. By comparing the cosmological constraints obtained from these two cases, we aim to evaluate the robustness of previous KiDS results with respect to these higher-order effects, which were not taken into account in the earlier shear calibration (FC17, K19).

4.3.2 Calibration uncertainties

Systematic uncertainties arising from redshift and shear calibrations can propagate into cosmological analyses, potentially leading to biased results. Therefore, it is crucial to adequately address these uncertainties in the analysis. In this section, we outline our approach to managing these calibration uncertainties.

The uncertainties in redshift calibration were addressed by introducing an offset parameter for the estimated mean redshift of galaxies in each tomographic bin. This offset parameter, described as correlated Gaussian priors, serves as a first-order correction to both the statistical and systematic uncertainties associated with redshift calibration. Table 4.1 lists the exact values for these parameters, which we obtained from vdB22 and Hildebrandt et al. (2021). They determined these prior values using spectroscopic and KiDS-like mock data generated by van den Busch et al. (2020). We consider the current priors to be conservative enough to account for any potential changes in the redshift biases from KiDS-1000-v1 to KiDS-1000-v2, given that both catalogues use the same photometric estimates. However, for the forthcoming KiDS-Legacy analysis, we plan to re-estimate these values based on the new SKiLLS mock data.

We improved our approach to handling uncertainties related to the shear calibration. In L23, nominal uncertainties were proposed for each tomographic bin based on sensitivity analyses. This aimed to ensure the robustness of the shear calibration within the specified uncertainties, but at the cost of reducing statistical power. In this work, we aim to improve this approach by separately accounting for the statistical and systematic uncertainties within the shear calibration.

The statistical uncertainties, as presented in Table 4.1, are computed directly from simulations and are limited only by the volume of the simulations, which can be increased with more computing resources⁷. These uncertainties are also easily propagated into the covariance matrix for cosmological inference. Although increasing the simulation volume could, in principle, reduce these uncertainties, we found that the current values already comfortably meet the KiDS requirements, thus further efforts in this direction were considered not necessary.

If the SKiLLS simulations perfectly match KiDS data, these statistical uncertainties would be the only contribution to the final uncertainty from the shear calibration. However, since our simulations are not a perfect replica of the real observations, residual shear biases may still be present in the data even after calibration. These biases, referred to as systematic uncertainties, are typically the primary source of error in shear calibration. Increasing the simulation volume cannot improve these uncertainties as they are determined by the realism of the image simulations. The level of these uncertainties can only be roughly estimated through sensitivity analyses.

Since the systematic residual shear biases directly scale the data vector, accurately quantifying their impact using the covariance matrix is challenging. Therefore, we use a forward modelling approach to capture the impact of these systematic uncertainties. Instead of incorporating these uncertainties into the covariance matrix, we examine how

⁷However, the finite volume of the input galaxy sample prevents an indefinite increase.

the final estimates of the cosmological parameters change due to the shift in signals caused by the systematic residual shear biases. This forward modelling approach can be easily implemented using simple optimisation algorithms since the shift is small, and the covariance remains unchanged. More details on how to determine residual shear biases and implement the forward modelling approach are provided in Appendix 4.C.

4.4 Cosmological inference

The cosmological inference in this study largely aligns with the approach used in the KiDS-1000-v1 analyses (A21; vdB22), with minor modifications primarily influenced by the recent joint DES Y3+KiDS-1000 cosmic shear analysis (DK23). In this section, we outline the configurations and reasoning behind these choices in our fiducial analysis. For certain notable changes, we also conduct extended analysis runs with different configurations to evaluate the impact of these modifications. Our analysis code is publicly accessible⁸.

Our code builds upon the `Cat_to_Obs_K1000_P1`⁹ and the KiDS Cosmology Analysis Pipeline (KCAP)¹⁰ infrastructure, as developed in Giblin et al. (2021), Joachimi et al. (2021), A21, Heymans et al. (2021) and Tröster et al. (2021). The `Cat_to_Obs_K1000_P1` pipeline converts KiDS shear and redshift measurements into various second-order statistics, with the assistance of the `TREECORR` code (Jarvis et al. 2004; Jarvis 2015). Meanwhile, KCAP estimates cosmological parameters using the CosmoSIS framework, which bridges the likelihood function calculation pipelines with MCMC samplers (Zuntz et al. 2015).

We measure the shear field using Complete Orthogonal Sets of E/B-Integrals (COSEBIs, Schneider et al. 2010). As reported by Asgari et al. (2020), COSEBIs offer enhanced robustness against small-scale effects on the shear power spectrum, which primarily stem from complex baryon feedback. Furthermore, we account for baryon feedback when modelling the matter-matter power spectrum, using `HMCODE-2020` (Mead et al. 2021) within the `CAMB` framework with the version 1.4.0 (Lewis et al. 2000; Howlett et al. 2012).

`HMCODE-2020`, an updated version of `HMCODE` (Mead et al. 2015, 2016), models the non-linear matter-matter power spectrum, incorporating the influence of baryon feedback through an enhanced halo-model formalism. This updated model is empirically calibrated using hydrodynamical simulations, following a more physically informed approach. Unlike its predecessor calibrated with OWLS hydrodynamical simulations (van Daalen et al. 2011), this newer version uses the updated BAHAMAS hydrodynamical simulations for calibration (McCarthy et al. 2017). These simulations, in turn, are calibrated to reproduce the observed galaxy stellar mass function and the hot gas mass fractions of groups and clusters. This calibration ensures that the simulation accurately reflects the impact of feedback on the overall distribution of matter (refer to McCarthy et al. 2017 for

⁸<https://github.com/lshuns/CSK1000LF321>

⁹https://github.com/KiDS-WL/Cat_to_Obs_K1000_P1

¹⁰<https://github.com/KiDS-WL/kcap>

further details). Furthermore, `HMCODE-2020` improves the modelling of baryon-acoustic oscillation damping and massive neutrino treatment, achieving an improved accuracy of 2.5% (compared to the previous version's 5%) for scales $k < 10h \text{ Mpc}^{-1}$ and redshifts $z < 2$ (Mead et al. 2021).

The model incorporates a single-parameter variant, T_{AGN} , representing the heating temperature of active galactic nuclei (AGN). Higher T_{AGN} values correspond to more intense AGN feedback, leading to a lower observed matter power spectrum. Following DK23, we use a uniform prior on $\log_{10}(T_{\text{AGN}})$ ranging from 7.3 to 8.0. This choice is motivated by the findings from the BAHAMAS hydrodynamical simulations (McCarthy et al. 2017; van Daalen et al. 2020).

Given the characteristics of COSEBIs and the implementation of the `HMCODE`, the KiDS-1000-v1 analyses included small-scale measurements down to $\theta_{\text{min}} = 0.5$. This strategy was, however, re-evaluated in DK23, which suggested more stringent scale cuts for the KiDS COSEBIs data vector, determined by the baryon feedback mitigation strategy proposed by Krause et al. (2021). Following this recommendation, we apply a scale cut of $\theta_{\text{min}} = 2'$ in our fiducial analysis.

We use the non-linear linear alignment (NLA) model to describe the intrinsic alignment (IA) of galaxies. This model combines the linear alignment model with a non-linear power spectrum and contains a single free parameter A_{IA} to describe the amplitude of IA signals (Hirata & Seljak 2004; Bridle & King 2007). It is also common to include a power law, with an index denoted as η_{IA} , to capture potential redshift evolution of the IA strength. To distinguish it from the redshift-independent NLA model, we refer to this variant as the NLA- z model.

In line with previous KiDS analyses, we take the redshift-independent NLA model as our fiducial choice since introducing η_{IA} has a minimal effect on the primary S_8 constraint (A21), and current direct observations of IA signals show little evidence of substantial redshift evolution (e.g., Joachimi et al. 2011; Singh et al. 2015; Johnston et al. 2019; Fortuna et al. 2021b; Samuroff et al. 2022). However, Fortuna et al. (2021a) suggests that the selection of galaxy samples resulting from the redshift binning may introduce a detectable redshift variation in the IA signal, although its impact remains negligible for current weak lensing analyses. To assess the impact of η_{IA} on our results, we perform an extended run using the NLA- z model, following the same prior selection as in DK23.

The KiDS-1000-v1 analyses adopted a broad and uninformative prior for A_{IA} , ranging from $[-6, 6]$, considering that the data can constrain it and that an incorrect informative prior could bias the final cosmological results. Although uncertainties regarding IA signals remain large, recent developments in the field have improved our knowledge of the expected IA signal strength. For instance, Fortuna et al. (2021a) used a halo model formalism, incorporating results from the latest direct IA measurements, and predicted $A_{\text{IA}} = 0.44 \pm 0.13$ for the redshift-independent NLA model targeted for KiDS-like mixed-colour lensing samples¹¹. This prediction aligns well with the

¹¹Fortuna et al. (2021a) also examined the NLA- z model under similar conditions, but found the fits were predominantly driven by the low-redshift bins, resulting in less accurate recovery of large-scale alignments

constraints from recent cosmic shear analyses (A21; Secco et al. 2022; Li et al. 2023c; Dalal et al. 2023). Moreover, recent studies revealed that other nuisance parameters in such analyses, especially those related to redshift calibration uncertainties, can result in misleading A_{IA} values (Hikage et al. 2019; Wright et al. 2020b; Li et al. 2021; Fischbacher et al. 2023).

Given these considerations, we consider it is necessary to explore the prior for the A_{IA} parameter. As an initial step towards a fully informed A_{IA} approach, we begin by simply narrowing the previously broad prior, leaving a more comprehensive exploration of the IA model setups for the forthcoming KiDS-Legacy analysis. In our fiducial analysis, we choose a flat yet narrower prior of $[-0.2, 1.1]$, which corresponds to the 5σ credible region of predictions by Fortuna et al. (2021a). We note that our new prior will not significantly impact the sampling results, provided that the final posterior distributions fall within the set prior range. For comparison purposes, we also conduct a test run using the wider $[-6, 6]$ prior.

Sampling the high-dimensional posterior distribution is a challenging task. In the KiDS-1000-v1 analyses, an ellipsoidal nested sampling algorithm, **MULTINEST** (Feroz et al. 2009), was used. However, recent studies demonstrated that **MULTINEST** systematically underestimates the 68% credible intervals for S_8 by about 10% in current weak lensing analyses (Lemos et al. 2023; DK23; Li et al. 2023c). A promising alternative is the sliced nested sampling algorithm, **POLYCHORD** (Handley et al. 2015a,b), which provides more accurate estimates of parameter uncertainties. However, **POLYCHORD** is almost five times slower than **MULTINEST**. As a result, we opt to use **POLYCHORD** for our main analysis, while retaining **MULTINEST** for our testing purposes.

When presenting point estimates and associated uncertainties for parameter constraints, we adhere to the recommendations of Joachimi et al. (2021). We derive our best-fit point estimates from the parameter values at the maximum of the joint posterior (MAP). Given that the MAP reported by the sampling code can be affected by noise due to the finite number of samples, we enhance the precision of the MAP by conducting an additional local optimisation step. This process initiates from the MAP reported by the sampling code and utilises the Nelder-Mead minimisation method (Nelder & Mead 1965), a method also employed by A21. To represent uncertainties linked to these estimates, we compute the 68% credible interval based on the joint, multi-dimensional highest posterior density region, projected onto the marginal posterior of the parameter of interest (PJ-HPD). This hybrid approach is more robust against projection effects stemming from high-dimensional asymmetric posterior distributions than traditional 1D marginal summary statistics (refer to Sect. 6 in Joachimi et al. 2021 for a comprehensive discussion). To facilitate comparison with results from other surveys, we also provide constraints based on the traditional mean and maximum of the 1D marginal posterior, along with their respective 68% credible intervals.

It is worth noting that, as systematic uncertainties from shear calibration are excluded in the construction of our covariance matrix (see Sect. 4.3.2), the uncertainties derived from the main sampling chains do not fully account for the true uncertainties. To

at high redshifts.

Table 4.2: Fiducial model parameters and their priors.

Parameter	Symbol	Prior
Density fluctuation amp.	S_8	[0.1, 1.3]
CDM density	ω_c	[0.051, 0.255]
Baryon density	ω_b	[0.019, 0.026]
Hubble constant	h	[0.64, 0.82]
Scalar spectral index	n_s	[0.84, 1.1]
AGN heating temperature	$\log_{10}(T_{\text{AGN}}[\text{K}])$	[7.3, 8.0]
Intrinsic alignment amp.	A_{IA}	[-0.2, 1.1]
Redshift offsets	δ_z	$\mathcal{N}(\mu; \sigma^2)$

The first section lists the primary cosmological parameters describing the Λ CDM model assumed, while the second section contains nuisance parameters related to baryon feedback, intrinsic alignments, and redshift biases. The values in square brackets indicate the limits of top-hat priors. The notation $\mathcal{N}(\mu; \sigma^2)$ refers to a normal prior with mean μ and (co-)variance σ^2 , as specified in Table 4.1.

compensate for the additional uncertainties arising from residual shear biases, we employ a forward modelling approach. This method involves shifting the data vector and subsequently the likelihood, based on the estimated residual shear biases, followed by recalculating the MAP. As the adjustment is minor and the covariance matrix remains static, it is not necessary to re-sample the posterior distribution. Instead, we simply need to repeat the previously mentioned local optimisation step. Starting with the original MAP and using the updated likelihood, we can determine the new MAP corresponding to each shift in the data vector. The variation in these MAP estimates represents additional uncertainties introduced by the systematic uncertainties arising from shear calibration. Further details on this process can be found in Appendix 4.C.

Table 4.2 summarises the model parameters and their priors as used in our fiducial analysis. These parameters can be broadly classified into two categories: the first category includes five cosmological parameters, which describe the spatially flat Λ CDM model we employ. We fix the sum of the neutrino masses to a value of $0.06 \text{ eV } c^{-2}$, where c is the speed of light. This choice is based on Hildebrandt et al. (2020)’s finding of the negligible influence of neutrinos on cosmic shear analyses. The second category encompasses three nuisance parameters, accounting for astrophysical and measurement uncertainties as previously discussed. We note that all parameters, with the exception of T_{AGN} and A_{IA} , retain the same priors as those used in the KiDS-1000-v1 cosmic shear analyses. The T_{AGN} parameter replaces the previous baryon feedback amplitude parameter associated with the preceding version of HMCODE, while the A_{IA} parameter adopts a narrower prior for reasons previously discussed.

4.5 Results

In this section, we present our cosmological parameter constraints and evaluate the robustness of our findings against a variety of systematic uncertainties. We begin by presenting the outcomes from our fiducial analysis in Sect. 4.5.1. We then assess the impact of shear biases in Sect. 4.5.2, by quantifying the shifts in final constraints resulting from different shear bias scenarios. This highlights the main development of our work. Additionally, since we implemented several changes to the cosmological inference pipeline, we evaluate the effects of these adjustments by comparing results from multiple setup variations in Sect. 4.5.3.

4.5.1 Fiducial analysis results

Our fiducial model has a total of twelve free parameters: five are cosmological parameters specifying the spatially flat Λ CDM model with a fixed total neutrino mass, and the remaining seven are nuisance parameters addressing astrophysical and redshift calibration uncertainties, as detailed in Sect. 4.4. However, not all of these parameters are constrained by the cosmic shear analysis. In this section, we focus on the primary parameters that our analysis constrains. Meanwhile, the posterior distributions for all free parameters are displayed as contour plots in Appendix 4.D for reference.

Table 4.3 provides the point estimates along with their corresponding 68% credible intervals for the primary parameter as constrained by our fiducial analysis using the POLYCHORD sampling code. We display results using three summary statistics: MAP and PJ-HPD, the mean of the 1D marginal posterior, and the maximum of the 1D marginal. As discussed in DK23, each of these approaches has its own advantages and limitations. Specifically, the accurate determination of MAP and PJ-HPD can be challenging, while marginal constraints for multi-dimensional posteriors are prone to projection effects. Aligning with the KiDS convention, we choose the MAP and PJ-HPD constraints as our headline results, but caution against direct comparisons with results from other surveys that might use different summary statistics. The uncertainties we report include additional contributions from the systematic uncertainties associated with our shear calibration, as detailed in Sect. 4.5.2. These additional uncertainties are overall small compared to the main sampling uncertainties, so when plotting the posterior distributions or conducting extended runs for test purposes, we do not incorporate these uncertainties.

Figure 4.4 shows the projected 2D posterior distributions for the parameters Ω_m and S_8 , as derived from our fiducial setups employing POLYCHORD and MULTINEST. We see that MULTINEST results yield a roughly 10% narrower width of the posterior distribution compared to POLYCHORD, aligning with previous findings (Lemos et al. 2023; DK23; Li et al. 2023c). However, as expected, the results from the two sampling codes show consistency in terms of best-fit values. In addition, we compare these results with those from the cosmic microwave background (CMB) analysis by the *Planck* satellite, using their baseline Λ CDM chains with the Plik likelihood from their most recent *Planck*-2018 results (Planck Collaboration et al. 2020). An offset is evident between

Table 4.3: Primary parameter constraints from our fiducial analysis, based on the KiDS-1000-v2 catalogue, as determined using the POLYCHORD sampling code.

Parameter	MAP & PJ-HPD	Marginal	
		Mean	Maximum
S_8	$0.776^{+0.029+0.002}_{-0.027-0.003}$	$0.765^{+0.029}_{-0.023}$	$0.769^{+0.027}_{-0.029}$
Ω_m	$0.259^{+0.115+0.001}_{-0.064-0.001}$	$0.302^{+0.062}_{-0.115}$	$0.273^{+0.102}_{-0.088}$
σ_8	$0.835^{+0.151+0.002}_{-0.158-0.003}$	$0.791^{+0.124}_{-0.163}$	$0.752^{+0.173}_{-0.129}$
A_{1A}	$0.348^{+0.350+0.009}_{-0.322-0.011}$	$0.400^{+0.330}_{-0.339}$	$0.397^{+0.340}_{-0.346}$

Our headline results, based on the MAP & PJ-HPD statistics, include additional uncertainties that account for systematic uncertainties within the shear calibration. These uncertainties, originating from minor deviations from realism in the image simulations and the shear measurement algorithm’s sensitivity to the morphology of the galaxy sample, are estimated using a forward modelling approach (as detailed in Sect. 4.5.2). On the other hand, the statistical uncertainties within the shear calibration, determined by the simulation volume, are folded into the main uncertainties through their inclusion in the covariance matrix used for the cosmological inference. The mean-marginal is determined through `postprocess` within `CosmoSIS` using the default settings (Zuntz et al. 2015); while the max-marginal is calculated using the `ChainConsumer` with the settings of `statistics=‘max’` and `kde=1.0` (Hinton 2016). The indicated uncertainties correspond to the 68% credible intervals.

our cosmic shear results and those from *Planck*-2018. Adopting the Hellinger distance tension metric (Beran 1977; Heymans et al. 2021; DK23), we detected a 2.35σ tension in the constrained S_8 values. For the constrained parameter set (S_8 , Ω_m), a similar level of tension, 2.30σ , was found using the Monte Carlo exact parameter shift method (Raveri et al. 2020; DK23).

Figure 4.5 presents our primary S_8 constraints and compares them with those from other contemporary cosmic shear surveys and the *Planck* CMB analysis. For ease of comparison, we show all three summary statistics for our fiducial results, while for other surveys, we display their headline values, as per their preferred summary statistics. Overall, our results align well with those from all major contemporary cosmic shear surveys.

We note that our fiducial analysis pipeline is similar to the DK23 Hybrid pipeline with one notable difference: while DK23 included a free neutrino parameter, we kept the total neutrino mass fixed. DK23 showed that this additional degree of freedom in the cosmological parameter space can slightly increase the projected marginal S_8 values relative to an analysis with a fixed neutrino mass. However, since we refer to their




Figure 4.4: Comparison of projected 2D posterior distributions for the parameters Ω_m and S_8 as derived from our fiducial setups using two sampling codes: POLYCHORD (solid black line) and MULTINEST (dashed grey line), against the *Planck*-2018 results (solid red line). The contours correspond to the 68% and 95% credible intervals and are smoothed using a Gaussian Kernel Density Estimation (KDE) with a bandwidth scaled by a factor of 1.5, made possible by the ChainConsumer package (Hinton 2016).

MAP & PJ-HPD results in Fig. 4.5, the comparison should not be influenced by these projection effects (for more details, refer to the discussion in DK23).

It is interesting to note that our fiducial results align almost identically with the KiDS-1000-v1 re-analysis conducted by DK23, which used the A21 redshift calibration. This alignment arises from a balance of several effects in our analysis. Our improved shear calibration tends to increase S_8 , while the enhanced vdB22 redshift calibration tends to lower it. Moreover, our analysis does not show a significant increase in S_8 when introducing scale cuts, as seen in the KiDS-1000-v1 Hybrid analysis. This helps reconcile the minor difference between our results and those of A21. We explore these changes in more detail in the following sections.




Figure 4.5: Marginalised constraints on S_8 derived from our fiducial analysis with POLYCHORD, compared with those from other contemporary cosmic shear surveys and the *Planck* CMB analysis. Three sections, separated by dotted horizontal lines, indicate results of different origins. The first section includes results from individual cosmic shear surveys with their own analysis pipelines. The second section presents results from a collaborative effort between DES and KiDS teams, which built a hybrid pipeline for analysing the data from both groups (DK23). The final section displays results from the *Planck* CMB analysis. Different labels are used for different statistical methods: the diamond represents results using the MAP & PJ-HPD statistics, the square denotes the mean-marginal statistics, and the circle shows the maximum-marginal statistics. The error bars correspond to the 68% credible intervals.

4.5.2 Impact of shear biases

The primary aims of this study are to assess the impact of higher-order shear biases on the final parameter constraints and to develop a methodology for effectively addressing shear calibration uncertainties. Both of these aims can be achieved by examining the shifts in the constrained cosmological parameters resulting from different shear bias scenarios. As discussed in Sect. 4.4 and Appendix 4.C, the residual shear biases have only a minor effect on the measured data vector. This allows us to determine the shifts in the best-fit values of the constrained parameters using a local minimisation algorithm, such as the Nelder-Mead method (Nelder & Mead 1965). These shifts in the best-fit values indicate the additional uncertainties stemming from systematic uncertainties in shear calibration.

Figure 4.6 shows shifts in our primary S_8 constraints for different residual shear bias scenarios. For comparison, we also include a shaded region denoting different levels

of PJ-HPD credible intervals, as derived from our fiducial POLYCHORD chain. Apart from the extreme case where no shear calibration is applied, all other residual shear bias scenarios result in shifts less than 10 per cent of the initial sampling uncertainties. Notably, neglecting the higher-order correction for the shear-interplay effect and uncertainties in PSF modelling results in a negligible shift of only -0.03σ . This finding reinforces the reliability of previous KiDS cosmic shear analyses, which did not consider these higher-order effects.

The S_8 shifts, resulting from the input morphology test simulations, indicate additional systematic uncertainties within our shear bias calibration. The generation of these test simulations is detailed in Appendix 4.C. Briefly, we generated six sets of test simulations, where the input values of three morphological parameters of the adopted Sérsic profile - the half-light radius (labelled as ‘size’ in the figure), axis ratio (labelled as ‘ q ’), and Sérsic index (labelled as ‘ n ’) - were shifted up and down. We observe that shifts in the input galaxy axis ratio lead to the most significant changes in S_8 : a -0.10σ shift for increased input axis ratio and a $+0.06\sigma$ shift for decreased input axis ratio. This behaviour aligns with our expectations for the *lensfit* code employed in our analysis. As it incorporates prior information on measured galaxy ellipticities during its Bayesian fitting process, it is more sensitive to changes in the distributions of sample ellipticities.

These S_8 shifts, obtained from the test simulations, provide a quantitative measure of the potential impact of inaccuracies in the input morphology and the sensitivity of the *lensfit* code to the underlying sample morphology distributions. When presenting the S_8 constraints, we account for these systematic uncertainties by including the maximum shifts into the reported uncertainties. In other words, we consider the shifts corresponding to the changes in input axis ratio (represented as dashed lines in Fig. 4.6), from the six sets of test simulations, as additional systematic uncertainties. These are reported alongside the original statistical uncertainties from the main sampling chain. It should be noted that these additional systematic uncertainties are specific to the SKiLLS image simulations and the *lensfit* shape measurement code used in our analysis. To reduce these uncertainties, future advancements in shear measurements should focus on improving the realism of image simulations and enhancing the robustness of the shear measurement algorithm.

4.5.3 Impact of altering inference setups

Although our main updates revolve around the shear measurement and calibration, we have also implemented several modifications to the cosmological inference pipeline, drawing upon recent developments from DK23. As such, it is beneficial to conduct some extended runs with various setup configurations.

For these test runs, we employ MULTINEST as our sampling code, as it operates approximately five times faster than POLYCHORD, but at the cost of underestimating the width of the posterior distributions and thus the reported uncertainties by about 10%. However, the best-fit values from MULTINEST are not biased (as evident in Fig. 4.4). Thus, comparisons made using MULTINEST will yield conservative but unbiased results.

Table 4.4: Point estimates for S_8 from different inference setups.

Setups	χ^2	MAP & PJ-HPD		Mean marginal		Maximum marginal	
		S_8	ΔS_8	S_8	ΔS_8	S_8	ΔS_8
Fiducial: POLYCHORD	62.7	$0.776^{+0.029+0.002}_{-0.027-0.003}$	-	$0.765^{+0.029}_{-0.023}$	-	$0.769^{+0.027}_{-0.029}$	-
Fiducial: MULTINEST	62.7	$0.774^{+0.027}_{-0.023}$	0.00σ	$0.765^{+0.028}_{-0.017}$	0.00σ	$0.771^{+0.023}_{-0.027}$	0.00σ
NLA: $A_{IA} : [-6, 6]$	62.7	$0.774^{+0.034}_{-0.020}$	0.02σ	$0.766^{+0.028}_{-0.020}$	0.03σ	$0.770^{+0.023}_{-0.026}$	-0.01σ
NLA-z: $A_{IA}, \eta_{IA} : [-5, 5]$	62.7	$0.776^{+0.022}_{-0.025}$	0.10σ	$0.765^{+0.027}_{-0.020}$	0.00σ	$0.770^{+0.023}_{-0.025}$	-0.05σ
Scales: $0.5 < \theta < 300'$	64.5	$0.769^{+0.014}_{-0.036}$	-0.17σ	$0.755^{+0.029}_{-0.022}$	-0.40σ	$0.763^{+0.022}_{-0.030}$	-0.31σ
KiDS-1000-v1 setups	64.9	$0.771^{+0.017}_{-0.029}$	-0.13σ	$0.755^{+0.026}_{-0.020}$	-0.44σ	$0.759^{+0.022}_{-0.024}$	-0.50σ

The ‘Fiducial: POLYCHORD’ denotes our headline results, which are the same as those presented in Table 4.3. The ‘Fiducial: MULTINEST’ represents the same parameter setup as our fiducial analysis, but employs the MULTINEST sampling code. We use this as the reference to assess test results because all test runs utilise the MULTINEST code for increased speed. When comparing the ‘Fiducial: MULTINEST’ results with the primary POLYCHORD results, we can conclude that the best-fit values obtained from MULTINEST are unbiased. The relative shift in S_8 , denoted as ΔS_8 , is calculated by comparing the best-fit values from the test runs to the reference ‘Fiducial: MULTINEST’ run. The ΔS_8 values are expressed as a fraction of σ , which signifies the standard deviation of estimates from the test run. We calculate ΔS_8 for different summary statistics separately for consistency. For MAP & PJ-HPD results, we also present the best-fit χ^2 values. For comparison, the best-fit χ^2 values from A21 and vdB22 are 82.2 and 63.2, respectively.




Figure 4.6: Shifts in best-fit values of S_8 under different residual shear bias scenarios. The shift, ΔS_8 , is calculated as $\Delta S_8 = S_8^{\text{test}} - S_8^{\text{fiducial}}$, where S_8^{test} represents the best-fit values in the test scenarios determined by a local minimisation method using the best-fit values from the fiducial analysis (S_8^{fiducial}) as a starting point. The grey shaded regions represent different percentiles of the credible intervals derived from our fiducial POLYCHORD run. From the innermost to the outermost region, these percentiles are 6.8%, 20.4%, and 34%, corresponding to 0.1, 0.3, and 0.5 fractions of the reported sampling uncertainties. The dashed lines display the maximum shifts encountered in the six sets of morphology test simulations. These maximum shifts are used as the additional uncertainties in the reported best-fit values to account for the systematic uncertainties arising from shear calibration.

Priors for the NLA model

We begin by testing the prior for the NLA model. As discussed in Sect. 4.4, our fiducial analysis implemented a redshift-independent NLA model with a narrow flat prior for the amplitude parameter A_{IA} . This model, motivated by the work of Fortuna et al. (2021a), serves as an alternative to the uninformative broad prior previously used. To investigate the impact of this change on our final results, we performed two additional runs: one employing a redshift-independent NLA model with a broad A_{IA} prior ranging from $[-6, 6]$, in line with KiDS-1000-v1 analyses, and another allowing for a redshift-dependent IA amplitude, i.e., the NLA- z variant. The redshift evolution is modelled using a power-law of the form $[(1+z)/(1+0.62)]^{\eta_{\text{IA}}}$, with priors of $[-5, 5]$ for both A_{IA} and η_{IA} , in line with DK23.

Figure 4.7 presents a comparison of the posterior distributions obtained from the different NLA prior setups, and Table 4.4 lists the point estimates for the critical S_8 parameter. We see consistent constraints on S_8 across all setups. The constrained A_{IA} under our narrower prior setup also aligns with those from the broad priors, albeit




Figure 4.7: Comparison of projected posterior distributions for the parameters S_8 , A_{IA} , and η_{IA} , derived from three different NLA prior setups. The contours correspond to the 68% and 95% credible intervals and are smoothed using Gaussian KDE with a bandwidth scaled by a factor of 1.5.

spanning a narrower range due to the constrained prior range, validating the prior range used in our fiducial analysis. Additionally, we observe that the η_{IA} parameter is not constrained by the data, suggesting that the use of the NLA-z model may not be necessary for current weak lensing analyses.

Different scale cuts

In our fiducial analysis, we adopted a scale cut for the measured data vectors, ranging from $2'$ to $300'$, as suggested by DK23. This is a change from the KiDS-1000-v1 analyses, which used a range of $0.5' < \theta < 300'$. A re-analysis of KiDS-1000-v1 with this new scale cut by DK23 led to a $0.7 - 0.8\sigma$ increase in the S_8 constraint. Using mock analyses, they found that this offset could arise from noise fluctuations 23% of the time.

In light of the updates to our shear measurement, we revisited this test. Interestingly, we observed a smaller difference between the two scale cuts than what was reported by DK23. Specifically, we observed shifts of -0.17σ , -0.40σ , and -0.31σ , corresponding

to the MAP & PJ-HPD, mean marginal, and maximum marginal summary statistics, respectively (refer to Table 4.4 for exact values).

We attribute this increased robustness against small scale fluctuations to our improved empirical corrections of the PSF leakages into shear measurement. This is supported by Figs. 4.2 and 4.3, where we see that the shear signals measured from the KiDS-1000-v2 catalogues exhibit overall smaller systematic errors. We note that Giblin et al. (2021) performed a mock test using the two-point correlation function and identified a change of less than 0.1σ in the S_8 constraints when the detected PSF residuals were incorporated into the KiDS-1000-v1 mock data. Nevertheless, it is plausible that these systematic effects have a more significant influence on COSEBIs, given their use of more sophisticated weighting functions (Schneider et al. 2010). To quantify the improvements brought about by the updated shear measurements regarding the robustness of the COSEBIs, a similar mock analysis based on the COSEBIs statistic is warranted. We consider this an important topic for future study. For the current analysis, the test results simply affirm the robustness of our primary S_8 constraints.

KiDS-1000-v1 setups

To draw a direct comparison with the KiDS-1000-v1 results and evaluate the impact of our improved shear measurements and calibration, we performed a test run using the same inference pipeline and parameter priors as in the KiDS-1000-v1 analyses conducted by A21 and vdB22. The differences compared to our fiducial analysis setup include: measurements from scales of 0.5 to $300''$, use of the older version of `HMCODE`, sampling with the `MULTINEST` code, and a broad A_{IA} prior ranging from $[-6, 6]$. As shown in Fig. 4.8, our test results are well-aligned with the outcomes of the analyses by A21 and vdB22. Notably, our new results show an increase in the S_8 value relative to vdB22, bringing it closer to the result obtained by A21.

We re-emphasise that our redshift calibration aligns with that of vdB22, who expanded the redshift calibration sample to more than double the size used by A21 (see Appendix 4.B for details). This means that our redshift-related selection function closely mirrors that used in the vdB22 sample. However, due to changes in the weighting and selection scheme between the KiDS-1000-v2 catalogue and the KiDS-1000-v1 catalogue, our sample cannot be considered as directly comparable to theirs.

To provide a more quantitative understanding of the sample differences among the three analyses, we compared the effective number density of the source sample in our analysis to those used in A21 and vdB22. The differences for each tomographic bin are 9.6%, 9.8%, 6.1%, 10.6%, and 2.8% when compared to A21; and -1.8% , -1.3% , -0.7% , 0.7% , and 3% when compared to vdB22. Here, positive values signify an increase, while negative values denote a decrease. The differences between our catalogue and that of A21 stem from both shear measurement and redshift calibration, whereas the difference between ours and that of vdB22 arises mainly from the shear measurement, as we used the same SOM for the ‘gold’ selection (see Appendix 4.B). As such, comparing our results directly with those of vdB22 can provide clearer insights into the impact of




Figure 4.8: Comparison of projected posterior distributions for parameters Ω_m , S_8 , and A_{IA} from our analysis (grey dashed lines) based on the KiDS-1000-v2 catalogue, to those from vdB22 (orange solid lines) and A21 (green dotted lines), both of which are based on the KiDS-1000-v1 catalogues. The cosmological inference pipeline and parameter priors are identical across all three analyses presented here. In terms of measurements, vdB22 and A21 use the same shear measurements and calibration, while vdB22 and our analysis share the same redshift calibration. The contours correspond to the 68% and 95% credible intervals and are smoothed using Gaussian KDE with a bandwidth scaled by a factor of 1.5.

our improvements in shear measurements. It is also worth noting that the increased effective number density in high redshift bins compared to vdB22 is largely due to the increased weighting of faint objects in the updated version of *lensfit* code. However, this comes at the cost of increased sample ellipticity dispersion, with a maximum increase of 6% found in the fifth bin. These subtle differences in the source catalogues change the noise properties of the samples. Consequently, even with perfect calibration in each study, we would not expect to derive identical cosmological constraints from each analysis.

4.6 Summary

We conducted a cosmic shear analysis using the KiDS-1000-v2 catalogue, which is an updated version of the public KiDS-1000(-v1) catalogue with respect to shear measurements and calibration. Under the assumption of a spatially flat Λ CDM cosmological model, we derived constraints on $S_8 = 0.776^{+0.029+0.002}_{-0.027-0.003}$ based on the MAP & PJ-HPD summary statistics. The second set of uncertainties were incorporated to account for the systematic uncertainties within our shear calibration. The mean-marginal and maximum-marginal values obtained from the same sampling chain are $0.765^{+0.029}_{-0.023}$ and $0.769^{+0.027}_{-0.029}$, respectively. Our results are consistent with earlier results from KiDS-1000-v1 and other contemporary weak lensing surveys, but show a $\sim 2.3\sigma$ level of tension with the *Planck* cosmic microwave background constraints.

The main improvements in our analysis, relative to the KiDS-1000-v1 cosmic shear analyses, are attributed to the enhanced cosmic shear measurement and calibration. These enhancements were achieved through the updated version of the *lensfit* shape measurement code, a new empirical correction scheme for PSF contamination, and the newly developed SKiLLS multi-band image simulations, as detailed in L23. We verified the reliability of the new measurement via a series of catalogue-level null tests proposed by Giblin et al. (2021). The results indicate that the KiDS-1000-v2 catalogue shows overall better control over measurement systematics compared to the KiDS-1000-v1 catalogues. This improvement in reducing measurement systematics assists in reducing noise in small scale measurements, thereby enhancing the robustness of our cosmological parameter constraints against varying scale cut choices.

Our methodology for shear calibration largely aligns with the one detailed in L23, where we account for higher-order blending effects that arise when galaxies from different redshifts are blended, as well as the uncertainties in PSF modelling. However, when comparing the outcomes from the shear calibration with and without these higher-order adjustments, we found that these effects have a negligible impact on the present weak lensing analysis, a conclusion that is in line with the findings of Amon et al. (2022).

We recommend treating the statistical and systematic uncertainties from the shear calibration separately, given their distinct origins. The statistical uncertainties, which are determined by the simulation volume, can be reduced and are readily incorporated into the covariance matrix used for cosmological inference. On the other hand, systematic uncertainties, associated with the realism of image simulations and sensitivity of the shape measurement algorithm, can be more effectively addressed when considered as residual shear biases post-calibration. Assuming these residual shear biases are small, a forward modelling approach, combined with a local minimisation method, can be used to estimate their impact on the final parameter constraints. In our analysis, these additional systematic uncertainties contribute roughly 8% of the final uncertainty on S_8 . However, ongoing efforts to enhance shear measurement and calibration, such as increasing the realism of image simulations through Monte-Carlo Control Loops (Refregier & Amara 2014) and leveraging new techniques like METACALIBRATION/METADETECTION (Huff & Mandelbaum 2017; Sheldon & Huff 2017; Sheldon et al. 2020; Hoekstra et al. 2021) to

improve measurement robustness against underlying sample properties, may well lead to a reduction in these additional systematic uncertainties.

In our fiducial analysis, we opted for a redshift-independent NLA model with a narrow flat prior for the IA amplitude parameter, A_{IA} , motivated by the work of Fortuna et al. (2021a). However, we also investigated two alternative scenarios: one with a broad A_{IA} prior for the redshift-independent NLA model, echoing the KiDS-1000-v1 analysis by A21, and the other, the NLA- z variant, allowing for redshift evolution of the IA amplitude, as per the recent joint DES Y3+KiDS-1000 cosmic shear analysis (DK23). In all three scenarios, we found fully consistent constraints for S_8 and A_{IA} , indicating that the impact of the variations among these scenarios is negligible. To better understand the IA signals and their impact on cosmic shear analyses, future tests need to implement more substantial variations in IA models, for instance, the halo model formalism introduced by Fortuna et al. (2021a). Such exploration would not only enhance our understanding of the measured IA signals, but also help mitigate correlations between nuisance parameters, thereby improving the precision of future cosmic shear analyses.

Acknowledgements

We acknowledge support from: the Netherlands Research School for Astronomy (SSL); the Netherlands Organisation for Scientific Research (NWO) under Vici grant 639.043.512 (HHo); the Royal Society and Imperial College (KK); the Polish National Science Center through grants no. 2020/38/E/ST9/00395, 2018/30/E/ST9/00698, 2018/31/G/ST9/03388 and 2020/39/B/ST9/03494 (MB); the Polish Ministry of Science and Higher Education through grant DIR/WK/2018/12 (MB); the Royal Society through an Enhancement Award (RGF/EA/181006) and the Royal Society of Edinburgh for support through the Saltire Early Career Fellowship with ref. number 1914 (BG); the European Research Council (ERC) under Grant number 647112 (CH) and Consolidator Grant number 770935 (HHi, JLvdB, AHW, RR); the Max Planck Society and the Alexander von Humboldt Foundation in the framework of the Max Planck-Humboldt Research Award endowed by the Federal Ministry of Education and Research (CH); the UK Science and Technology Facilities Council (STFC) under grant ST/V000594/1 (CH) and ST/V000780/1 (BJ); Heisenberg grant of the Deutsche Forschungsgemeinschaft grant Hi 1495/5-1 (HHi); CMS-CSST-2021-A01 and CMS-CSST-2021-A04, NSFC of China under grant 11973070 (HYS); Key Research Program of Frontier Sciences, CAS, Grant No. ZDBS-LY-7013 (HYS); Program of Shanghai Academic/Technology Research Leader (HYS). 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.A3017, 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.

Author Contributions: All authors contributed to the development and writing of this paper. The authorship list is given in three groups: the lead authors (SSL, HHo, KK) followed by two alphabetical groups. The first alphabetical group includes those who are key contributors to both the scientific analysis and the data products. The second group covers those who have either made a significant contribution to the data products, or to the scientific analysis.

4.A Shear bias in two-point statistics

When calibrating the shear measurements in the two-point correlation function, it is usually assumed that the correlations involving the shear bias can be ignored, which includes correlations between different tomographic and spatial angular bins. This simplification leads to the following relationship between the true correlation function of cosmic shear in tomographic bins i, j , denoted as ξ^{ij} , and the measured signal $\hat{\xi}^{ij}$:

$$\begin{aligned}\hat{\xi}^{ij}(\theta) &= \langle \hat{\gamma}^i(\theta') \hat{\gamma}^j(\theta' + \theta) \rangle \\ &= \langle [1 + m^i(\theta')] [1 + m^j(\theta' + \theta)] \gamma^i(\theta') \gamma^j(\theta' + \theta) \rangle \\ &= \langle [1 + m^i(\theta')] [1 + m^j(\theta' + \theta)] \rangle \xi^{ij}(\theta) \\ &\simeq (1 + \overline{m^i}) (1 + \overline{m^j}) \xi^{ij}(\theta),\end{aligned}\tag{4.4}$$

where $\overline{m^i}$ is estimated by averaging over all sources in a given tomographic bin i , and we use $\langle \cdot \rangle$ to denote the correlation function. We also assumed that the shear bias is independent of the underlying shear to simplify the equation. The result of Eq. (4.4) allows us to average the multiplicative biases over all the galaxies in a given tomographic bin to mitigate the individual noisy bias estimation.

However, in principle, the shear bias can be scale dependent due to spatial fluctuations in source density (e.g. Samuroff et al. 2018). With SKiLLS, we can directly examine these correlations by measuring the shear bias in the two-point estimators. We measured the shear correlation function in the SKiLLS mock catalogue using Eq. (4.2). Since we know the true $\xi_+^{ij}(\theta) = \gamma_{\text{input}}^2$ in simulations, where γ_{input} is the amplitude of the constant input shear, we can estimate the shear bias in the two-point correlation function directly by comparing the measured $\hat{\xi}_+^{ij}$ to the input ξ_+^{ij} following Eq. (4.4).

Figure 4.9 shows the difference between the shear biases with and without considering its correlations, defined as $\Delta m_\xi \equiv \langle [1 + m^i(\theta')] [1 + m^j(\theta' + \theta)] \rangle - (1 + \overline{m^i}) (1 + \overline{m^j})$. It shows that the difference is negligible across all scales and tomographic bins, in agreement with the statistical uncertainties of our shear calibration, which are represented by the shaded regions. These findings confirm that we can neglect the correlations between shear biases in the current KiDS weak lensing analysis.

4.B SOM redshift calibration

This appendix provides information on the redshift calibration reference sample and SOM configurations used in our analysis. For a more comprehensive overview and validation of the SOM redshift calibration method in the KiDS analysis, we refer to Wright et al. (2020b), Hildebrandt et al. (2021) and vdB22.

We employed the fiducial spectroscopic sample described in vdB22 as our calibration reference sample. This sample comprises spectroscopic redshift estimates (spec- z s) from various spectroscopic surveys that overlap with KiDS fields, enabling us to assign KiDS photometric measurements to objects in the reference sample. In cases where an




Figure 4.9: Two-point correlations between the multiplicative shear biases. The correlation is estimated as $\Delta m_\xi \equiv \hat{\xi}_+^{ij} / \gamma_{\text{input}}^2 - (1 + \overline{m}^i) (1 + \overline{m}^j)$. The 15 panels represent the different combinations of the five redshift bins utilised in our cosmic shear analysis. The shaded regions within each panel denote the statistical uncertainties of our shear calibration for each tomographic bin, as outlined in Table 4.1.

object had multiple spectroscopic measurements, vdB22 defined a specific hierarchy to select the most reliable redshift estimates based on the quality of the measurements. For further details on the adopted spectroscopic samples and the compilation procedure, readers are referred to Appendix A of vdB22.

For our calibration, we used a 101×101 hexagonal SOM trained on the r -band magnitude and 36 colours derived from the PSF-matched, list-driven nine-band $ugriZYZJHK_s$ photometry from the KiDS+VIKING surveys. This SOM is identical to the fiducial SOM constructed in vdB22. We segregated the reference and target samples into the trained SOM cells separately for each tomographic bin, allowing us to create comparable groupings between the spectroscopic and photometric sources in each bin. During this process, we further categorised the original SOM cells using a hierarchical cluster analysis implemented by the ‘hclust’ function within the R Stats Package¹² to increase the number of galaxies per grouping. We adopted the same number of clusters per

¹²<https://www.rdocumentation.org/packages/stats/versions/3.6.2/topics/hclust>

bin (4000, 2200, 2800, 4200, and 2000) as Wright et al. 2020b, who determined these numbers using simulations produced by van den Busch et al. (2020).

To mitigate the effects of photometric noise and the incompleteness of the reference sample, we applied an additional selection step to the SOM groupings. We excluded any grouping where the mean spectroscopic redshift of the reference sample $\overline{z_{\text{spec}}}$ and the mean photometric redshift of the target sample $\overline{z_B}$ exhibited a significant discrepancy, defined as $|\overline{z_{\text{spec}}} - \overline{z_B}| > 5\sigma_{\text{mad}}$. Here, σ_{mad} represents the normalised median absolute deviation of all SOM groupings, which we calculated to be 0.122 in our case. This step allowed us to define the KiDS ‘gold’ sample, which we used to compute the redshift distributions and perform the cosmic shear analysis.

4.C Systematic uncertainties from the shear calibration

In this appendix, we outline our approach to address the systematic uncertainties arising from shear calibration. Our methodology involves two primary steps: In Sect. 4.C.1, we quantify the potential residual biases after implementing our simulation-based shear calibration. In Sect. 4.C.2, we propagate these systematic uncertainties into the final uncertainties of the estimated cosmological parameters.

We propose a separate accounting of the shear calibration uncertainties, as it is considered more accurate and informative than the traditional approach, which uses nominal shear calibration uncertainties that are deliberately overestimated to encompass potential systematic uncertainties arising from shear calibration. Our approach clearly illustrates the extent to which the final cosmological parameters of interest are influenced by these systematic uncertainties from shear calibration.

Furthermore, as mentioned in Sect. 4.3.2, these systematic uncertainties have fundamentally different origins from the statistical uncertainties incorporated in the covariance matrix. They represent the fundamental limitations of current simulation-based shear calibration methods. The limitations inherent in these systematic uncertainties cannot be eliminated by merely increasing the scale of image simulations. However they can be mitigated by empirically enhancing the realism of the image simulations, for example, using the Monte-Carlo Control Loop method (Refregier & Amara 2014), or by improving the robustness of the shear measurement algorithm, such as the METACALIBRATION/METADETECTION method (Huff & Mandelbaum 2017; Sheldon & Huff 2017; Sheldon et al. 2020; Hoekstra et al. 2021).

4.C.1 Quantifying residual shear biases with sensitivity analysis

Residual biases may persist after simulation-based shear calibration due to imperfect alignment between simulations and data, as elucidated by K19. These discrepancies pose challenges for shear calibration methods dependent on image simulations and underscore the need for re-weighting simulations to more closely align with the data. However, given that intrinsic galaxy properties in real data are unknown, this re-weighting process relies on noisy measured properties, rendering it vulnerable to calibration selection




Figure 4.10: Changes in input morphological parameter values for assessing residual biases after simulation-based shear calibration. From left to right, the order is half-light radius, axis ratio, and Sérsic index. *Top panels:* The relationship between relative fitting uncertainties and the r -band magnitude, as measured from the catalogues of Griffith et al. (2012). The values shown are calculated using the median values of the binned samples. *Middle panels:* The overall distributions of input morphological parameters. Comparisons are drawn among the fiducial (grey shades), the test sample with input values increased by an amount corresponding to the relative uncertainties (red lines), and the test sample with input values decreased by the same amount of relative uncertainties (blue lines). *Bottom panels:* The shifts in multiplicative biases in tomographic bins resulting from changes in input morphological parameter values. Both fiducial and test simulations are re-weighted using the same data set, with the method detailed in Sect. 4.3.1.

biases as discussed by FC17. The uncertainties linked with the measured properties cause galaxies to be intermixed among defined bins, leading to the up-weighting or down-weighting of certain galaxies. As a result, even if the re-weighted sample aligns with the data in terms of the distribution of measured properties, it does not ensure identity in terms of intrinsic properties. In other words, shear biases can still vary between two samples with identical distributions of apparent measured properties. Our aim is to quantify these residual biases and incorporate them into the final uncertainties of cosmological parameters.

The SKiLLS multi-band image simulations used in this analysis incorporate several enhancements, informed by insights gathered from previous KiDS simulation studies (FC17, K19). These improvements include: reproducing variations in star density, PSF, and noise background across the KiDS footprint; incorporating faint galaxies down to an r -band magnitude of 27 to account for correlated noise from undetected objects (e.g. Hoekstra et al. 2017); including realistic clustering from N -body simulations to address blending effects (e.g. K19); and adopting an end-to-end approach for photo- z estimation to account for photo- z measurement uncertainties. These improvements augment the robustness of the shear biases estimated from SKiLLS against various observational conditions.

In an investigation on the propagation of observational biases in shear surveys, Kitching et al. (2019) demonstrated that the measured shear power spectrum is, to first order, predominantly influenced by the mean of the multiplicative bias field across a survey. This suggests that if the shear bias estimated from simulations accurately reflects the mean value of the targeted sample, the shear calibration will be robust enough for KiDS-like cosmic shear analyses. Therefore, we conclude that potential residual biases related to observational conditions have negligible influence on our shear calibration, and we focus on systematic uncertainties arising from galaxy morphology uncertainties, specifically the assumed Sérsic profile and its parameters derived from *Hubble* Space Telescope observations (Griffith et al. 2012). For a model-fitting shape measurement code like *lensfit*, these galaxy morphology uncertainties are the main sources of residual shear biases after implementing the simulation-based shear calibration.

The deviation from the Sérsic profile is challenging to address for the current SKiLLS simulations, as our copula-based learning algorithm requires a parameterised model for its application. However, the Sérsic model has been validated as sufficient for KiDS-like analyses by K19, who used the same morphology catalogue as our work. Thus, we focus on the measurement uncertainties of the Sérsic parameters: half-light radius, axis ratio, and Sérsic index. We first examined the fitting uncertainties reported by Griffith et al. (2012) to assess the accuracy of these parameters in our input catalogue. We found that the median relative uncertainties for these parameters are a smooth function of galaxy magnitude, as shown in the top panels of Fig. 4.10. This allows us to capture these correlations through simple linear interpolation.

We interpreted these relative uncertainties as indicators of the systematic uncertainties in our input morphology. We assumed the most extreme scenarios, in which these measured statistical uncertainties are all caused by a coherent bias in the same direction.

Consequently, we adjusted all galaxies in our input sample in the same direction, with the amplitude of the adjustment determined based on their r -band magnitude using a simple linear interpolation of the measured median correlations. We examined shifts towards both larger and smaller values and considered the three Sérsic parameters separately. This resulted in six test simulations corresponding to the six different sets of variations in input morphology parameter values. The input parameter distributions for these test simulations, as shown in the middle panels of Fig. 4.10, are compared to the distributions of the fiducial simulations. A clear shift of the entire distribution is evident, suggesting that our test simulations represent the most extreme scenarios in which the measured statistical uncertainties are coherently biased in the same direction, a situation that is unlikely in reality. Therefore, the residual biases we identified from these test simulations provide a conservative estimate.

We applied the same data analysis procedures to the test simulations as we did to the fiducial simulations, including shear and redshift estimates. We also followed the same re-weighting procedure for the test simulations as for the fiducial simulations, ensuring that the calibration selection biases are also captured. The differences in shear biases between these test simulations and our fiducial simulation are illustrated in the bottom panels of Fig. 4.10. The small differences indicate that the residual shear biases, after implementing our fiducial shear bias calibration, are insignificant.

4.C.2 Propagating residual shear biases with forward modelling

Accurately incorporating the systematic uncertainties from shear calibration into the covariance matrix presents a challenge, as residual shear biases directly scale the data vector, as shown in Eq. (4.4). A more direct approach is to assess the shift in the measured shear signal caused by the residual shear biases and evaluate how these data vector shifts influence the constrained cosmological parameters. Given the minor residual shear biases illustrated in Fig. 4.10 and the unchanged covariance, it is not necessary to reiterate the sampling of the posterior distributions for each shift. Instead, we can implement a local minimisation algorithm to find nearby best-fit values for each shift, using starting points from the fiducial sampling chain. The range of these new best-fit values, each associated with a shift, indicates the additional systematic uncertainties introduced by the residual shear biases.

This approach naturally integrates with our existing cosmological inference method, as outlined in Sect. 4.4, which already requires an additional local optimisation step to refine the best-fit values identified by the sampling code. We simply replicated this optimisation step, using the original best-fit value as the starting point and the shifted likelihood to determine the best-fit values associated with various alterations in measured signals. The variability in these test best-fit values provides an expanded credible region for the inferred parameters, thereby representing the systematic uncertainties from shear calibration. We included these additional uncertainties when presenting the point estimates of our primary parameters (see Sect. 4.5.2 for details).




Figure 4.11: Posterior distributions of cosmological and baryonic parameters from our fiducial analyses, as generated by POLYCHORD (black solid lines) and MULTINEST (grey dashed lines). The contours represent the 68% and 95% credible intervals and are smoothed using a Gaussian KDE with a bandwidth scaled by a factor of 1.5. We note that S_8 is the only parameter that our data robustly constrains.

4.D Contour plots for all free parameters

In this appendix, we provide two supplementary contour plots that display the posterior distributions of all twelve free parameters from our fiducial analyses, as produced by both the POLYCHORD and MULTINEST sampling codes. The overall concordance between the results generated by POLYCHORD and MULTINEST is evident.




Figure 4.12: Posterior distributions of S_8 and nuisance parameters from our fiducial analyses, as generated by POLYCHORD (black solid lines) and MULTINEST (grey dashed lines). The contours represent the 68% and 95% credible intervals and are smoothed using a Gaussian KDE with a bandwidth scaled by a factor of 1.5. We note that the Gaussian priors we have set, as outlined in Table 4.2, strongly influence the redshift offset parameters. The dotted lines represent the central values of these Gaussian priors.