

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

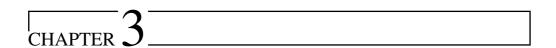
Version: Publisher's Version

Licence agreement concerning inclusion of doctoral thesis License:

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).



KiDS-Legacy calibration: Unifying shear and redshift calibration with the SKiLLS multi-band image simulations

Abstract

We present SKiLLS, a suite of multi-band image simulations for the weak lensing analysis of the complete Kilo-Degree Survey (KiDS), dubbed KiDS-Legacy analysis. The resulting catalogues enable joint shear and redshift calibration, enhancing the realism and accuracy over previous efforts. By integrating cosmological simulations with high-quality imaging data, we created a large volume of simulated galaxies, faithfully mirroring the realistic galaxy properties to a sufficient depth. We accounted for point spread function (PSF) variability across CCD images, while also considering variations in stellar density and noise levels between different pointings. Using variable shear fields, we studied blended systems at different redshifts. Our results show a discernible redshiftshear bias correlation, underscoring this subtle yet noticeable higher-order blending effect. Additionally, we detected a minor but noteworthy impact of PSF modelling errors on shear bias. Finally, we conducted sensitivity tests, affirming the robustness of our fiducial shape measurement algorithm, lensfit, within the KiDS weak lensing analysis requirement. For future, more stringent weak lensing surveys, we recommend further study into blending effects, refinement of PSF modelling, and improvement of shape measurement techniques to minimise their sensitivity to variations in galaxy properties.

3.1 Introduction

Weak gravitational lensing, the small deflection of light rays caused by inhomogeneous matter distributions, is a powerful tool for observational cosmology as an unbiased tracer of gravity (see Bartelmann & Schneider 2001, for a review). It allows us to study the underlying distribution of both baryonic and dark matter (see Refregier 2003; Hoekstra & Jain 2008; Kilbinger 2015, for some reviews). Together with redshift estimates for the sources, the cosmological lensing signal can even quantify the growth of the cosmic structure and infer the properties of dark energy (e.g. Hu 1999; Huterer 2002). Recent weak lensing surveys, including the Kilo-Degree Survey + VISTA Kilo-degree INfrared Galaxy (KiDS+VIKING) survey (de Jong et al. 2013; Edge et al. 2013)¹, the Dark Energy Survey (DES, Dark Energy Survey Collaboration et al. 2016)2, and the Hyper Suprime-Cam (HSC) survey (Aihara et al. 2018)³, provided some of the tightest cosmological constraints on the clumpiness of matter in the local Universe (Heymans et al. 2021; Abbott et al. 2022; Hamana et al. 2020). The upcoming so-called Stage IV surveys, such as the ESA Euclid space mission (Laureijs et al. 2011)⁴, the Rubin Observatory Legacy Survey of Space and Time (LSST, Ivezić et al. 2019)5, and the NASA Nancy Grace Roman space telescope (Spergel et al. 2015)6, will advance the field significantly by increasing the statistical power of weak lensing measurements by more than an order of magnitude.

While promising, measuring the weak lensing signals to the desired accuracy in practice is demanding (see Mandelbaum 2018, for a recent review). In particular, the observed images of distant galaxies are smeared by the point spread function (PSF) and contain pixel noise, biasing the measurements of galaxy shapes (e.g. Paulin-Henriksson et al. 2008; Massey et al. 2013; Melchior & Viola 2012; Refregier et al. 2012). These issues drove the early development of many shape measurement methods and triggered a series of community-wide blind challenges based on image simulations, including the Shear TEsting Programme (STEP, Heymans et al. 2006; Massey et al. 2007) and the Gravitational LEnsing Accuracy Testing (GREAT, Bridle et al. 2010; Kitching et al. 2012; Mandelbaum et al. 2015). These early efforts illuminated some crucial issues and paved the way to calibrate the systematic biases for an actual survey using image simulations.

Early applications of simulation-based calibration have already demonstrated that the calibration accuracy depends on how well the simulation matches the survey under consideration, especially the observational conditions and the galaxy properties (e.g. Miller et al. 2013; Hoekstra et al. 2015, 2017; Samuroff et al. 2018). Therefore, recent implementations carefully mimic the data processing procedures and use morphological measurements from deep imaging surveys to reproduce the measured galaxy properties

https://kids.strw.leidenuniv.nl

²https://darkenergysurvey.org

³https://hsc.mtk.nao.ac.jp/ssp/

⁴https://sci.esa.int/web/euclid/

⁵https://www.lsst.org/

⁶https://roman.gsfc.nasa.gov/

3

for a specific survey (e.g. Mandelbaum 2018; Kannawadi et al. 2019 hereafter K19; MacCrann et al. 2022). Alternately, newer methods, such as the Bayesian Fourier Domain (Bernstein & Armstrong 2014) and Metacalibration (Huff & Mandelbaum 2017; Sheldon & Huff 2017), seek an unbiased estimate of the shear either using deeper data as a prior or directly calibrating the measurements using the observed data.

Recent studies highlighted the effect of blending. The blending effect occurs when two or more objects are close together in the image plane, so their light distributions overlap. It introduces biases during both the selection and measurement processes. For example, Hartlap et al. (2011) found that the rejection of recognised blends alters the selection function of the final sample (see also Chang et al. 2013). In some circumstances, blended systems are so close that they appear as single objects. These unrecognised blends increase the shape noise by decreasing the number density and widening the measured ellipticity dispersion (e.g. Dawson et al. 2016; Mandelbaum et al. 2018). Even if the blended objects are below the detection limit, they still introduce correlated noise that affects the detection and measurement of the adjacent bright galaxies (e.g. Hoekstra et al. 2015, 2017; Samuroff et al. 2018), an effect that becomes even more dramatic when the clustering of galaxies is considered (Euclid Collaboration et al. 2019). Given all of these concerns, it is essential for image simulations to contain faint objects and physical clustering features.

More concerns arise when considering a tomographic analysis, which is at the core of current and future weak lensing surveys. From the shear estimate side, the tomographic binning approach introduces further selections that link the shear bias to redshift estimates (K19, MacCrann et al. 2022). From the redshift estimate side, redshift calibration methods need mock photometric catalogues to verify their performance. These mock catalogues must resemble the target data in object selections and photometric measurements, which are challenging to address at the catalogue level (Hoyle et al. 2018; Wright et al. 2020a; van den Busch et al. 2020; DeRose et al. 2022).

All these issues become even more challenging for the KiDS-Legacy analysis, the weak lensing analysis of the complete KiDS. It covers the entire 1350 deg² survey area, a \sim 35% increase over the latest KiDS release (KiDS-DR4, Kuijken et al. 2019). More importantly, thanks to the deeper *i*-band observations and dedicated observations in spectroscopic survey fields, the KiDS-Legacy analysis aims to unleash the power of high-redshift samples (up to a redshift of $z\sim$ 2). The improved statistical power, however, makes a higher demand on the shear and redshift calibrations, including an assessment of the cross-talk between the systematic errors in the shear and redshift estimates.

In this chapter, we present SKiLLS (SURFS-based KiDS-Legacy-Like Simulations), the third generation of image simulations for KiDS following SCHOol (Simulations Code for Heuristic Optimization of *lens*fit, Fenech Conti et al. 2017, hereafter FC17) and COllege (COSMOS-like lensing emulation of ground experiments, K19). By simulating multi-band imaging that includes realistic galaxy evolution and clustering in terms of colour, morphology and number density, SKiLLS allows for the simultaneous measurement of shear and photometric redshifts from the same simulation. This study, therefore, provides the first joint calibration of these two key observables for cosmic

shear analyses. With our approach, we provide a natural solution to address the expected cross-talk between shear and redshift bias, accounting for the impact of blends that carry different shears (Dawson et al. 2016; Mandelbaum et al. 2018; MacCrann et al. 2022). We also release our simulation pipeline, which contains customisable features for general use by other surveys⁷.

The remainder of this chapter is structured as follows. In Sect. 3.2, we build input catalogues for image simulations. Then in Sect. 3.3, we detail the creation and processing of the KiDS-like multi-band images, starting from instrumental setups and ending with photometric catalogues. Section 3.4 reviews our fiducial shape measurement algorithm, *lens*fit (Miller et al. 2007; Kitching et al. 2008; Miller et al. 2013), with some improvements introduced for the KiDS-Legacy analysis. The shear calibration results for the updated *lens*fit measurements are presented in Sect. 3.5, and the sensitivity test is conducted in Sect. 3.6. Finally, we conclude in Sect. 3.7.

Throughout the chapter, we define the complex ellipticity of an object as

$$\epsilon \equiv \epsilon_1 + i\epsilon_2 = \left(\frac{1-q}{1+q}\right) \exp(2i\phi),$$
 (3.1)

where q and ϕ denote the axis ratio and the position angle of the major axis, respectively. In terms of the quadrupole moments of the measured surface brightness Q_{ij} , this definition equals

$$\epsilon = \frac{Q_{11} - Q_{22} + 2iQ_{12}}{Q_{11} + Q_{22} + 2(Q_{11}Q_{22} - Q_{12}^2)^{1/2}}.$$
(3.2)

As stated by Bartelmann & Schneider (2001), this ellipticity definition is convenient because it directly links to the weak lensing shear signal γ via the estimator

$$\gamma = \frac{\sum_{i} w_{i} \epsilon_{i}}{\sum_{i} w_{i}}, \qquad (3.3)$$

where w_i is a weight assigned per object to account for individual measurement uncertainties⁸. Although the cosmic shear analysis uses higher-order statistical measures, such as the two-point correlation functions (e.g. Kaiser 1992), the simple estimator presented in Eq. (3.3) is commonly used for constraining the shear bias from image simulations (e.g. Heymans et al. 2006).

3.2 Input mock catalogues

To generate mock images, we need input catalogues of galaxies and stars with realistic morphology, photometry and clustering. We detail our procedure for building these catalogues in this section. Section 3.2.1 describes how we create the mock galaxy catalogue by combining deep observations with up-to-date cosmological and galactic

⁷https://github.com/KiDS-WL/MultiBand_ImSim.git

⁸Strictly speaking, the expectation value of the ellipticity is $\gamma/(1-\kappa)$, where κ is the convergence. But as $\kappa \ll 1$ in the weak lensing regime, we can safely neglect this term.

simulations. Section 3.2.2 shows how we generate stellar multi-band magnitude distributions from a population synthesis code.

3.2.1 Galaxies: SURFS-Shark simulations with COSMOS morphology

Our input galaxy catalogue, a compilation of simulations and observations, is crafted to achieve a balance between considerable sample volume and realistic galaxy morphology. In the subsequent parts of this section, we first review the simulation part, which includes the clustering and multi-band photometry. Next, we introduce our specifically designed algorithm, which learns galaxy morphology from observations and applies this knowledge to simulated galaxies — a crucial step towards ensuring accurate shear calibration.

Generating synthetic galaxies from simulations

To jointly calibrate the shear and redshift estimates, we must base the image simulations on wide and deep (z > 2) cosmological simulations, where the true redshift is known. In the previous KiDS redshift calibration, van den Busch et al. (2020) used the MICE Grand Challenge (MICE-GC) simulation, an N-body light-cone simulation that covers an octant of the sky (Fosalba et al. 2015a). However, the MICE simulation has a redshift limit of $z\sim1.4$, preventing its use for calibrating the high-redshift samples in the KiDS-Legacy analysis (up to $z\sim2$). Therefore, we switched to another public N-body simulation from the Synthetic UniveRses For Surveys (SURFS, Elahi et al. 2018).

The SURFS simulation we adopted has a box size of $210 \ h^{-1} \rm cMpc$ (cMpc stands for comoving megaparsec), containing 1536^3 particles with a mass of $2.21 \times 10^8 \ h^{-1} \rm M_{\odot}$, and a softening length of $4.5 \ h^{-1} \rm ckpc$ (ckpc stands for comoving kiloparsec). It assumes a $\Lambda \rm CDM$ cosmology with parameters from Planck Collaboration (2016). The final halo catalogues and merger trees are constructed from 200 snapshots starting at redshift z=24, using the phase-space halo-finder code VELOCIRAPTOR (Cañas et al. 2019; Elahi et al. 2019a) and the halo tree-builder code TREEFROG (Elahi et al. 2019b). We refer to Lagos et al. (2018) for details on the building and Poulton et al. (2018) for validating the halo catalogues and merger trees.

The galaxy properties, including the star formation history and the metallicity history, are from an open-source semi-analytic model named Shark⁹ (Lagos et al. 2018). The model parameters are tuned to reproduce the z=0,1 and 2 stellar-mass functions (Wright et al. 2018), the z=0 black hole-bulge mass relation (McConnell & Ma 2013) and the mass-size relations at z=0 (Lange et al. 2016). Any other observables are predictions of the model, which also match well with observations (see Lagos et al. 2018 for more details). As for weak lensing calibration, the most crucial property is the redshift evolution of the galaxy number density (e.g. Hoekstra et al. 2017), which we checked in detail in Appendix 3.A and found it to be sufficient for KiDS.

The light cones from the Shark outputs are created using the code stingray

⁹https://github.com/ICRAR/shark

(Chauhan et al. 2019), an improved version of the code used by Obreschkow et al. (2009). It first tiles the simulation boxes together to build a complex 3D field along the line of sight, then draws galaxy properties from the closest available time-step, resulting in spherical shells of identical redshifts. A possible issue would be the same galaxy appearing once in every box but with different intrinsic properties due to cosmic evolution. To avoid this problem, STINGRAY randomises galaxy positions by applying a series of operations consisting of 90 deg rotations, inversions, and continuous translations. We refer to Chauhan et al. (2019) for more details about the light-cone construction.

The final mock-observable sky covers $\sim 108~\text{deg}^2$ with minimum repetition of the large-scale structure. The sample variance bias caused by the replicating structure is negligible for our direct shear and photometric redshift calibration. Since we learn galaxy morphology from deep observations, our input galaxy sample is still limited mainly by the observational data we have, which only covers $\sim 1~\text{deg}^2$ (see Sect. 3.2.1 for details). We test the robustness of our calibration results against this sample variance bias using the sensitivity analysis detailed in Sect. 3.6.

The multi-band photometry is drawn from a stellar population synthesis technique implemented in the ProSpect¹⁰ and Viperfish¹¹ packages. ProSpect (Robotham et al. 2020) is a high-level package combining the commonly used stellar synthesis libraries with physically motivated dust attenuation and re-emission models; while Viperfish is a light wrapper to aid the interface with the Shark outputs. We refer to Lagos et al. (2019) for detailed predictions, validations and a demonstration that the predicted results agree with observations in a broad range of bands from the far-ultraviolet to far-infrared, without any fine-tuning with observations.

For our purpose, we care most about the nine-band photometry covered by the KiDS+VIKING data, so we compared the synthetic near-infrared and optical magnitude distributions to observations from the COSMOS2015 catalogue (Laigle et al. 2016). Figure 3.1 shows the magnitude distributions of eight filters available in both Shark and COSMOS2015 catalogues, together with an analytical fitting result from Eq. (4) of FC17. The counts in the original simulations are ~35% lower than the observations with some variation between filters. As this affects the blending level and then the shear bias (Hoekstra et al. 2015, 2017), we calibrated the original synthetic photometry for a better agreement. The technical details are presented in Appendix 3.A. In short, we found that the differences in the magnitude distributions stem from the difference in stellar mass-to-light ratio between the simulations and observations. Therefore, we scaled the original Shark magnitudes using a modification factor derived from the stellar mass-to-light ratio difference. The modification is the same for all bands, preserving the intrinsic colours of individual galaxies. The modified magnitudes now agree with the observations within ~3%.

We later noticed that Bravo et al. (2020) proposed a similar fine-tuning method when working with the panchromatic Galaxy And Mass Assembly (GAMA) survey. They used an abundance matching method by comparing the number counts between Shark

¹⁰https://github.com/asgr/ProSpect

[&]quot;https://github.com/asgr/Viperfish

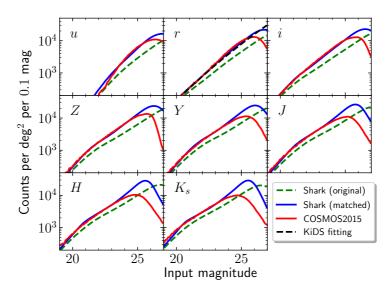


Figure 3.1: Number of galaxies per square degree per 0.1 magnitude in the input apparent magnitudes. The green dashed lines represent data from the original SURFS-Shark mock catalogue, while the blue solid lines indicate the modified results. The red solid lines correspond to the COSMOS2015 observations, factoring in flags for the UltraVISTA area within the COSMOS field after the removal of saturated objects and problematic areas (yielding a 1.38 square degree effective area, as per Table 7 in Laigle et al. 2016). The black dahsed linne depicts the analytical fitting result in the *r*-band, derived from FC17. The *g*-band photometry, absent from the COSMOS2015 catalogue, is not displayed here. We note that the COSMOS2015 catalogue lacks completeness at $K_s \gtrsim 24.5$ (Laigle et al. 2016).

and GAMA after fine binning in redshift and r-band apparent magnitude. They tuned magnitudes for all Shark galaxies with r<21.3 to match the number counts in GAMA. Their modifications are consistent with our results, albeit targeting different magnitude ranges.

Learning galaxy morphology from observations

Simulating galaxies with realistic morphology is essential for accurate shear calibration. Following K19, we represent the galaxy morphology using the Sérsic profile (Sérsic 1963) with three parameters: the effective radius determining the galaxy size (also known as the half-light radius), the Sérsic index describing the concentration of the brightness distribution, and the axis ratio determining the galaxy ellipticity. We learned these structural parameters from deep observations accounting for their mutual correlations and their correlations to galaxy photometry and redshift. Figure 3.2 shows the workflow for the learning algorithm.

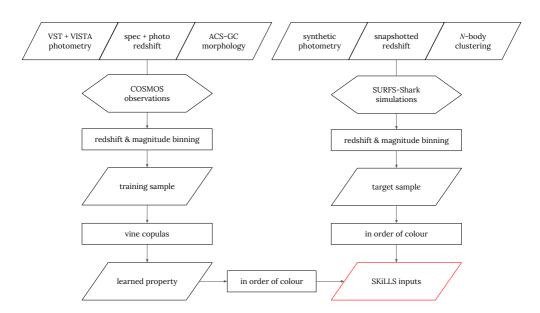


Figure 3.2: Flowchart summarising the algorithm employed in constructing the SKiLLS input mock catalogue. The SKiLLS galaxies inherit the synthetic multi-band photometry and 3D positions from the SURFS-SHARK simulations, while their morphology is learned from the observations in the COSMOS field using a vine-copula modelling-based algorithm (see the detailed explanation in Sect. 3.2.1)

We start with a 'reference' sample comprising morphology, photometry and redshifts from several deep observations. The structural parameters are adopted from the catalogue produced by Griffith et al. (2012), who fitted Sérsic models to the galaxy images taken by the Advanced Camera for Surveys (ACS) instrument on the *Hubble* Space Telescope (HST). We used their results derived from the COSMOS survey and cleaned the sample by only preserving objects with a good fit (FLAG_GALFIT_HI = 0) and reasonable size (half-light radius between 0'.'01 and 10'') to avoid contamination. We note that this catalogue was also used by K19 and proved to be sufficient for KiDS-like simulations.

The r-band photometry is derived from a deep VST-COSMOS catalogue using 24 separate VST observations of the COSMOS field taken from KiDS and the SUpernova Diversity And Rate Evolution (SUDARE) survey (Cappellaro et al. 2015; De Cicco et al. 2019). These observations have a maximum seeing of 0".82, close to the KiDS r-band image qualities. To ensure consistent measurements, we conducted the stacking and detection processes using the same pipeline as the standard KiDS data processing. The stacked image has an average seeing of 0".75 and a total exposure time of 42 120 seconds, which is a factor of ~23 over a standard KiDS observation. The limiting magnitude of the final deep catalogue is more than one magnitude deeper than usual KiDS catalogues. To include colour information, we also used the K_s -band photometry from the COSMOS2015 catalogue (Laigle et al. 2016), as it originates from the UltraVISTA project (McCracken et al. 2012) that shares the same instruments with the VIKING

near-infrared observations.

The redshifts are taken from the catalogue compiled by van den Busch et al. (2022). It contains observations from several spectroscopic and high-quality photometric surveys in the COSMOS field. The spectroscopic redshifts were collected from G10-COSMOS (Davies et al. 2015), DEIMOS (Hasinger et al. 2018), hCOSMOS (Damjanov et al. 2018), VVDS (Le Fèvre et al. 2013), LEGA-C (van der Wel et al. 2016), FMOS-COSMOS (Silverman et al. 2015), VUDS (Le Fèvre et al. 2015), C3R2 (Masters et al. 2017, 2019; Euclid Collaboration et al. 2020; Stanford et al. 2021), DEVILS (Davies et al. 2018) and zCOSMOS (*private comm. from M. Salvato*), while the photometric redshifts were from the PAU survey (Alarcon et al. 2021) and COSMOS2015 (Laigle et al. 2016). For sources with multiple measurements, a specific 'hierarchy' was defined with orders based on the quality of measured redshifts to choose the most reliable redshift estimates (see Appendix A in van den Busch et al. 2022, for details). Given the high quality of the redshift estimates, we treated them as true redshifts.

All catalogues mentioned above overlap in the COSMOS field, so we can combine them by cross-matching objects based on their sky positions. The final reference catalogue has 75 403 galaxies with all the necessary information. It has a limiting magnitude of 27 in the r-band but suffers incompleteness after $m_r \gtrsim 24.5$. We verified that the incompleteness at the faint end does not bias the overall morphological distribution by comparing it to measurements from the Hubble Ultra Deep Field observations (Coe et al. 2006).

We aim to inherit not only the individual distributions of structural parameters but also their mutual dependence and possible correlations with redshifts and magnitudes. To achieve this goal, we developed a learning algorithm based on a statistical inference technique, dubbed vine-copulas (e.g. Joe 2014; Czado 2019). A brief introduction to the technique is presented in Appendix 3.B. In short, a copula-based method models joint multi-dimensional distributions by separating the dependence between variables from the marginal distributions. It is popular in studies concerning dependence modelling, given its flexibility and reliability. In practice, we first divided galaxies into 30×40 bins based on their redshifts and r-band magnitudes. Each bin contains a similar number of reference galaxies. Then in each bin, we built a data-driven vine-copula model from the measured $r - K_s$ colour and morphological parameters using the public pyvinecopulib package¹². The learned vine-copula model can be sampled to produce an arbitrary number of vectors of parameters from the constrained multi-dimensional distributions. We decided to generate the same number of vectors as the available Shark galaxies and assign them to the Shark galaxies in the order of $r - K_S$ colour. This approach allows us to mimic observations from the underlying distributions rather than repeatedly sampling from the measured values.

Figure 3.3 shows the correlations between the magnitude and the two critical structural parameters: half-light radius and ellipticity, in several redshift bins. We see that the learned sample follows the average trends of the reference sample. Figure 3.4 presents two-dimensional contour plots in several magnitude bins to better inspect the

¹²https://github.com/vinecopulib/pyvinecopulib

underlying distributions of morphological parameters. We again see agreements in correlations between the size and ellipticity and between the size and concentration, proving that our copula-based algorithm captures the multi-dimensional dependence from the reference sample.

3.2.2 Stars: Point objects with synthetic photometry

We treated stars as perfect point objects. Their multi-band photometry was obtained from the population synthesis code, Trilegal (Girardi et al. 2005, with version 1.6 and the default model from its website¹³). We generated six stellar catalogues at galactic coordinates evenly spaced across the KiDS footprint to capture the variation of stellar densities between KiDS tiles. Each catalogue spans $10 \, \mathrm{deg^2}$. When simulating a specific tile image covering $1 \, \mathrm{deg^2}$, we selected the stellar catalogue whose central pointing is closest to the target tile, then randomly drew ten per cent of stars from that catalogue as the input. Figure 3.5 shows the r-band magnitude distributions of the six stellar catalogues compared to the catalogue used by the COllege simulations. The broader coverage of stellar densities is noticeable, marking one of the improvements in SKiLLS. Also, stars in SKiLLS have nine-band magnitudes consistently predicted from a library of stellar spectra (see Girardi et al. 2005, for details), while in COllege, stars only have r-band magnitudes.

3.3 KiDS+VIKING nine-band image simulations

This section details the creation and processing of the multi-band mock images. We start with the creation of KiDS-like optical images (Sect. 3.3.1) and VIKING-like infrared images (Sect. 3.3.2), then summarise the SKiLLS fiducial setups in Sect. 3.3.3. We end the section with the measurement of colours and photometric redshifts (Sect. 3.3.4).

3.3.1 KiDS-like optical images

Each KiDS pointing consists of four-band optical images taken with the OmegaCAM camera at the VLT Survey Telescope (Kuijken 2011): u, g, r and i. The r-band images are the primary products used for the shear measurement, while the remaining bands are only for photometric measurements. The science array of the OmegaCAM camera has a $\sim 1^{\circ} \times 1^{\circ}$ field of view covered by 8×4 CCD images, each of size 2048×4100 pixels with an average resolution of 0"214. Although the CCDs are mounted as closely as possible, a narrow gap between the neighbouring CCDs is technically inevitable. The average gap sizes between the pixels of neighbouring CCDs are:

- between the long sides of the CCDs: 1.5 mm (100 pixels)
- central gap along the short sides: 0.82 mm (55 pixels)

¹³http://stev.oapd.inaf.it/cgi-bin/trilegal

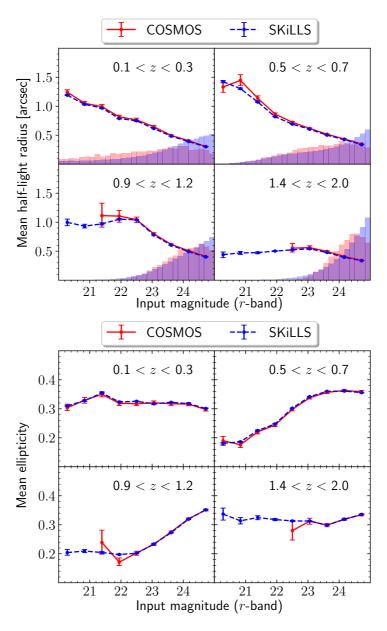


Figure 3.3: Comparison of the overall magnitude-morphology relationships across several redshift bins. The training and target samples are represented by red solid and blue dashed lines, respectively. The top panel illustrates the average half-light radius as a function of r-band magnitude, while the bottom panel depicts the average ellipticity as a function of r-band magnitude. Statistical uncertainties displayed are calculated from 500 bootstrap iterations. The top panel also includes histograms of normalised magnitude distributions, indicating that the bright galaxies at higher redshifts from the simulation contribute minimally to the overall population.

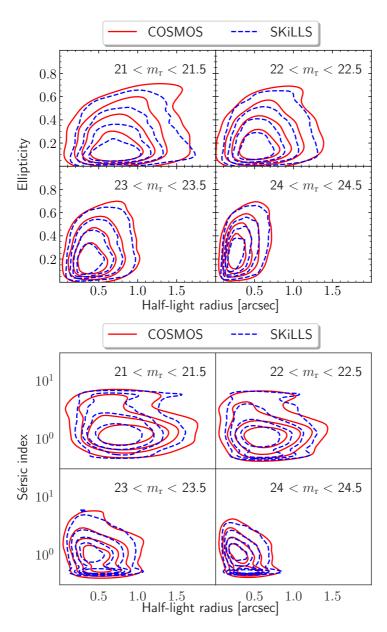


Figure 3.4: Two-dimensional kernel density plots of morphological parameters within distinct magnitude bins. The red solid lines represent the training samples, while the blue dashed lines correspond to the target samples. In the top panel, the correlation between galaxy size and ellipticity is displayed, while the bottom panel elucidates the correlation between galaxy size and Sérsic index. The depicted contour levels correspond to 20%, 40%, 60%, 80%.

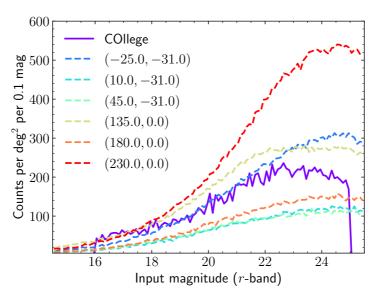


Figure 3.5: The *r*-band input magnitude distributions for the six stellar catalogues utilised by SKiLLS. Labels represent the pointing centres (RA, Dec), with the exception of 'COllege', which signifies the stellar catalogue employed by K19.

• wide gap along short sides: 5.64 mm (376 pixels)

To avoid 'dead zones' caused by these gaps, each tile image incorporates multiple dithered exposures (five in the g, r and i bands, four in the u band). The dithers form a staircase pattern with steps of 25" in RA and 85" in declination to match the gaps between CCDs (de Jong et al. 2013).

KiDS raw observations are processed with two independent pipelines: the Astro-WISE pipeline designed for the photometric measurements (McFarland et al. 2013; de Jong et al. 2015)¹⁴, and the theli pipeline optimised for the shape measurements (Erben et al. 2005; Schirmer 2013; Kuijken et al. 2015)¹⁵. While the former is applied to all four-band observations, the latter is only used for the *r*-band observations, as KiDS only measures galaxy shapes for lensing in the *r*-band images. The main difference between the Astro-WISE and theli pipelines is in the co-addition process, where the former resamples all exposures to the same pixel grid with a uniform 0″20 pixel size, while the latter preserves the original pixels to maintain image fidelity as much as possible.

We kept all these features in mind when generating SKiLLS optical images. We created raw exposures using the GalSim pipeline¹⁶ (Rowe et al. 2015), with galaxies and stars from the mock catalogues described in Sect. 3.2. The underlying canvas mimicked the science array of the OmegaCAM camera, including pixels and gaps.

¹⁴http://www.astro-wise.org/

¹⁵https://www.astro.uni-bonn.de/theli/

¹⁶https://github.com/GalSim-developers/GalSim

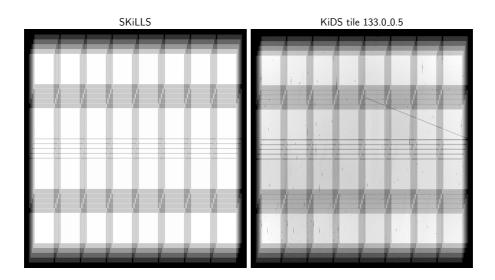


Figure 3.6: Comparison of the THELI weight image generated by SKiLLS (shown in the left panel) with a randomly chosen example from KiDS (depicted in the right panel). The 8 × 4 CCDs span a sky area of approximately one square degree. The shallow regions are attributed to gaps in individual exposures. A similar level of agreement is also observed in the Astro-WISE co-added images.

Galaxies and stars were mapped to the canvas using the gnomonic (TAN) projection of their original sky coordinates. Following the KiDS image processing, we stacked exposures using the SWARP software (Bertin 2010), with the identical setups as in the KiDS pipelines, including Astro-WISE-like images re-gridded to a uniform 0"20 pixel size and theli-like images preserving the original 0"214 pixel size. Figure 3.6 compares a co-added theli weight image from SKiLLS to a randomly selected tile from KiDS. It shows that the SKiLLS images contain the main features of KiDS images, including the gaps and dither patterns, albeit lacking subtle features, such as the inhomogeneous backgrounds between CCDs and masks of satellites.

Besides the image layout, we need information on the pixel noise and point spread function (PSF) to mimic observational conditions. We extracted this information from the fourth public data release of KiDS (KiDS-DR4, or DR4 for short, Kuijken et al. 2019). It has a total of 1006 square-degree survey tiles with stacked ugri images along with their weight maps, masks and source catalogues. We selected a representative sample of 108 tiles and replicated their properties in our image simulations (see Sec. 3.3.3 for details). For the raw pixel noise, we adopted Gaussian distributions with variances estimated from the Astro-WISE weight maps corrected with a boost factor of ~ 1.145 [= $(0.214/0.2)^2$] to account for the re-gridding effect. For the PSF, we used two approaches, depending on the different usages of the images.

For the r-band images from which galaxy shapes are measured, we used the position-dependent PSF models for individual exposures. These PSF models, constructed from

well-identified stars, are in the form of two-dimensional polynomial functions and can recover a PSF image in the pixel grid for any given image position (see Miller et al. 2013; Kuijken et al. 2015; Giblin et al. 2021 for details). In practice, we recovered 32 PSF images for each exposure using the centre positions of the CCD images. The recovered PSF images contain modelling uncertainties, which can introduce artificial spikes when being used to simulate bright stars. Therefore, we applied a cosine-tapered window to the original PSF image to suppress the modelling noise at its outskirts. The two edges of the window function are defined at 5 and 10 times the full-width half-maximum (FWHM) of the target PSF to preserve features in the central region as much as possible. With these recovered PSF images, we can treat the 32 CCD images separately using their own PSFs, a significant improvement from the constant PSF used in previous work. The recovered PSF image is also superior to a Moffat profile as it captures more delicate features of complex PSFs, such as ellipticity gradients.

For other optical bands where only photometry is measured, we still adopted the Moffat profile, given that the photometric measurement is insensitive to the detailed profile of PSF. We estimated the Moffat parameters by modelling bright stars identified in the Astro-WISE images. Since the photometry is measured from the stacked images and is less sensitive to the gentle PSF variation within a given tile, we kept the PSF model invariant for all exposures for simplicity. To alleviate the Moffat fitting bias introduced by the pixelisation of CCD images, we applied the first-order correction to the measured Moffat parameters using image simulations. Specifically, we simulated the pixelated PSF image using measured Moffat parameters and then remeasured them with the same fitting code. The difference between the remeasured and input values is the correction factor and is subtracted from the initially measured value. Our test shows that this correction can suppress the original percent-level bias down to a sub-percent level, which is sufficient for our photometry-related purpose.

3.3.2 VIKING-like infrared images

To improve the accuracy of photometric redshifts, KiDS includes near-infrared (NIR) measurements from the VISTA Kilo-degree Infrared Galaxy (VIKING) survey (Edge et al. 2013). The two surveys share an almost identical footprint. We refer to Wright et al. (2019) for details of the VIKING imaging and its usage in KiDS. Briefly, the VIKING data have three levels of products: exposures, paw-prints, and tiles. Given the complex NIR backgrounds, the VIKING survey first takes multiple exposures in quick succession with small jitter steps for reliable estimation of the noisy background. These exposures are then stacked together to create the second level of product: the 'paw-print'. A paw-print still contains gaps between individual detectors, so six paw-prints with a dither pattern are used to produce a contiguous tile image. However, these co-added tiles have non-contiguous PSF patterns caused by the large dithers between successive paw-prints. Therefore, in the KiDS+VIKING analyses, photometry is done on individual paw-prints instead of the co-added tiles. The dither pattern of paw-prints causes multiple flux measurements per source (typically four in the case of the *J*-band and two in the other

bands). The final flux estimate for each source is a weighted average of the individual measurements with the weights derived from individual flux errors.

Given the complexity of the VIKING observing strategy, we simplified the NIR-band observations in SKiLLS with single images per square degree of KiDS tile. To compensate for the simplified images, we considered the overlap between individual paw-prints when estimating the observational conditions. As we show in Sect. 3.3.4, this simplified approach can still achieve realistic photometry, which is the only important quality we seek from the NIR-band images.

Specifically, we created a 'flat-field image' for each paw-print with the same size and pixel scale. Its pixel value equals the absolute standard deviation of the background pixel values on the corresponding paw-print. For each KiDS pointing, we selected all VIKING paw-prints that overlap in the given one square-degree sky area and stacked their flat-field images with shifts accounting for the different sky pointings of the paw-prints. We took the median pixel value of the co-added flat-field image as the final pixel noise of the corresponding KiDS pointing. In doing so, we captured various overlapping VIKING paw-prints in individual KiDS pointings. Following the typical situations of the KiDS+VIKING data (Wright et al. 2019), we only preserved KiDS pointings with at least two paw-prints in the $ZYHK_s$ -bands and at least four paw-prints in the J-band. This requirement reduced the number of pointings from 1006 to 979, which is still plentiful for our purpose. As for the PSF, we employed a constant Moffat profile for each KiDS pointing. The PSF FWHM is a weighted average from overlapping VIKING paw-prints with the weights determined by their noise levels. In order to determine the Moffat concentration index for a given FWHM value, we fitted Moffat profiles to bright stars in some representative paw-prints. The Moffat fitting bias introduced by the pixelisation is corrected using the same method introduced in Sect. 3.3.1. We found the relationship between the Moffat index n and FWHM (arcsec) in VIKING images to be: ln(n) = 66.56 exp(-6.36 FWHM) + 0.90. This empirical formula is used to pair each FWHM with a unique Moffat index.

3.3.3 SKiLLS fiducial setup

Since we have 108 deg² of Shark galaxies as described in Sect. 3.2.1, we selected 108 KiDS pointings for the SKiLLS fiducial run. Figure 3.7 shows the sky locations of the selected 108 tiles along with the 979 KiDS-DR4 tiles that have the nine-band noise and PSF information. Clusters of the selected blocks pair with the six stellar catalogues generated from Trilegal so that SKiLLS captures the stellar density variation across the whole KiDS survey (see Sect. 3.2.2).

Figure 3.8 compares the *r*-band noise and PSF properties between the SKiLLS selected tiles and all usable KiDS-DR4 tiles. We measured the PSF size and ellipticity using the weighted quadrupole moments with a circular Gaussian window of dispersion 2.5 pixels, the typical galaxy size in the KiDS sample. The PSF size is defined as

$$r_{\text{PSF}} \equiv (Q_{11}Q_{22} - Q_{12}^2)^{1/4},$$
 (3.4)

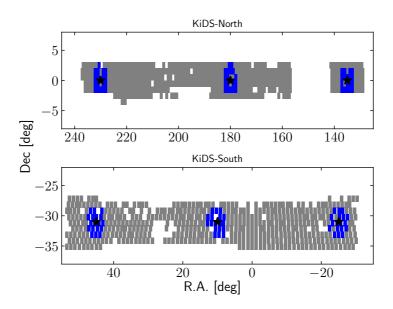


Figure 3.7: Sky distribution of the KiDS-DR4 tiles. The tiles marked in blue are incorporated in the SKILLS fiducial run, accounting for a total of 108 tiles. The grey blocks represent all KiDS-DR4 tiles, amounting to 979, for which nine-band noise and PSF information are available. The black stars signal the centres of the stellar catalogues generated using Trilegal (Girardi et al. 2005).

where Q_{ij} are the weighted quadrupole moments, and the PSF ellipticity is defined by Eq. (3.2). Figure 3.8 shows that the selected tiles represent the KiDS-DR4 data well. Because we vary PSF for individual CCD images and exposures, the 108 SKiLLS images cover 17 280 different PSF models, a significant extension of the 65 PSF models used by FC17 and K19. That also explains the smooth distributions of the PSF parameters. Figure 3.9 shows similar comparisons for other bands. Again we see fair agreements across all bands. As KiDS-DR4 already covers ~75% of the whole survey, we expect a similar agreement to the KiDS-Legacy data. The wide coverage of the noise and PSF properties also makes the SKiLLS results more robust than previous simulations and simplifies sensitivity tests (see Sect. 3.6 for details).

3.3.4 Photometry and photometric redshifts

With the simulated multi-band images, we can measure colours and estimate photometric redshifts (photo-zs) for simulated galaxies using the same tools developed in KiDS with minor adjustments.

For galaxy colours, we used the GAAP (Gaussian Aperture and PSF) pipeline (Kuijken et al. 2015, 2019). It provides accurate multi-band colours by accounting for PSF differences between filters and optimises signal-to-noise ratio (S/N) by down-

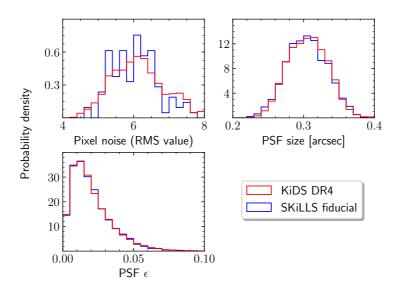


Figure 3.8: Comparison of normalised histograms for pixel noise (top left), PSF size (top right), and PSF ellipticity (bottom left) between KiDS-DR4 (depicted in red) and SKiLLS (in blue) within the *r*-band images. Both PSF size and ellipticity are measured from the recovered PSF image employing a circular Gaussian window of a sigma value of 2.5 pixels.

weighting the noise-dominated outskirts. The latter is possible because the photo-z estimation only needs the ratio of the fluxes from the same part of a galaxy in the given bands rather than the total light. A prerequisite for the GAAP pipeline is a detection catalogue with source positions and aperture parameters, which we measured from the THELI-like r-band images using the SEXTRACTOR code (Bertin & Arnouts 1996). Once the detection catalogue is ready, we can obtain the list-driven photometry by running the GAAP algorithm on the u, g, r and i Astro-WISE-like images and the Z, Y, J, H and K_s simple images. In short, the GAAP method includes three major steps:

- 1. Homogenising PSFs by convolving the whole image with a spatially variable kernel map modelled from high S/N stars. The resulting image has a simple Gaussian PSF, for which estimating the PSF-independent Gaussian aperture flux is possible. The main side effect is that the convolution process introduces correlated noise between neighbouring pixels, complicating the estimation of measurement uncertainties. GAAP handles this by tracking the noise covariance matrix through the whole process.
- 2. Defining an elliptical Gaussian aperture function for each source using the size and shape parameters measured by SEXTRACTOR on the *r*-band detection images. In practice, users must customise the minimum and maximum GAAP aperture sizes

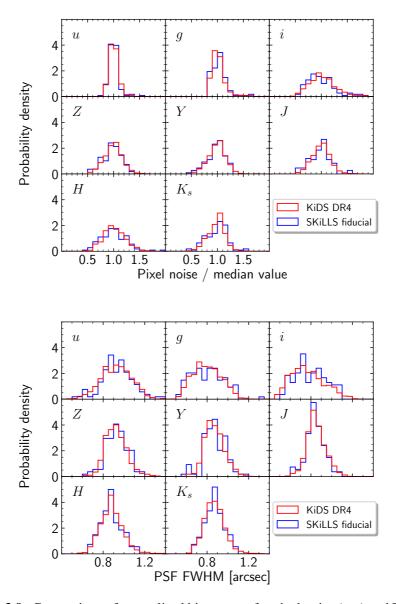


Figure 3.9: Comparison of normalised histograms for pixel noise (top) and PSF FWHM (bottom) across bands used solely for photometry, contrasting KiDS-DR4 (red) with SKiLLS (blue). The equivalent comparisons for the lensing r-band images are illustrated in Fig. 3.8. To facilitate comparison in the same range, pixel noise values are normalised by their median values in each band.

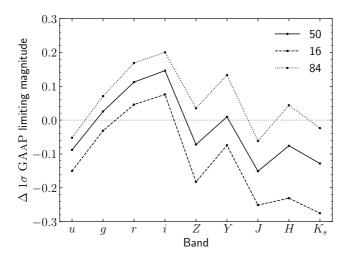


Figure 3.10: Differences in the median 1σ GAAP limiting magnitudes across the nine bands (simulation - data). The three lines represent the 16, 50 and 84 percentiles derived from the 108 tiles incorporated in the SKiLLS fiducial run. The larger scatters in the NIR bands can be partially attributed to the adoption of a simplified simulation strategy.

to balance the S/N and the effect of blending. Following the KiDS fiducial setup, we set the maximum aperture to 2" to avoid contamination from neighbouring sources. We conducted two separate runs by setting the minimum aperture to 0".7 and 1".0. When used as the input for the photo-z estimation, a source-by-source decision was made to optimise the flux errors across the nine bands (see Kuijken et al. 2019 for details).

3. Performing the aperture photometry on the PSF-Gaussianised images for each band using the defined aperture functions. It is worth stressing that GAAP aims to provide robust colours for the high S/N parts of galaxies; it underestimates the total fluxes for extended sources by design.

Figure 3.10 compares the nine-band 1σ GAAP limiting magnitudes between the KiDS-DR4 data and SKiLLS fiducial results. We calculated the median limiting magnitudes for tiles in both KiDS and SKiLLS and then compared their differences. We see a general agreement for all the bands, verifying our noise and PSF modelling. Noticeably, even for the NIR bands where we simplified the VIKING observations with single images, the differences are still tolerable, albeit with larger uncertainties. Figure 3.11 compares the GAAP photometric distributions between the simulation and data. Once again, we see a decent agreement in both magnitude and colour distributions.

For the photo-z estimation, we implemented the public Bayesian Photometric Redshift (BPZ, Benítez 2000) code with the re-calibrated template set from Capak (2004) and the Bayesian redshift prior from Raichoor et al. (2014). We closely followed the

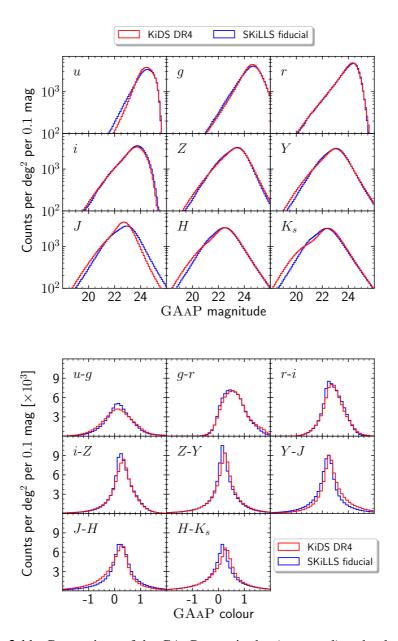


Figure 3.11: Comparison of the GAAP magnitudes (top panel) and colours (bottom panel) for KiDS-DR4 (red) and SKiLLS (blue). The displayed results incorporate all galaxies possessing valid photometric measurements (with the GAAP flags equating to zero across nine bands). No shape-measurement-related filters have been applied at this stage.

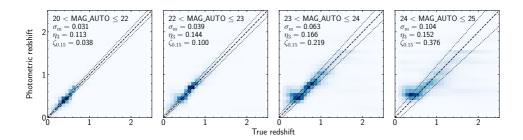


Figure 3.12: Photometric versus true redshifts distributed across several apparent r-band magnitude bins. The annotated statistics are: the normalised median-absolute-deviation (σ_m) of the variable $\Delta z/(1+z)$, the proportion of sources with $|\Delta z/(1+z)| > 3\sigma_m$ (η_3) and the proportion of sources with $|\Delta z/(1+z)| > 0.15$ $(\zeta_{0.15})$. The dashed lines represent the one-to-one correlation, while the dotted lines indicate $|\Delta z/(1+z)| = \sigma_m$.

settings in the KiDS-DR4 analysis (Kuijken et al. 2019) unless it conflicts with the simulation input. For example, we set ZMAX to 2.5, the limiting redshift of SKiLLS galaxies, instead of 7.0 as in the data. We tested the choice of ZMAX in the simulations and found that only 0.1% of the test sample resulted in estimates differing more than 0.1, which means most of the objects have similar photo-z estimates and end up in the same tomographic bins for these two choices. Moreover, the Shark photometry in the u, g, r, i and Z bands is based on the Sloan Digital Sky Survey (SDSS) photometric system, which is slightly different from the KiDS/VIKING system (Kuijken et al. 2019). We corrected these slight differences in the measured GAAP magnitudes in order to use the KiDS/VIKING filters to run the BPZ code. The detailed procedures and comparisons are described in Appendix 3.C. Overall, the modification is minor and has a negligible impact on the magnitude, colour distributions, and final shear biases. Still, it improves the agreement between the simulation and the data in the photo-z distributions. Unless specified otherwise, we base our fiducial results on the transformed photometry.

Figure 3.12 compares the estimated photo-z to the true redshift from the input SURFS-Shark simulations in several measured magnitude bins. It shows the photo-z vs. true redshift distributions, along with annotated statistics based on the distributions of $(z_{\rm B}-z_{\rm true})/(1+z_{\rm true}) \equiv \Delta z/(1+z)$ values. We see the BPZ code works well in SKiLLS and is at the same level as in KiDS (Wright et al. 2019). More detailed verification of the SKiLLS photo-z performance is presented in the companion redshift calibration paper (van den Busch et al., in prep).

As for the redshift calibration, our end-to-end approach, which starts with image simulation followed by object detection, PSF homogenisation, forced multi-band photometry, and photo-z estimation, is a significant improvement compared to previous catalogue-level simulations (e.g. Hoyle et al. 2018; van den Busch et al. 2020; DeRose et al. 2022). The image-simulation-based approach not only yields more realistic observational uncertainties but also naturally accounts for the blending effect, which

is hard to address at the catalogue level. As for the shear calibration, these photo-*z* estimates are essential for performing tomographic selections (K19). Our approach that directly measures the photo-*z*s from simulated images accounts for various measurement uncertainties of photo-*z*s, hence a tomographic selection consistent with how it is done in the data. Moreover, using the same mock catalogue in both shear and redshift calibration unites these two long-separated processes in the KiDS-Legacy analysis.

3.4 Shape measurements with the updated *lens* fit

The primary task of any weak lensing survey is to measure the shapes of galaxy images. Previous KiDS analyses tackled this task using a likelihood-based code, dubbed *lens*fit (Miller et al. 2007; Kitching et al. 2008; Miller et al. 2013). It is the default shape measurement algorithm for the KiDS-Legacy analysis, with some updates described in this section. We test SKiLLS using this updated *lens*fit code¹⁷.

3.4.1 The self-calibration version of *lens* fit

The *lens*fit code, first developed for CFHTLenS (Heymans et al. 2012), follows a Bayesian model-fitting approach. We refer to Miller et al. (2013) for its detailed formalism. In brief, it first performs a joint fit to individual exposures using a PSF-convolved galaxy model, which yields a likelihood distribution of seven parameters: 2D position, flux, scalelength, bulge-to-total flux ratio and complex ellipticity. Then it deduces the ellipticity parameters from the likelihood-weighted mean values by marginalising other parameters with priors as described by Miller et al. (2013). For each ellipticity estimate, an inverse-variance weight is also determined from (Miller et al. 2013)

$$w_{i} \equiv \left[\frac{\sigma_{\epsilon, i}^{2} \epsilon_{\max}^{2}}{2\epsilon_{\max}^{2} - 4\sigma_{\epsilon, i}^{2}} + \sigma_{\epsilon, \text{pop}}^{2} \right]^{-1}, \qquad (3.5)$$

where $\sigma_{\epsilon, i}$ is the uncertainty of the measured ellipticity, $\sigma_{\epsilon, pop}$ is the ellipticity dispersion of the galaxy population (intrinsic shape noise), and ϵ_{max} is the maximum allowed ellipticity in the *lens*fit model-fitting. As for KiDS data, we adopted $\sigma_{\epsilon, pop} = 0.253$ and $\epsilon_{max} = 0.804$.

The code has evolved as KiDS progressed. The most significant is a self-calibration scheme for noise bias, as detailed in FC17. The pixel noise in a given image skews the likelihood, which biases the estimate of individual galaxy ellipticities. It is a complex function of the signal-to-noise ratio, galaxy properties and PSF morphology, making it difficult to predict accurately. Thus, *lens*fit conducts an approximate correction using the measurements themselves, that is a self-calibration. The basic idea is to simulate a test galaxy with parameters measured from the first run, then remeasure the test galaxy using the same pipeline. The difference between the remeasured and input values serves as a

¹⁷Nevertheless, we note that SKiLLS can also calibrate other algorithms, such as the KiDS Metacalibration catalogue (Yoon et al., in prep.).

correction factor for the corresponding parameter. Since its introduction, self-calibration has been a standard part of *lens*fit, given its promising overall performance (Mandelbaum et al. 2015; FC17; K19). We keep this feature for the KiDS-Legacy analysis.

3.4.2 Updates for KiDS-Legacy analysis

A long-standing mystery of all previous *lens*fit analyses has been the presence of a small but significant residual bias in ϵ_2 that is uncorrelated with the PSF and the underlying shear (Miller et al. 2013; Hildebrandt et al. 2016; Giblin et al. 2021). We now understand that this feature arises from an anisotropic error in the original likelihood sampler, which has been corrected in our algorithm. However, we found that this correction inadvertently increases the fraction of residual PSF contamination in the weighted average signal (see the discussion in Giblin et al. 2021). Besides, object selection and galaxy weights are also known to introduce bias (e.g. Kaiser 2000, Bernstein & Jarvis 2002, Hirata & Seljak 2003, Jarvis et al. 2016 and FC17). These selection biases can be more severe than the raw measurement bias and hence cannot be ignored even for a perfect self-calibration measurement algorithm.

FC17 presented a method to isotropise weights using an empirical correction scheme, which has been adopted in previous KiDS studies to mitigate these biases. Unfortunately, we found this approach to be insufficient for the improved *lens*fit algorithm. Furthermore, we found the approach to be sensitive to the sample volume, and therefore hard to apply consistently to the data and simulations. So, we introduce a new empirical correction scheme that mitigates the PSF contamination to the weighted shear signal.

Weight correction

We begin by investigating the impact of PSF leakages on the reported weight. For galaxies of similar surface brightness, those oriented in the same direction as the PSF tend to exhibit a higher integrated signal-to-noise ratio compared to those perpendicular to the PSF. This orientation bias introduces an asymmetry in the measurement variance (the $\sigma_{\epsilon,i}^2$ term in Eq. 3.5), which can be quantified using a first-order linear function:

$$S_i = \alpha_S \epsilon_{\text{PSF}, i, \text{proj}} + \mathcal{N} \left[\langle S \rangle, \sigma_S \right] . \tag{3.6}$$

Here, $S_i \equiv \sigma_{\epsilon, i}^2$ signifies the measurement variance, while $\epsilon_{\text{PSF}, i, \text{proj}}$ denotes the scalar projection of the PSF ellipticity along the direction of the galaxy ellipticity. The term α_S quantifies the degree of PSF contamination in the measurement variance. The noise, denoted by $\mathcal{N}\left[\langle S \rangle, \sigma_S\right]$, is assumed to follow a Gaussian distribution with a mean of $\langle S \rangle$ and a standard deviation of σ_S .

Following FC17, we estimate the PSF contamination as a function of the integrated signal-to-noise ratio (ν_{SN}) reported by *lens*fit and the resolution, which is defined as

$$\mathcal{R} \equiv \frac{r_{\text{PSF}}^2}{r_{\text{PSF}}^2 + r_{ab}^2} \,, \tag{3.7}$$

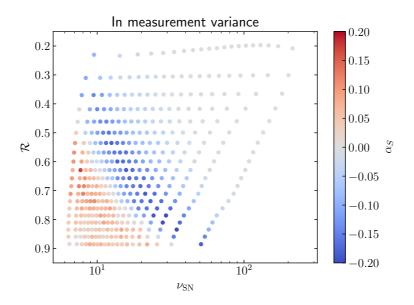


Figure 3.13: PSF leakage in the measurement variance as a function of signal-to-noise ratio and the resolution parameter \mathcal{R} . We note that, according to the definition in Eq. (3.7), a larger \mathcal{R} corresponds to a lower resolution.

where $r_{ab} \equiv r_{\rm e}\sqrt{q}$ is the circularised galaxy size with $r_{\rm e}$ and q denoting the scalelength along the major axis and the axis ratio, respectively. The PSF size $r_{\rm PSF}$ is defined by Eq. (3.4). By construction, the resolution $\mathcal R$ has a value between 0 and 1, with a larger value corresponding to a more poorly resolved object.

When estimating α_S , we first divide galaxies into an irregular 20×20 grid of ν_{SN} and \mathcal{R} , each containing the same number of objects. Then in each bin, we perform a linear regression using Eq. (3.6) to measure α_S . Figure 3.13 shows the measurements for the KiDS-DR4 re-run with the updated *lens*fit. It demonstrates a clear correlation between the estimated α_S and the ν_{SN} and \mathcal{R} . We derive the corrected measurement variance for individual galaxies through $\sigma_{\epsilon, i, \text{corr}}^2 = \sigma_{\epsilon, i}^2 - \alpha_S \epsilon_{PSF, i, \text{proj}}$, where the value of α_S is determined based on which ν_{SN} - \mathcal{R} bin the target galaxy is assigned to. The corrected *lens*fit weight is then calculated with

$$w_{\text{corr, }i} \equiv \left[\frac{\sigma_{\epsilon, i, \text{ corr}}^2 \epsilon_{\text{max}}^2}{2\epsilon_{\text{max}}^2 - 4\sigma_{\epsilon, i, \text{ corr}}^2} + \sigma_{\epsilon, \text{ pop}}^2 \right]^{-1}, \tag{3.8}$$

following Eq. (3.5). We verified that this approach is sufficient to remove the overall weight bias and is robust against the binning scheme.

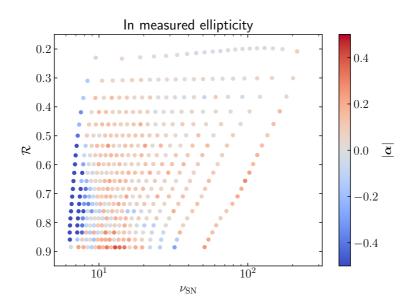


Figure 3.14: PSF leakage in the measured ellipticity after the weight calibration, plotted as a function of signal-to-noise ratio and the resolution parameter \mathcal{R} . We note that, according to the definition in Eq. (3.7), a larger \mathcal{R} corresponds to a lower resolution.

Ellipticity correction

In addition to the weight bias, there is still some residual PSF leakage in the measured ellipticity because of the residual noise bias and selection effects. To first order, this residual PSF bias can be formulated as

$$\epsilon_{\text{obs}, i} = \epsilon_{\text{true}, i} + \alpha \epsilon_{\text{PSF}, i} + c + \mathcal{N} [0, \sigma_{\epsilon}],$$
 (3.9)

where $\epsilon_{\text{obs}, i}$ is the measured ellipticity, $\epsilon_{\text{true}, i}$ is the underlying true ellipticity, α is the fraction of the PSF ellipticity $\epsilon_{\text{PSF}, i}$ that leaks into the measured ellipticity, and c is an additive term uncorrelated with the PSF. $\mathcal{N}\left[0, \sigma_{\epsilon}\right]$ denotes the noise in individual shape measurements, which are assumed to follow a Gaussian distribution of mean 0 and standard variation σ_{ϵ} . We note that all parameters in Eq. (3.9) are complex numbers ($\alpha = \alpha_1 + i\alpha_2$). We focus on the α term, as the c term with the improved likelihood sampler is now small in practice, and the $\mathcal{N}\left[0, \sigma_{\epsilon}\right]$ vanishes for an ensemble of galaxies.

Like the weight bias correction, we first estimate α in the 20×20 grid of ν_{SN} and \mathcal{R} using a linear regression of Eq. (3.9). Figure 3.14 shows the amplitude of α in the 2D ν_{SN} and \mathcal{R} plane. We see modest values in most situations, except for the low ν_{SN} cases, where it drops abruptly to negative values. We confirmed that the negative tail is mainly from the selection effects by measuring the PSF leakage using the input ellipticity in simulations. This non-trivial negative tail prevents us from using the direct correction

approach introduced in the weight bias correction section. Therefore, we propose a hybrid approach, with a fitting procedure for the overall trend and a direct correction for residuals. Specifically, we first fit the measured α as a function of ν_{SN} and \mathcal{R} , using a function of the form

$$\alpha_{\rm p}(\nu_{\rm SN}, \mathcal{R}) = a_0 + a_1 \nu_{\rm SN}^{-2} + a_2 \nu_{\rm SN}^{-3} + b_1 \mathcal{R} + c_1 \mathcal{R} \nu_{\rm SN}^{-2},$$
 (3.10)

whose coefficients are constrained using the weighted mean results from the 20×20 grid. Then, we correct the raw measurements of individual galaxies using $\epsilon_{\text{obs, }i, \text{ tmp}} = \epsilon_{\text{obs, }i} - \alpha_p(\nu_{\text{SN, }i}, \mathcal{R}_i)$ $\epsilon_{\text{PSF, }i}$, where the polynomial $\alpha_p(\nu_{\text{SN, }i}, \mathcal{R}_i)$ is determined from the target galaxy's $\nu_{\text{SN, }i}$ and \mathcal{R}_i . After removing the overall trend, we use the corrected $\epsilon_{\text{obs, }i, \text{ tmp}}$ to measure the residual α_r , which changes mildly across the 2D ν_{SN} and \mathcal{R} plane. Therefore, we can conduct the direct correction through $\epsilon_{\text{obs, }i, \text{ corr}} = \epsilon_{\text{obs, }i, \text{ tmp}} - \alpha_r \epsilon_{\text{PSF, }i}$, where the values of α_r for individual galaxies are determined based on which ν_{SN} - \mathcal{R} bin they are assigned. This two-step approach balances performance and robustness. We verified that the corrected measurements have negligible PSF leakages and the results are robust against the binning scheme.

3.4.3 Comparison between KiDS and SKiLLS

We applied the updated *lens*fit code to KiDS-DR4 and SKiLLS *r*-band images. The object selections after the measurements are detailed in Appendix 3.D. In short, we largely followed the selection criteria proposed in Hildebrandt et al. (2017), with an additional resolution cut introduced to mitigate the PSF contamination. We applied the same selections to the KiDS data and SKiLLS simulated catalogue to ensure a consistent selection effect, even though SKiLLS does not contain artefacts like asteroids and binary stars.

Figure 3.15 compares the weighted distributions of some critical observables reported by the updated *lens*fit. The SKiLLS results match the KiDS-DR4 data reasonably well. We also checked the properties of the close pairs. Specifically, we show the magnitude difference and the projected distance between close pairs in the measured catalogues. Both properties agree well between the data and simulations, implying SKiLLS has realistic clustering features. These realistic neighbouring properties are essential for an accurate shear calibration, especially when considering the shear interference between blended objects (see Sect. 3.5 for details).

3.5 Shear biases for the updated lensfit

The central task of image simulations is to quantify the average shear bias for a selected source sample. This is done by comparing the inferred shear γ_{obs} , to the input shear γ_{input} , which have a linear correlation to the first order (Heymans et al. 2006)

$$\gamma_{\text{obs}} = (1 + m) \gamma_{\text{input}} + c , \qquad (3.11)$$

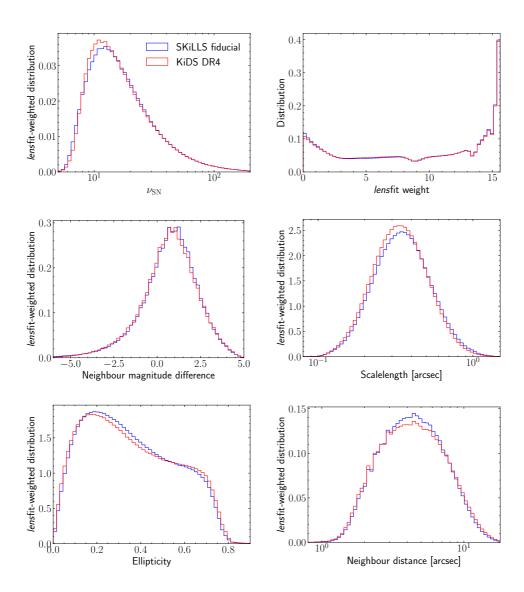


Figure 3.15: Comparison of the updated *lens*fit measurements between KiDS (represented in red) and SKiLLS (in blue). All distributions are normalised using *lens*fit weights, with the exception of the *lens*fit weight distribution itself. Neighbour properties are based on the nearest neighbour identified in the measured catalogue. The magnitude difference is calculated by subtracting the primary target's neighbour magnitude from that of the neighbour. The absence of closely paired galaxies with a distance less than ~1 arcsecond results from the conservative blending cut used by KiDS (refer to Appendix 3.D). This cut effectively mitigates the most severe blending bias.

3

where m is known as the multiplicative bias, and c is the additive bias. The simulation-based calibration focuses on the multiplicative bias, as the additive bias is usually corrected empirically (for example, the correction scheme proposed in Sect. 3.4.2). So we use the term 'shear bias' and 'multiplicative bias' interchangeably throughout the paper. We note that all parameters in Eq. (3.11) are in complex forms, such as $m = m_1 + im_2$. However, we found m_1 and m_2 to be consistent in our analysis, so unless specified, we only report the amplitude m.

The shear calibration methodology keeps evolving as our understanding of systematics deepens. Early studies demonstrated that the shear bias correlated with galaxy properties and PSFs, especially the signal-to-noise ratio and resolution (e.g. Miller et al. 2013; Hoekstra et al. 2015; Mandelbaum et al. 2018; Samuroff et al. 2018). So the first lesson is to avoid using one averaged result from the whole simulation as a scalar calibration to the entire data unless the simulations perfectly represent the data. A natural procedure then attempts to estimate the shear bias as a function of the galaxy and PSF properties (e.g. Miller et al. 2013; Jarvis et al. 2016). Nevertheless, we can only derive the relation of the bias to the noisy, measured properties, as the true properties are unknown in actual data. FC17 found that the relation derived from the measured properties introduces biases because of the correlations between observed quantities. an effect referred to as the 'calibration selection bias'. So the second lesson is that we should be cautious about object-based shear calibrations that rely on the relation to the noisy properties. That is why the recent simulations try to resemble the data and only provide a mean correction for an ensemble of galaxies (e.g. K19). The latest lesson, stressed by MacCrann et al. (2022), is the interplay between shear estimates of blended objects at different redshifts, a higher-order effect that the traditional constant shear simulations cannot capture. It becomes more important as the precision of surveys improves.

Our shear calibration method builds on all these lessons. We created constant shear simulations following the previous KiDS tomographic calibration method but with improvements to the photo-*z* estimates by taking advantage of the simulated multi-band images (Sect. 3.5.1). Using additional blending-only variable shear simulations, we applied a correction to account for the interplay between blends containing different shears (Sect. 3.5.2). When testing the PSF modelling algorithm in image simulations, we detected a small but noticeable change of shear bias, which was also corrected in our fiducial results (Sect. 3.5.3).

3.5.1 Results from the constant shear simulations

Our constant shear simulations largely followed FC17 and K19 with some simplifications for better usage of computational resources. Table 3.1 lists the main changes we made compared to our predecessor. Given the 108 deg² of unique synthetic galaxies we built in Sect. 3.2, we mimicked 108 KiDS pointings, where we vary the PSF, noise level and stellar density as detailed in Sect. 3.3. To reduce the shape noise, we copied each tile image with galaxies rotated by 90 degrees. We created four sets of constant shear

simulations with input shear: (0.0283, 0.0283), (0.0283, -0.0283), (-0.0283, -0.0283), (-0.0283, 0.0283). The total simulated area is 864 (= $108 \times 4 \times 2$) deg², which is equivalent to ~5170 deg² after accounting for the shape noise cancellation (= $864 \times (\sigma_{\epsilon,raw}/\sigma_{\epsilon,SNC})^2$, where $\sigma_{\epsilon,raw}$ and $\sigma_{\epsilon,SNC}$ denote the weighted dispersion of the mean input ellipticities before and after the shape noise cancellation), which is roughly four times the final KiDS-Legacy area.

For a tomographic analysis, we need to estimate the bias for each redshift bin separately, given that the galaxy properties vary between bins. This requires photo-z estimates for the simulated galaxies. For SKiLLS, we can follow the KiDS processing steps to directly measure photo-zs, thanks to the simulated nine-band images. We conducted the detection from the Theli-like r-band images, the PSF Gaussianisation and forced multi-band photometry using the GAAP pipeline, and the photo-z estimates with the BPZ code (see Sect. 3.3.4 for details). This consistent data processing ensures that SKiLLS embraces realistic photometric properties, marking one of the most significant improvements over the previous image simulations.

Table 3.1: Differences between the COllege (K19) and SKiLLS simulations.

		COllege (K19)	SKiLLS (this work)
Galaxies	Morphology	Sérsic models with parameters taken directly from the HST-ACS measurements (Griffith et al. 2012)	Sérsic models with parameters learned from the HST-ACS measurements (Sect. 3.2.1)
	Photometry	Single-band magnitudes from the Subaru r^+ -band observations	Nine-band synthetic magnitudes based on a semi-analytic model (Sect. 3.2.1)
	Depth	Limited by the HST-ACS measurements	Extending to 27th magnitude in the <i>r</i> band
	Position	Based on the observed locations in the COS-MOS field	Based on the SURFS <i>N</i> -body simulations (Elahi et al. 2018)
Stars	Photometry	Single-band synthetic magnitudes from the Besançon model (Robin et al. 2003; Czekaj et al. 2014)	Nine-band synthetic magnitudes from the TRILEGAL model (Gi- rardi et al. 2005)
Images	Band	the r-band images only	the full nine-band images
			Continued on next page

Continued on next page

Table 3.1 continued from previous page

Tuble 5.	1 continued from previous	
	COllege (K19)	SKiLLS (this work)
Layout	32 CCDs with even gaps in between	32 CCDs with variable gaps as in the actual camera (Fig. 3.6)
PSF	13 sets of spatially constant Moffat profiles, with each containing five different models corresponding to the five exposures	108 sets of spatially varying polynomial models, with each containing 5 × 32 different models
Noise	One fixed noise level for all tiles	108 different noise levels
Stack	Only THELI-like stacks for shape measurements	Both THELI-like and As- TRO-WISE-like stacks for shape and photomet- ric measurements, re- spectively
Shape	From the self-calibration version of <i>lens</i> fit with the weight bias correction of FC17	From the updated <i>lens</i> fit with the AlphaRecal method detailed in Sect. 3.4.2
photo-z	Assigned with the KiDS observations of the COSMOS field	Measured from the simulated nine-band images following the KiDS photometric processing steps (Sect. 3.3.4)
nce	Identical input catalogues of galaxies and stars for all the 13 realisations	Different galaxy catalogues for the 108 realisations and six stellar catalogues for the selected sky blocks (Fig. 3.7)
	Layout PSF Noise Stack Shape photo-z	COllege (K19) Layout 32 CCDs with even gaps in between PSF 13 sets of spatially constant Moffat profiles, with each containing five different models corresponding to the five exposures Noise One fixed noise level for all tiles Stack Only THELI-like stacks for shape measurements Shape From the self-calibration version of lensfit with the weight bias correction of FC17 photo-z Assigned with the KiDS observations of the COSMOS field Identical input catalogues of galaxies and stars for all the 13 reali-

Continued on next page

Table 3.1 continued from previous page		
	COllege (K19)	SKiLLS (this work)
Input shears ^a	Eight sets of constant shears	Four sets of constant shears in the baseline simulations and a variable shear field for the blended objects (Appendix 3.E)
Shape noise cancellation ^b	Each tile has three counterparts with galaxies rotated by 45, 90 and 135 degrees	Each tile has one counterpart with galaxies rotated by 90 degrees
Total simulated area	416 deg ²	864 deg ² in the constant shear simulations plus 7776 deg ² of blending- only simulations for the correction of the 'shear interplay' effect (Sect. 3.5.2)

As shown in Fig. 3.15, SKiLLS matches KiDS generally well but not perfectly. K19 argued that an accurate estimate of the shear bias must account for any mismatches between the simulations and the target data. Therefore, we followed FC17 and K19 to reweight the simulation estimates using the *lens*fit reported v_{SN} and resolution factor R(Eq. 3.7). Specifically, for each tomographic bin, we first divided simulated galaxies into 20 × 20 bins of v_{SN} and R, each containing equal lensfit weight. Then we estimated the multiplicative bias for each ν_{SN} - \mathcal{R} bin using Eq. (3.11). Galaxies in the target data were assigned the bias based on the ν_{SN} -R bin they fall in, and the final bias for each tomographic bin was the lensfit-weighted average of these individual assignments. This procedure ensures the estimated bias accounts for any v_{SN} and R differences between the simulations and the data while also minimising the impact of the calibration selection bias.

Table 3.2 and Figure 3.16 show the multiplicative bias estimates for the KiDS-DR4 re-run with the updated *lens* fit from our constant shear simulations. The quoted errors only contain the statistical uncertainties from the linear fitting. Compared to Table 2 of K19, we reduced the statistical uncertainties by about half because of the larger sky area

⁽a) We verified that the four sets of input shears are sufficient to recover the previous results.

⁽b) Although more rotations suppress shape noise more efficiently (FC17), the selection effects diminish the actual performance of the shape noise cancellation (K19).

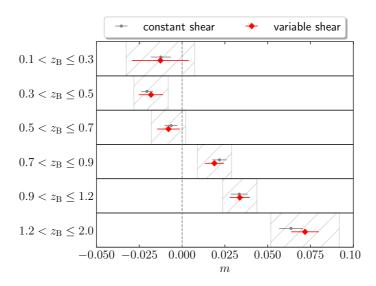


Figure 3.16: Multiplicative bias across tomographic bins for KiDS-DR4, with the updated *lens*fit. The red diamonds represent our final results, including corrections for the shear-interplay effect (Sect. 3.5.2) and PSF modelling bias (Sect. 3.5.3). The grey points, on the other hand, display the raw results from the idealised constant shear simulations (Sect. 3.5.1). The hatched regions denote the proposed nominal error budgets, intended for comparison (refer to Sect. 3.6 for more details).

simulated. Direct comparisons between the calibration values quoted in Table 3.2, cannot be made to those in K19 and Giblin et al. (2021). We updated the shape measurement algorithm *lens*fit and calibrated the raw measurement against PSF contamination in our analysis (see Sect. 3.4.2). These changes modify the effective size and signal-to-noise ratio distribution of the samples and hence the overall calibration in each tomographic bin. Furthermore, Giblin et al. (2021) accounts for the Wright et al. (2020a) 'gold' selection for photometric redshifts, which reduces the effective number density by $\sim 20\%$, compared to the sample simulated in this analysis.

3.5.2 Impact of blends at different redshifts

MacCrann et al. (2022) recently highlighted a complication that arises from blended objects at different redshifts, which are, therefore, sheared by different amounts. It stems from the fact that when objects are blended, a shear measurement of one object responds to the shear of the neighbouring object. This higher-order effect, which we refer to as 'shear interplay' through this paper, cannot be captured by the aforementioned constant shear simulations. So, we built an extra suite of variable shear simulations to account for this effect.

Since the shear interplay only happens when objects are blended, we built a blendingonly input catalogue for these additional simulations to save some computing time. This blending-only catalogue only contains bright galaxies with bright neighbours, assuming

Table 3.2: Shear bias for the six tomographic bins

ZB range	Ratio of $N_{\rm eff}$ (blending / whole)	$\Delta ar{m}$ blending	$\Delta m_{ m PSF}$	$m_{ m raw}$	<i>m</i> final
$0.1 < z_{\rm B} \le 0.3$	0.345	-0.012 ± 0.034	$+0.002 \pm 0.001$	-0.012 ± 0.006	-0.013 ± 0.017
$0.3 < z_{\rm B} \le 0.5$	0.332	-0.003 ± 0.014	$+0.004 \pm 0.001$	-0.021 ± 0.004	-0.018 ± 0.007
$0.5 < z_{\rm B} \le 0.7$	0.365	-0.021 ± 0.012	$+0.004 \pm 0.001$	-0.006 ± 0.004	-0.008 ± 0.007
$0.7 < z_{\rm B} \le 0.9$	0.366	-0.018 ± 0.008	$+0.003 \pm 0.001$	$+0.022 \pm 0.004$	$+0.019 \pm 0.006$
$0.9 < z_{\rm B} \le 1.2$	0.370	-0.013 ± 0.007	$+0.005 \pm 0.001$	$+0.033 \pm 0.005$	$+0.034 \pm 0.006$
$1.2 < z_{\rm R} \le 2.0$	0.358	$+0.000 \pm 0.008$	$+0.007 \pm 0.002$	$+0.064 \pm 0.007$	$+0.072 \pm 0.008$

general performance of the updated lensht simulations. All results are based on the KiDS-DR4 re-run with the updated lensfit before any redshift calibration. They only indicate the quoted along with individual m values are reported by the linear regression fitting, thus only reflecting the statistical power of SKiLLS correlation with the signal-to-ratio and resolution (see Sect. 3.5.2 for details). The Δm_{PSF} is the residual bias introduced by the PSF The ratio of $N_{\rm eff}$ between the blending-only simulation and the whole simulation is calculated from the measured catalogue with the the $m_{\rm final}$ are our final estimates with the corrections for the shear-interplay effect and PSF modelling bias (Sect. 3.5.4). The uncertainties modelling errors (see Sect. 3.5.3 for details). The m_{raw} results are derived from the idealised constant shear simulations (Sect. 3.5.1), and blending-only simulations (see Sect. 3.5.2 for details). The correction to the whole sample should also account for the $N_{\rm eff}$ ratio and the lensfit weight taken into account. The $\Delta m_{
m blending}$ is the mean residual bias introduced by the shear-interplay effect, estimated from the

that the blending effects caused by the faint objects are sufficiently accounted for by our main constant shear simulations, which include galaxies down to magnitude 27. It means we only ignore the higher-order shear-interplay effect from the faint objects, which is valid as long as the excluded faint galaxies are below the measurement limit of the survey. In practice, we selected all galaxies with an input r-band magnitude < 25. The choice of this magnitude cut meets the overall sensitivity of the KiDS survey. We further discarded those isolated galaxies whose nearest neighbour is 4'' away based on their input positions (see Fig. 3.20). The final selected sample covers $\sim 10\%$ of the entire input catalogue. But after the *lens*fit measurements, this blending-only simulation covers $\sim 35\%$ of the objects measured in the whole simulation (see Table 3.2 for the exact values). The higher fraction in the measured catalogue is because most objects fainter than 25 in the r-band magnitude are not measurable for KiDS.

To properly account for the shear-interplay effect, we need realistic shear fields with proper correlations between the shear and the environment of galaxies. We refer to Appendix 3.E for technical details of our approach to creating such variable shear fields. In short, we considered two primary contributions to the weak lensing signal: the cosmic shear due to the large-scale structure and the tangential shear induced by the foreground objects (also known as the galaxy-galaxy lensing effect). The cosmic shear was learned from the MICE Grand Challenge (MICE-GC) simulation (Fosalba et al. 2015b), whilst the tangential shear was calculated analytically by assuming Navarro-Frenk-White (Navarro et al. 1995) density profiles for the underlying dark matter halos. Figure 3.17 shows the average shear signals as a function of redshift. We see a roughly linear relationship between the mean signals and redshift. On average, the cosmic shear contributes more than the tangential shear. However, we note that the importance of the tangential shear varies between systems depending on the host halo mass of the foreground galaxies.

To increase the constraining power, we used 32 variable shear fields generated from the same learning algorithm but with different choices for the direction of the shear. Specifically, we created four variable shear fields with directions of the cosmic shear that differ by 90° . Then, we made eight copies for each shear field by rotating the final shear by 45° each time. We also created an extra suite of blending-only constant shear simulations to serve as a reference. The final sky area of these additional simulations is $7776 \, \text{deg}^2 (= 108 \times 36 \times 2)$. Except for the input shear, these blending-only simulations use the same pipeline, observational conditions and random seeds as the full simulations detailed in Sect. 3.5.1 so that we can directly correct the constant shear results using the extra bias estimated from these additional simulations.

While estimating the shear bias for constant shear simulations is straightforward by directly conducting the linear least squares fitting to all measurements using Eq. (3.11), given that the input shear values do not depend on the underlying sample. The situation is more complicated for variable shear simulations. The crucial caveat is that the shear bias is now correlated with redshift $[m_{\text{blending}}^{\text{varShear}}(z_{\text{true}})]$ due to the shear-interplay effect. Owing to the realistic shear field we built, we can measure $m_{\text{blending}}^{\text{varShear}}(z_{\text{true}})$ directly from simulations by performing the least squares fitting to sub-samples of

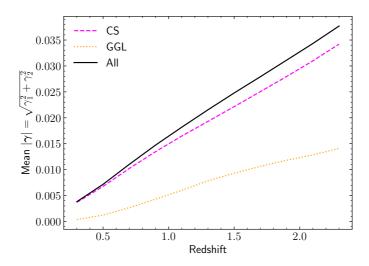


Figure 3.17: Variation of the shear field with respect to redshift. The solid black line represents the average amplitude of the final utilised shears, comprising two components: the cosmic shear (dashed magenta line) and the tangential shear (dotted orange line). For more information, see Appendix 3.E.

galaxies split based on their true redshift. The same approach can also be applied to the blending-only constant shear simulations to get $m_{\rm blending}^{\rm constShear}(z_{\rm true})$; only, in that case, we would expect a negligible correlation with the true redshift, except for some fluctuations stemming from the different signal-to-noise ratios between true redshift bins. Figure 3.18 shows the difference $\Delta m_{\rm blending}(z_{\rm true}) \equiv m_{\rm blending}^{\rm varShear}(z_{\rm true}) - m_{\rm blending}^{\rm constShear}$, which is a direct measure of the impact of the shear-interplay effect, as the only difference between the simulations is the input shear value. It demonstrates evident residuals that correlate with redshift, indicating the non-trivial impact of the shear-interplay effect. Interestingly, the high-redshift outliers, which have an estimated photo-z much lower than their true redshifts, show the most noticeable residuals across all tomographic bins, implying that the blends with objects from different redshifts are likely responsible for those outliers. This coupling between the photo-z and shear biases in blended systems warrants a dedicated future study.

To correct the raw shear bias derived in Sect. 3.5.1, an average correction $\Delta \bar{m}_{\text{blending}}$ is necessary, which takes into account z_{true} via the equation:

$$\Delta \bar{m}_{\text{blending}} = \int_0^\infty dz_{\text{true}} \ n(z_{\text{true}}) \ \Delta m_{\text{blending}}(z_{\text{true}}) \ . \tag{3.12}$$

Here, $n(z_{\text{true}})$ signifies the weighted number density related to redshift (as represented by the dashed lines in Fig. 3.18). The averaged results for individual tomographic bins are presented in Table 3.2 and Fig. 3.19.

In practice, the blending fraction, which is associated with the signal-to-noise ratio

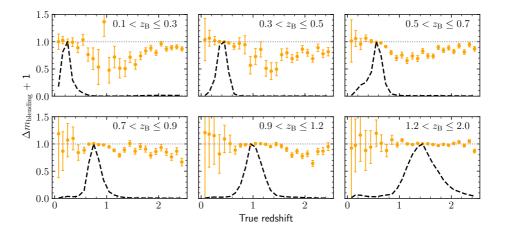


Figure 3.18: Shear bias residuals caused by the shear-interplay effect (orange points) as a function of true redshift, as estimated from the blending-only simulations. The residuals are derived from $\Delta m_{\rm blending} \equiv m_{\rm blending}^{\rm varShear} - m_{\rm blending}^{\rm constShear}$, where $m_{\rm blending}^{\rm varShear}$ refers to the shear bias from the blending-only variable shear simulations, and $m_{\rm blending}^{\rm constShear}$ refers to the shear bias from the blending-only constant shear simulations. The error bars correspond to the fitting uncertainties reported by the linear regression. These uncertainties are influenced by two factors: the number of objects used in the fitting and the amplitude of the input shear value. The dashed lines represent the normalised number density with respect to redshift.

and resolution—much like the bias itself—should also be accounted for. Hence, we implement the correction in each $\nu_{\rm SN}$ - $\mathcal R$ bin, adhering to the binning approach proposed for reweighting the simulation (refer to Sect. 3.5.1). Specifically, within each $\nu_{\rm SN}$ - $\mathcal R$ bin, we calculate the average correction $\Delta \bar{m}_{\rm blending}$ and the blending fraction. The blending fraction is measured as the ratio of the effective number counts between the blending-only simulation and the complete simulation. Subsequently, we adjust the raw bias in each $\nu_{\rm SN}$ - $\mathcal R$ bin with the product of $\Delta \bar{m}_{\rm blending}$ and blending fraction. The final corrected bias is the *lens*fit-weighted average of these adjusted biases. This correction methodology can be easily merged with the reweighting procedure, as they employ the same binning strategy.

Another more direct way to inspect the blending effect is to check the relation between the shear bias and the nearest neighbour distance in the input catalogue. Figure 3.20 demonstrates such estimations for both constant shear and variable shear simulations. We see a clear correlation between the bias and the neighbour distance in both simulations, indicating the significant impact of the blending effect. It also confirms our choice of 4" to define blended systems, as we barely see any correlation after this threshold. The other important finding is that the traditional constant shear simulations can already capture the dominant contributions from the blending effect. The higher-order impact

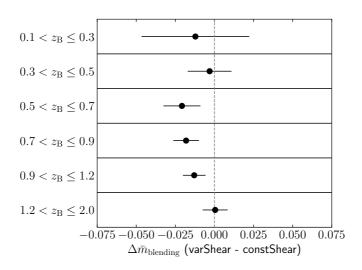


Figure 3.19: Mean residual multiplicative bias induced by the shear-interplay effect, as computed from Eq. (3.12). We emphasise that these results are derived from the blending-only simulations. When applying to the entire sample, we must also account for the blending fraction (the third column of Table 3.2).

we study in this section, shown as the bias difference between the variable shear and constant shear simulations, contributes relatively minor except for the very close blends. The aggressive treatment of the blending in *lens*fit can partially explain this finding, as it throws away most of the recognised blends (Hildebrandt et al. 2017).

We note that our variable shear simulations and the correction methodology differ from those of MacCrann et al. (2022). In their study, the simulated shear changes as a function of redshift, but, per redshift slice, it remains constant across the field of view. The chosen redshift intervals and adjusted shear have no physical meaning in their setups. But they built four sets of simulations by choosing different redshift intervals, so they were able to fit a smooth model to the simulated results, obtaining a continuous redshift-bias relation. In our approach, we computed the variable shear fields using a more physical model that accounts for the shear correlations to both the redshift and clustering of galaxies (see Appendix 3.E). Thanks to these realistic shear fields, we can measure the redshift-bias relation directly from the simulations without additional model fitting procedures. Our direct measurements confirmed the non-trivial impact of the shear-interplay effect (see Fig. 3.18). By design, our method results in large uncertainties for low redshift bins due to the small input shear values. Fortunately, these low redshift bins carry little cosmic shear signals, making the overall downgrade tolerable. Albeit following a different approach, our final result is consistent with MacCrann et al. (2022) finding that the overall correction due to the shear-interplay effect is negligible for the current weak lensing surveys. However, it will potentially impact the next-generation surveys.

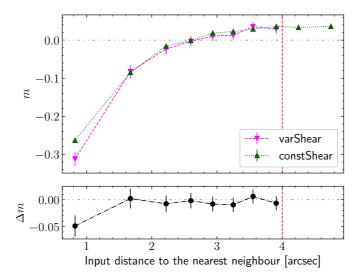


Figure 3.20: Multiplicative bias as a function of the distance to the nearest neighbour. The distance is measured in the input catalogue after exclusion of faint galaxies with r-band input magnitude >25. The values on the x-axis correspond to the weighted average of each sub-sample selected for the estimation of the multiplicative bias shown on the y-axis. The top panel displays the individual biases measured from the blending-only variable shear simulations (magenta points) and the blending-only constant shear simulations (dark green points). The vertical dashed lines denote the threshold we set when creating the blending-only simulations. Two additional dark green points beyond the threshold are calculated from the complete constant shear simulations. The bottom panel shows the difference between these two estimates (varShear - constShear).

3.5.3 PSF modelling bias

So far, we have ignored the PSF modelling errors, given the expected accuracy of PSF models relative to the requirement of the current weak lensing surveys (see e.g. Giblin et al. 2021). We used the input PSF for shape measurements (i.e. assuming perfect PSF modelling). However, as the requirement of systematics becomes more stringent, it becomes necessary to check the impact of PSF modelling errors. This section quantifies this impact by including the PSF modelling procedure in the simulations.

The SKiLLS images have realistic stellar populations and variable PSFs across the field, so we can apply the PSF modelling code directly to the simulated images using similar setups as for the data. We refer to Kuijken et al. (2015) for detailed descriptions of the PSF modelling algorithm used by KiDS. In short, it describes the position-dependent PSFs at the detector resolution using a set of amplitudes on a 48×48 pixel grid. The spatial variation of each pixel value is fitted with a two-dimensional polynomial of order n, with additional flexibility for allowing the lowest order coefficients to differ from CCD to CCD. This extra freedom allows for a more complex PSF variation between CCDs and, in principle, allows for discontinuities in the PSF between adjacent CCDs. When fitting

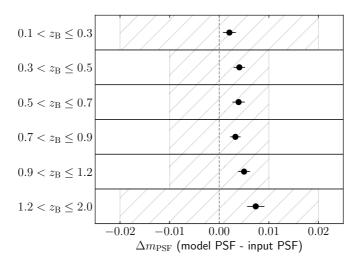


Figure 3.21: Adjustments in multiplicative bias when adopting modelled PSFs instead of the input PSFs. The hatched regions represent the nominal error budgets designated for comparison (refer to Sect. 3.6 for additional details).

to individual stars, the flux and centroid of each star are allowed to change, and a sinc function interpolation is used to align the PSF model with the star position. Following Giblin et al. (2021), we set n = 4 and allow the polynomial coefficients up to order 1 to vary between CCDs. We skipped the complicated star-galaxy separation procedure with an implicit assumption that the point-source sample used by KiDS is sufficiently pure as verified using NIR colours in Giblin et al. (2021). Instead, we built a perfect star sample by cross-matching the detected catalogue with the input star catalogue. However, we still applied the same magnitude and signal-to-noise ratio cuts as used in the data to ensure a similar noise level in the modelled stars.

We selected 30 tiles from the available 108 fiducial tiles to test the influence of PSF modelling uncertainty on the multiplicative bias. These selected tiles cover the whole range of the PSF size, including the minimum and maximum. We performed the PSF modelling on the selected tiles and re-ran *lens*fit using the modelled PSFs. Since all the images and detection catalogues are unchanged, the shift of the shear bias directly quantifies the contribution of the PSF modelling errors. Figure 3.21 and Table 3.2 show the shifts for the six tomographic bins. We find the PSF modelling procedure does introduce small yet noticeable biases. Our fiducial results take these additional biases into account.

3.5.4 Results

The final results after accounting for both the shear-interplay effect and PSF modelling errors are listed as m_{final} in Table 3.2 and shown as the red points in Fig. 3.16. Within the current statistical uncertainties, the average shifts due to the shear-interplay effect and

PSF modelling errors are insignificant across all redshift bins, as indicated in Fig. 3.16 between the grey points and the red points. A more noticeable change is the increased uncertainty introduced by the correction of the shear-interplay effect, especially in the low redshift bins where the input shear values are overall small in the variable shear simulations. Our proposed systematic error budgets account for these additional uncertainties (the hatched regions in Fig. 3.16).

3.6 Sensitivity analysis

Given the resemblance between the SKiLLS and KiDS images and the reweighting in the signal-to-noise ratio and \mathcal{R} when estimating the shear biases, it is reasonable to assume that the estimates from SKiLLS can be used to correct the actual measurements. Nevertheless, it is still worth testing the robustness of SKiLLS results and accounting for any potential systematic uncertainties. We start with tests proposed by FC17 and K19 in Sect. 3.6.1. Thanks to the broad coverage of observational conditions in SKiLLS, we can quickly achieve these analyses without dedicated test runs. Additionally, we test how sensitive the *lens*fit results are to the changes in the input galaxy morphology (Sect. 3.6.2). For comparison reasons, we propose some nominal error budgets based on the general performance of SKiLLS and the overall requirements of lensing analyses with KiDS. Specifically, we set an error budget of 0.02 for the first and sixth tomographic bins and 0.01 for the remaining bins. We found these nominal error budgets are conservative enough that our results are robust within them. Nevertheless, we note that these nominal error budgets can be over-conservative for cosmic shear analyses. In which case, we can estimate more accurate systematic uncertainties following other more aggressive approaches proposed by previous KiDS analyses (Giblin et al. 2021; Asgari et al. 2021).

3.6.1 Impact of observational conditions

When developing SKiLLS, we improved most of the critical sources of uncertainty in the previous KiDS simulations. For instance, we based our input galaxy catalogue on *N*-body simulations, so it has reasonable clustering features and is complete down to 27 in the *r*-band magnitude. We learned realistic morphologies from observations using a powerful technique, dubbed vine copulas, which captures the multi-dimensional correlations between ellipticities and other galaxy properties. We included six stellar catalogues to account for the varying stellar densities across the survey sky. We covered more variations of the PSF models and background noise levels. Above all, we measured photo-zs directly from the simulated multi-band images to properly account for the correlation between the measurement uncertainties on the redshift and shear estimates. Consequently, most of the sensitivity analyses proposed by FC17 and K19 are either trivial or redundant for the SKiLLS results.

Still, we examine the robustness of the *lens*fit results against some crucial properties by comparing between sub-samples. The basic idea is to split the fiducial simulations into three sub-samples based on a targeted property and examine the consistency of

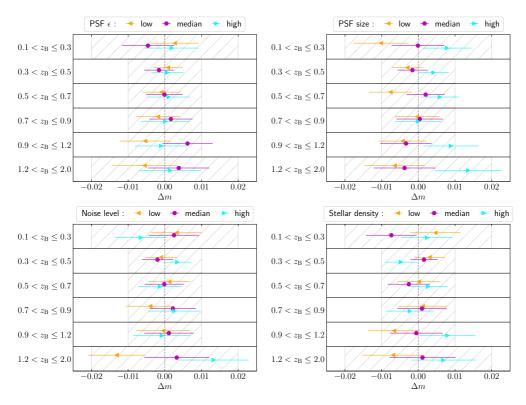


Figure 3.22: Variations in multiplicative bias when the fiducial simulations are divided into three subsets, each reflecting different observational conditions. From left to right and top to bottom, the panels present results when the division is based on PSF ellipticity, PSF size, background noise level in r-band images, and stellar density. The hatched regions denote the proposed nominal error budgets for comparison (refer to Sect. 3.6 for additional details). We note that these shifts represent the maximum possible systematic biases in our results (see Sect. 3.6.1 for further clarification).

their bias estimates to the fiducial results. These sub-samples contain roughly equal numbers of measured objects while covering different ranges of the targeted property. After applying the overall shear correction from the whole sample to the sub-samples, we calculate their residual biases to quantify the impact of the variations of the targeted property. We note that the estimated residuals are not systematic biases in our fiducial results, but they indicate the robustness of the shape measurement algorithm against the tested properties. Ideally, if the simulations fully match the data in the distributions of the targeted property, we would still expect an accurate bias estimate even if the estimated residuals are large. For that account, the estimated residuals are conservative upper limits of the systematic biases in our results.

Figure 3.22 shows the estimated residuals for the variations in four critical properties of the simulated images: the PSF ellipticity, PSF size, background noise level in r-band images, and stellar density. It indicates that our fiducial results are robust within the

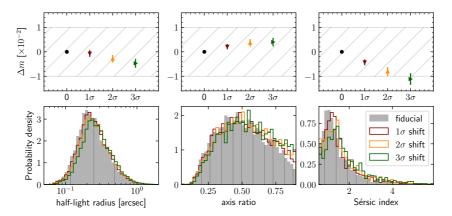


Figure 3.23: Variations in multiplicative bias due to increased morphological parameters of all input galaxies. Top panels show shifts in bias for increased factors of $1 + 1\sigma$ (dark red), $1 + 2\sigma$ (dark orange), and $1 + 3\sigma$ (dark green), where σ indicates median relative uncertainties from Griffith et al. (2012). Hatched regions represent the nominal 0.01 error budget. Bottom panels show normalised histograms of parameter distributions before and after the alterations. The shift of entire distributions depicts extreme cases, as mean shifts across all galaxies would typically be smaller.

nominal error budgets, considering the shifts shown in the plots are the upper limits of possible deviations.

3.6.2 Impact of the input galaxy morphology

We learned the galaxy morphology from Griffith et al. (2012) based on Sérsic models fitted to the HST observations. We have shown that our copula-based learning algorithm captures the properties of the reference sample (see Sect. 3.2.1). However, the reference sample itself contains measurement errors. This section examines how sensitive the *lens*fit measurements are to the changes in the input galaxy morphology.

We focus on the three morphological parameters used to describe the Sérsic profile: the half-light radius, axis ratio and Sérsic index. To get some indication of the overall accuracy of the reference sample, we first checked the fitting uncertainties. We found that the median relative uncertainties for these three parameters are $\lesssim 5\%$, $\lesssim 5\%$ and $\lesssim 10\%$, respectively. We took these values (quoted as σ below) as the benchmark for changing the input galaxy morphology. We built new input catalogues by increasing a certain parameter with 1σ , 2σ and 3σ each time while keeping the other parameters unchanged. We generated test simulations using these new input catalogues and measured the bias difference with respect to the fiducial simulations.

Figure 3.23 presents the test results from 10 tiles of simulations. We find minor residuals in most cases, with the most significant shifts seen when changing the Sérsic index. We note that we shifted all galaxies with the same amount of fractions, resulting

in an overall shift of the whole distribution, as shown in the bottom panels of Fig. 3.23. Given that the entire distribution's uncertainty is much smaller than the individual measurement uncertainties, we are testing the most extreme cases. Hence, the measured residuals only indicate the sensitivity of *lens*fit towards the input galaxy properties but cannot be seen as systematics in our fiducial results. To achieve tighter requirements for future surveys, we will need a shape measurement method that is less susceptible to the galaxy properties, as the fidelity of image simulations will always be limited by the realism of the input galaxy catalogue. For the upcoming KiDS-Legacy analysis, we will, therefore, also explore an alternative method based on the Metacalibration technique (Huff & Mandelbaum 2017; Sheldon & Huff 2017), which is expected to be more robust against the galaxy properties (Yoon et al., in prep.).

3.7 Discussion and conclusions

Achieving an unbiased measurement of the ensemble shear signal is crucial for maintaining the integrity of precision cosmology studies using weak lensing surveys. Contemporary shape measurement techniques attained an accuracy level of one percent, or even a fraction thereof. However, as the statistical powers of weak lensing surveys continue to expand, the demand for systematics control intensifies. This resulted in an increased focus on higher-order effects that differ from the shape estimation bias, such as selection bias, PSF modelling errors, and shear-interplay bias. These effects present a challenge to eradicate merely by refining shape measurement algorithms. Alternatively, image simulations offer promising capabilities for calibrating these higher-order effects.

In this chapter, we introduced the third-generation image simulations for the KiDS survey, termed SKiLLS, following SCHOol (FC17) and COllege (K19). These simulations incorporate several substantial enhancements to meet the calibration requirements of the KiDS-Legacy analysis, which utilises an updated *lens*fit. Among the key improvements are the simulation of full nine-band images and the creation of a joint shear-redshift mock catalogue. Balancing the sample volume and realism of galaxy morphology, we combined cosmological simulations with deep imaging observations as input. Additionally, we augmented the image realism by incorporating variations in PSF between CCDs, stellar density, and noise levels between pointings. We closely mirrored the entire KiDS procedure for photometric measurements, encompassing *r*-band detection, PSF Gaussianisation, forced multi-band photometry, and photo-*z* estimation. The large volume of simulated galaxies and their realistic photometric properties will enhance not only shear calibration but also redshift calibration, providing a valuable resource for further study (van den Busch et al., in prep.).

We extended our investigations to consider the impact of galaxy blends at varying redshifts by generating realistic shear fields that take into account both redshift and galaxy clustering. We also accounted for PSF modelling errors by implementing the PSF modelling procedures on the image simulations. Additionally, we conducted sensitivity tests, including changing the input galaxy properties, demonstrating the robustness of SKiLLS-calibrated measurements for future KiDS lensing studies. The final shear

3

calibration results for the updated *lens*fit are outlined in Table 3.2 and Fig. 3.16. Based on our statistical uncertainties and sensitivity tests, we believe that the shear bias estimated from SKiLLS is within the nominal error budget of 0.02 for the first and sixth tomographic bins and 0.01 for the others. Furthermore, our studies provide valuable insights for calibrating future weak lensing surveys.

The success of image simulations heavily depends on the realism of the input galaxy population in terms of photometry, morphology, and clustering. Recent image simulations used high-quality imaging observations as input, but these have limitations in sample volume and depth, which may soon prove inadequate for next-generation weak lensing surveys. While there is an alternative approach that uses input galaxy populations from cosmological simulations, these simulations are currently unable to fully replicate observed galaxy morphology – a crucial feature in image simulations.

In our study, we explored the feasibility of an integrated approach that combines the advantages of cosmological simulations with high-quality imaging observations. We introduced a copula-based learning algorithm designed to mimic and establish a link between observed morphology and synthetic galaxies from cosmological simulations. The results indicate that this hybrid methodology shows promise for future image simulations requiring a substantial volume of galaxies.

Recent studies have already indicated the necessity for shear calibration to consider redshift-related selections. This requires simulating multi-band observations to account for the measurement of photometric redshifts (e.g. K19; MacCrann et al. 2022). We extended this to demonstrate that multi-band image simulations, with a sufficiently large volume of galaxies, not only improve shear calibration but also redshift calibration. By performing the full procedure for photometric measurements, we achieved realistic photometric properties in the mock catalogue. This comprehensive approach advances over previous catalogue-level simulations (e.g. Hoyle et al. 2018; van den Busch et al. 2020; DeRose et al. 2022). In addition, image simulations allow us to examine the impact of blending on redshift estimates, which is challenging to account for at the catalogue level. Given the importance of blending, we believe that integrating shear and redshift calibrations with multi-band image simulations will be crucial for future high-accuracy tomographic analyses.

MacCrann et al. (2022) recently investigated the effects of blended systems in which galaxies experience varying shears, a phenomenon we refer to as 'shear interplay' throughout this paper. We extended their work by creating realistic variable shear fields that account for both the redshift and clustering of galaxies, explicitly including galaxy-galaxy lensing contributions. Although our final results confirmed its relatively minor impact on current weak lensing surveys (see Fig. 3.16), we detected a significant correlation between redshift and shear bias from our blending-only variable shear simulations. This correlation underscores the presence of the shear-interplay effect and its contributions (see Fig. 3.18). Furthermore, we observed that photo-z outliers demonstrate the most pronounced shear interplay, implying a mutual origin of the shear and redshift biases. A focused study is needed to further investigate this correlation in blended systems, as it will become increasingly important for the next generation of

weak lensing surveys.

Typically, image simulations bypass the PSF modelling process, due to the PSF validation conducted in data (see e.g. Giblin et al. 2021). However, thanks to the realistic SKiLLS images, we are able to evaluate the impact of PSF modelling errors by applying the PSF modelling code directly on simulated images. By comparing the shear biases derived from runs with and without PSF modelling, we discerned residual biases from PSF modelling errors, albeit below one percent. Despite being negligible for current requirements, this will be a point of concern for upcoming weak lensing surveys. Therefore, we underscore the necessity of refining the PSF modelling algorithm or incorporating it into image simulations for future surveys.

Lastly, we investigated the sensitivity of our simulation to the characteristics of the input galaxy population. By altering the input values of morphological parameters, we determined that our current standard shape measurement method, *lens*fit, is somewhat sensitive to the input galaxy shapes, but this sensitivity falls within an acceptable range for KiDS analysis. Despite this, we aim to implement an alternative approach based on the Metacalibration technique (Huff & Mandelbaum 2017; Sheldon & Huff 2017) for KiDS-Legacy analysis, which has proven to be more robust against variations in galaxy properties (Yoon et al., in prep.). For future weak lensing surveys, it will be crucial to develop methods less susceptible to galaxy properties, as image simulations may never fully capture the observed galaxy population due to limitations in the input catalogue.

Acknowledgements

We thank Fedor Getman for providing the deep VST-COSMOS catalogue and Arun Kannawadi for reading the manuscript and providing useful comments. We also wish to thank other members of the KiDS-Legacy Calibration Working Group (especially Benjamin Joachimi, Benjamin Stölzner and Anna Wittje) for informative discussions and suggestions through numerous teleconferences. This work used the compute resources from the Academic Leiden Interdisciplinary Cluster Environment (ALICE) provided by Leiden University. We acknowledge support from: the Netherlands Research School for Astronomy (SSL); the Royal Society and Imperial College (KK); the Netherlands Organisation for Scientific Research (NWO) under Vici grant 639.043.512 (HHo); and the UK Science and Technology Facilities Council (STFC) under grant ST/N000919/1 (LM) and ST/V000594/1 (CH). We also acknowledge support from the European Research Council (ERC) under grant agreement No. 647112 (CH) and No. 770935 (HHi, JLvdB, AHW); 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, MY); the Deutsche Forschungsgemeinschaft Heisenberg grant Hi 1495/5-1 (HHi); 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 (MBi); the Polish Ministry of Science and Higher Education through grant DIR/WK/2018/12 (MBi); the University of Western Australia through a Scholarship for International Research Fees and Ad Hoc Postgraduate

Scholarship (MBr); and the ARC Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D) through project number CE170100013 (CL). The results in this paper are based on observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs: 088.D-4013, 092.A-0176, 092.D-0370, 094.D-0417, 177.A-3016, 177.A-3017, 177.A-3018 and 179.A-2004, and on data products produced by the KiDS consortium. The KiDS production team acknowledges support from: Deutsche Forschungsgemeinschaft, ERC, NOVA and NWO-M grants; Target; the University of Padova, and the University Federico II (Naples). Contributions to the data processing for VIKING were made by the VISTA Data Flow System at CASU, Cambridge and WFAU, Edinburgh. The SURFS-SHARK simulations were produced at the Pawsey Supercomputing Centre with funding from the Australian Government and the Government of Western Australia.

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, KK, HHo, LM) 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.

3.A An empirical modification to the synthetic photometry

We detail the proposed empirical modification of the Shark photometry in this appendix. It intends to improve the agreement of the magnitude counts between the simulations and observations, which is critical for the redshift and shear calibrations.

We took the COSMOS2015 catalogue as the benchmark under an implicit assumption that the COSMOS field is representative. The COSMOS2015 catalogue is a near-infrared-selected photometric catalogue containing 30-band photometry, precise photometric redshifts and stellar masses for more than half a million objects (Laigle et al. 2016). We note that measurement uncertainties and modelling errors are inevitable for observations, especially for faint objects. Therefore, the COSMOS2015 catalogue cannot, in principle, be treated as the truth. Nevertheless, these uncertainties are tolerable for calibrating a KiDS-like sample. Following this reasoning, we tuned the simulated properties solely based on the COSMOS2015 measurements for the sake of simplicity, but caution any physical interpretation of our modified results.

First of all, we must locate the cause of the discrepancy. As the Shark free parameters were tuned using the observed stellar mass functions, we would expect the number density of the Shark galaxies is realistic. This is confirmed by Figure 3.24, where we see a good agreement of the stellar mass distributions between the data and simulations. As a next step, we inspected the stellar mass-to-light ratio (Υ_{\star}) , for which took the K_s -band photometry as an indicator of the total luminosity as it is least affected by the dust extinctions. Figure 3.25 shows the comparing results as a function of the stellar mass in several redshift bins. Noticeably, the Shark Υ_{\star} is systematically higher than the COSMOS2015 one, especially in the low stellar mass and low redshift ranges. It can, at least partially, explain the discrepancy seen in the magnitude distributions. Fortunately, this Υ_{\star} difference is easy to calibrate without changing other intrinsic properties, such as the colours, redshifts, and positions.

We, therefore, conducted an empirical modification of the simulated magnitudes to account for the Υ_{\star} difference. We divided Shark and COSMOS2015 galaxies into 24 × 23 evenly spaced small bins based on their redshifts and stellar masses. In each bin, we calculated the median Υ_{\star} for the Shark and COSMOS2015 galaxies, separately. To mitigate the observational uncertainties, we only used the COSMOS2015 galaxies with good stellar mass estimations ($\delta M_{\star} < 0.15 M_{\star}$). For bins that lack observations, we extrapolated $\Upsilon_{\star, \text{ obs}}$ as a function of M_{\star} for each redshift slice. After inspecting the general trend, we found a good fit by combining an exponential descending function in the low M_{\star} end and a linear ascending function in the high M_{\star} end. From these estimates, we constructed a magnitude modification factor Δ mag as

$$\Delta \text{mag} = -2.5 \log_{10} \left(\frac{\text{median}[\Upsilon_{\star, \text{Shark}}]}{\text{median}[\Upsilon_{\star, \text{obs}}]} \right). \tag{3.13}$$

Figure 3.26 demonstrates the estimated Δ mag values in the 2D redshift-stellar mass plane. Following the difference seen in Fig. 3.25, substantial modifications happen in the low mass and low redshift bins. Therefore, the magnitude modification reduces the

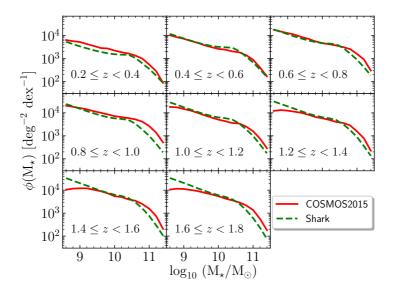


Figure 3.24: Comparison of stellar mass functions. The COSMOS2015 catalogue (represented by red solid lines) utilises the median values derived from the marginalised likelihood distributions. On the other hand, the Shark catalogue (represented by green dashed lines) operates under the assumption that the total stellar mass is equivalent to the combined stellar masses of the bulge and the disc.

range of magnitudes of Shark galaxies. We note that the different bands share the same Δ mag values, so the colours of individual galaxies are preserved.

3.B Modelling multivariate distributions with vine copulas

We outline some necessary background on the vine-copula modelling in this appendix. For a comprehensive introduction, we refer to Joe (2014) and Czado (2019).

A copula is simply a multivariate cumulative distribution function (CDF) with uniformly distributed margins. The Sklar (1959) theorem states that any d-dimensional CDF F(x), with univariate margins $F_1(x_1), ..., F_d(x_d)$, can be described as $F(x) = C_{1,...,d}(F_1(x_1), ..., F_d(x_d))$, where $C_{1,...,d}$ is the corresponding copula function. Therefore, given a joint probability distribution function (PDF) f(x) with d-dimensional variables $x = (x_1, ..., x_d)$, we can always find a copula density $c_{1,...,d}$ that is the partial differentiation of the copula $C_{1,...,d}$, such that

$$f(\mathbf{x}) = c_{1,\dots,d}(F_1(x_1), \dots, F_d(x_d)) \cdot f_1(x_1) \cdot \cdot \cdot f_d(x_d) . \tag{3.14}$$

It means we can divide the modelling of any joint multi-dimensional PDF into two parts: one for the independent distributions of the individual random variables

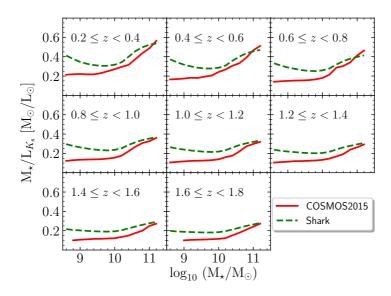


Figure 3.25: K_s -band stellar mass-to-light ratio as a function of stellar mass. The red and green lines correspond to the galaxies in the COSMOS2015 and Shark catalogues, respectively.

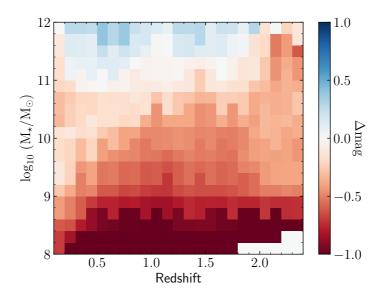


Figure 3.26: Distribution of the magnitude modification factor Δ mag in the redshift-stellar mass plane. Negative values are denoted in red, while positive values are denoted in blue. The definition of Δ mag is given in Eq. (3.13). For each galaxy, the same Δ mag value is added to the apparent magnitudes across all available bands.

 $\{f_i(x_i)\}\$, and the other for their mutual dependence captured by the copula density $c_{1,\dots,d}(F_1(x_1),\dots,F_d(x_d))$.

The restriction of the classical copula method is that most of the flexible copula families available in the literature are bivariate, making it tricky to deal with high-dimensional distributions. In this aspect, the vine copula method stands out as an effective approach (Bedford & Cooke 2002; Aas et al. 2009). A vine copula is a graphical model organising a set of bivariate copulas, called pair-copulas. The chain rule states that any PDF f(x) can be decomposed as

$$f(\mathbf{x}) = f(x_d) \cdot f(x_{d-1}|x_d) \cdot f(x_{d-2}|x_{d-1}, x_d) \cdot \cdot \cdot f(x_1|x_2, ..., x_d) , \qquad (3.15)$$

with f(.|.) being the conditional PDF. Aas et al. (2009) further states that each term in Eq. (3.15) can be decomposed into an appropriate pair-copula times a conditional marginal density as described by the following general formula

$$f(x|v) = c_{xv_j|v_{-j}}(F(x|v_{-j}), F(v_j|v_{-j})) \cdot f(x|v_{-j}), \qquad (3.16)$$

where v stands for a d-dimensional vector, v_j is an arbitrary component of v, and v_{-j} denotes the v-vector excluding this component. Therefore, the multiple dependence can be captured by a product of pair-copulas acting on underlying conditional probability distributions. Since the decomposition shown in Eq. (3.15) is not unique, there is a significant number of possible pair-copula constructions. These possibilities are organised by the graphical models, that is the vines.

3.C Transformation of the SDSS filters to the KiDS/VIKING filters

This appendix details the transformation of the Sloan Digital Sky Survey (SDSS) photometric system to the KiDS/VIKING system. The SDSS photometric system comprises five colour bands (u, g, r, i, z) that cover wavelengths ranging from ultra-violet at 3000 to near-infrared at 11 000 (Fukugita et al. 1996), whilst the KiDS/VIKING system contains optical filters (u, g, r, i) mounted on the VST OmegaCAM camera (Kuijken 2011) and near-infrared filters (Z, Y, J, H, K_s) mounted on the VISTA infrared camera (González-Fernández et al. 2018). Figure 3.27 compares the filter curves from these two systems. The differences are noticeable, especially for the Z filter, where the KiDS/VIKING system cuts the tail towards long wavelengths. We used the following relation to correct these differences:

$$X_{\text{KiDS/VIKING}} = X_{\text{SDSS}} + j(z_{\text{true}}) (X_{\text{SDSS}} - W_{\text{SDSS}}) + h(z_{\text{true}}),$$
 (3.17)

where X corresponds to the target filter, whilst W is another filter, helping to define the colour. Given the superior depth of the r-band measurement, we picked it as the Y filter whenever possible. When the r band is the target filter, we chose the g band as the Y filter. The coefficients $j(z_{\text{true}})$ and $h(z_{\text{true}})$ are correlated with the redshift, for which we

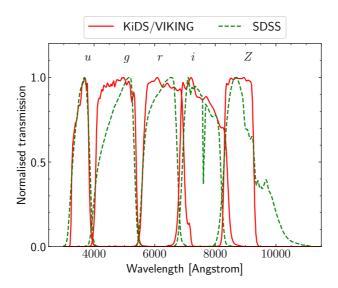


Figure 3.27: Comparison of normalised transmission curves. The curves represent the *ugriZ* filters in the SDSS photometric system (green dashed lines) and the KiDS/VIKING system (red solid lines).

took values from the ProSpect web-portal¹⁸. For the redshift, we used the true redshift from the input SURFS-Shark simulations.

As for the SKiLLS measured photometry, we need to correct six measurements: the five u-, g-, r-, i-, Z-band magnitudes measured in the Astro-WISE images (MAG_GAAP_X) and the r-band magnitudes measured in the thellimages (MAG_AUTO). There is no need to correct the remaining $YJHK_s$ bands as SKiLLS also uses VISTA filters for them. Figure 3.28 shows the distributions of the magnitude modification as a function of the initially measured magnitude. The modifications are generally small, especially for the u and g bands. Even for the r and g bands with the most significant differences, the majority of objects has a modification ≤ 0.05 . Accordingly, the changes in the overall magnitude and colour distributions are negligible. Still, we get a better agreement with the data in the photo-g distributions after transforming to the KiDS/VIKING filters, as shown in Fig. 3.29.

3.D Selection criteria for the updated lensfit catalogue

This appendix details all selections we propose to the updated *lens*fit shear catalogue. Most of the selection criteria were taken from earlier KiDS analyses, documented in Hildebrandt et al. (2017). These include:

- Several lensfit fitclass cuts to discard:
 - (a) objects without sufficient data, for example, those fall near the image edge

¹⁸https://transformcalc.icrar.org/

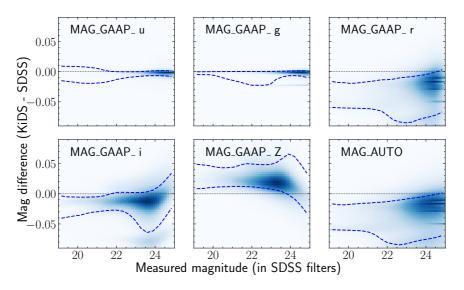


Figure 3.28: Joint distributions of the initially measured magnitude and the magnitude modifications. The dashed lines indicate the 16 and 84 percentiles. 'MAG_GAAP_X' magnitudes are those measured by the GAAP in the Astro-WISE images, while 'MAG_AUTO' represents magnitudes measured by SEXTRACTOR in the *r*-band THELI images (refer to Sect. 3.3 for more details).

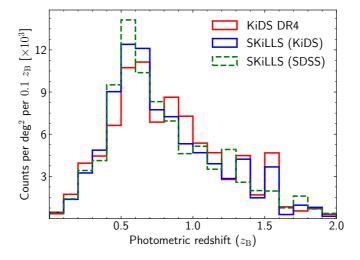


Figure 3.29: Distributions of the photometric redshift estimates. The KiDS-DR4 results are represented by the red histogram. The initial measurements in the SDSS filters are shown by the green histogram, while the blue histogram presents results corrected to the KiDS/VIKING filters. Most of the improvement is visible around $z_B \sim 0.55$ and 1.55.

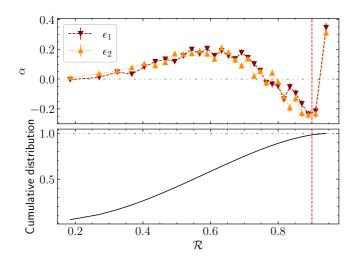


Figure 3.30: PSF leakage and effective number density based on the resolution factor. Top panel: PSF leakage; Bottom panel: effective cumulative distribution. Measurements for weighted average ellipticity ϵ_1 (dark-red triangle) and ϵ_2 (dark-orange triangle) were taken before PSF contamination correction. The vertical red dashed line indicates the proposed resolution cut (\Re < 0.9), resulting in a 2% loss in effective number density.

or a defect (fitclass = -1),

- (b) objects classified as duplicates (fitclass = -10),
- (c) objects poorly fitted by the given bulge plus disc galaxy model (fitclass = -4),
- (d) objects identified as stars and star-like point sources (fitclass = 1 and 2),
- (e) objects whose fitted centroid is more than 4 pixels away from the input centroid (fitclass = -7),
- (f) objects that are unmeasurable, usually because of being too faint (fitclass = -3).
- 2. A magnitude cut to remove bright objects (MAG_AUTO > 20).
- 3. A contamination radius cut to mitigate blending effects (contamination_radius > 4.25 pixels)
- 4. Removing asteroids based on the object colours (MAG_GAAP_g MAG_GAAP_r \leq 1.5 or MAG_GAAP_i MAG_GAAP_r \leq 1.5).
- 5. Removing unresolved binary stars by requiring objects with ellipticity > 0.8 to have a measured scalelength

$$\geq 0.5 \times 10^{(24.2 - \text{MAG_GAAP_r})/3.5} \; \text{pixels} \; .$$

- 6. A non-zero weight cut using the weight bias corrected weight (Sect. 3.4.2).
- 7. A resolution cut to remove poorly resolved objects ($\mathcal{R} < 0.9$).

The resolution cut is a new criterion proposed in this work. When developing our empirical correction method for the PSF contamination (Sect. 3.4.2), we noticed that objects with poor resolution contain very high PSF leakages, as demonstrated in Fig. 3.30. These poor-resolution outliers contribute little to the effective number density but introduce significant bias. So we propose a new selection using the resolution factor defined in Eq. (3.7). We found the proposed cut of $\mathcal{R} < 0.9$ can remove most outliers while only decreasing the effective number density by \sim 2 per cent.

3.E Building the variable shear field

In this appendix, we detail the creation of a realistic shear field accounting for the shear dependence on the redshift and clustering of galaxies. We considered the two main contributions to the weak lensing signals: the cosmic shear from the large-scale structure, and the tangential shear from the foreground objects (also known as the galaxy-galaxy lensing effect).

We split the blending-only sample into two classes based on their relative line-of-sight distances to their brightest neighbours. Those more distant than their brightest neighbours are referred as the background galaxies, whilst the remaining are the foreground galaxies. This classification is necessary to quantify the shear correlations within the blended systems. We found a roughly equal number of foreground and background galaxies in our sample.

For the cosmic shear effect, we learned it from the galaxy lensing mocks associated with the MICE Grand Challenge (MICE-GC) simulation (Fosalba et al. 2015b). The MICE-GC simulation is a large volume N-body light-cone simulation developed by the Marenostrum Institut de Ciències de l'Espai (MICE) collaboration (Fosalba et al. 2015a). It contains $\sim 6.9 \times 10^{10}$ dark matter particles with a mass of $\sim 2.9 \times 10^{10}$ $h^{-1} \rm M_{\odot}$ and a softening length of 50 $h^{-1} \rm kpc$, in a box of 3072 $h^{-1} \rm Mpc$ aside. The simulation starts at $z_i = 100$ and produces the light-cone in 265 steps from z = 1.4 to 0. It builds halo catalogues using the Friends-of-Friends algorithm (Crocce et al. 2015), and subsequently populates galaxies using halo occupation distribution recipes along with the subhalo abundance matching technique (Carretero et al. 2015). The construction of all-sky lensing maps follows the Onion Universe approach, which reaches a sub-arcminute spatial resolution up to z = 1.4 (Fosalba et al. 2015b). Here we used the second version of the catalogue, named MICECAT2, from the CosmoHub web-portal (Carretero et al. 2017; Tallada et al. 2020)¹⁹.

Following the building of the blending-only sample for SKiLLS, we selected blended objects and classified foreground and background galaxies for MICECAT2 under the same conditions expect for the magnitude cut. We first estimated the relationship between

¹⁹https://cosmohub.pic.es/

the mean cosmic shear amplitude and redshifts by averaging individual shear values of galaxies in redshift bins defined with a width of 0.1. These redshift-dependent mean amplitudes are good approximations for cosmic shears experienced by the foreground galaxies. It is more intricate to get proper cosmic shears for the background galaxies. Because of the overlapping line-of-sights of the blended objects, we expect the cosmic shear experienced by the background galaxy (γ_B) to correlate with that in its neighbour (γ_F). Based on our tests, the correlation can be described by a linear formula

$$\gamma_{\rm B}(z_{\rm B}, z_{\rm F}) = A(z_{\rm B}, z_{\rm F}) \cdot \gamma_{\rm F} + \gamma_{\rm I}(z_{\rm B}, z_{\rm F}) , \qquad (3.18)$$

with the scaling factor

$$A(z_{\rm B}, z_{\rm F}) \equiv \frac{D_{\rm c,B} - 0.5D_{\rm c,F}}{D_{\rm c,B}} \cdot \frac{D_{\rm c,F}}{D_{\rm c,F} - 0.5D_{\rm c,F}},$$
(3.19)

and an offset $\gamma_{\rm I}(z_{\rm B},z_{\rm F})=\mathcal{N}\left[0,\;\sigma_{\rm I}(z_{\rm B},z_{\rm F})\right]$ following the Gaussian distribution with a mean of zero and variance depending on redshifts of both galaxies. The $D_{\rm c,B}$ and $D_{\rm c,F}$ denote the comoving distances to the background galaxy and its neighbour, respectively. The scaling factor A reflects the geometrical relation between the blended objects; whilst the offset $\gamma_{\rm I}$ specifies contributions from the intermediate structures between blended galaxies. We estimated the redshift-dependent variance of $\gamma_{\rm I}$ again from MICECAT2 by measuring the dispersion of $\gamma_{\rm B} - A \cdot \gamma_{\rm F}$ in each redshift bin. Because the MICECAT2 stops at z=1.4, we linearly extrapolated measured values to z=2.5, which is the limit of SKiLLS. Figure 3.31 shows the learned cosmic shear as a function of redshift. The black solid line indicates the mean amplitude of the $\gamma_{\rm F}$ component; whilst the coloured lines present the dispersion of the $\gamma_{\rm I}$ component. It illustrates that the linear extrapolation captures the general trends towards the high redshift for both components.

We note that MICECAT2 assumes a Λ CDM cosmology with parameters from the Wilkinson Microwave Anisotropy Probe five-year data (WMAP5, Dunkley et al. 2009), whilst our base SURFS-Shark simulation uses cosmological parameters from Planck Collaboration (2016). Therefore, the cosmic shear field we learned from MICECAT2 does not necessarily match the galaxy mock we are using. But, since the current calibration still adopts one-point statistics (see Eq. 3.11), our calibration results are robust against detailed galaxy populations or underlying cosmologies and even more so to the higher-order correlation between galaxy populations and cosmology. We defer the proper treatment using a ray-tracing approach with consistent properties from the underlying cosmological simulations to future studies.

Besides the cosmic shear, a background galaxy also suffers from the tangential shear induced by the host dark matter halo of its neighbour. We calculated this effect analytically by assuming Navarro-Frenk-White (NFW) density profiles for dark matter halos presented in the SURFS-Shark simulation. The NFW profile, proposed by Navarro et al. (1995), is the most popular analytical model for dark matter halos, given its ability to describe the radial matter distribution of dark matter halos over a wide range

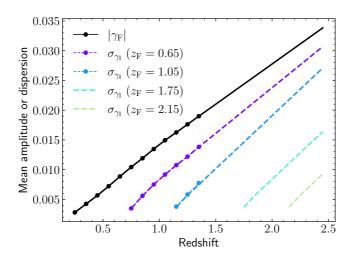


Figure 3.31: Cosmic shear signals learned from MICECAT2 (as described in Eq. 3.18). The black solid line and points represent the mean amplitude of the γ_F component, while the coloured lines and points denote the γ_I dispersion for various redshifts of the foreground galaxies. The points are direct measurements from MICECAT2, while the lines represent linear extrapolations.

of masses (Navarro et al. 1996, 1997). Its mass density is described by the formula

$$\rho(r) = \frac{\rho_{\rm cr} \,\delta_c}{(r/r_{\rm s})(1 + r/r_{\rm s})^2} \,, \tag{3.20}$$

where δ_c and r_s are two free parameters known as the characteristic overdensity and the scale radius, respectively. We set the normalisation to the critical density at the redshift of the halo $\rho_{\rm cr} \equiv 3H^2(z)/(8\pi G)$ with H(z) the Hubble parameter at that same redshift and G the gravitational constant. With the definition of the virial radius, r_{200c} , the radius inside which the mean mass density of the halo equals $200\rho_{\rm cr}$, we can construct a so-called concentration parameter $c \equiv r_{200c}/r_s$ and relate it to δ_c through

$$\delta_c = \frac{200}{3} \frac{c^3}{\ln(1+c) - c/(1+c)} \,. \tag{3.21}$$

In practice, we used mvir_subhalo, the virial mass of the subhalo from the SURFS-SHARK simulation²⁰, to calculate the virial radius for each lens. For the concentration parameter, we adopted the concentration—mass relation from Duffy et al. (2008)

$$c = 7.85 \left(\frac{M_{\text{vir}}}{2 \times 10^{12} \ h^{-1} \text{M}_{\odot}} \right)^{-0.081} (1+z)^{-0.71} \ . \tag{3.22}$$

²⁰https://shark-sam.readthedocs.io/en/latest/output_files.html

We note that Eq. (3.22) is estimated from *N*-body simulations based on a WMAP5 cosmology (Komatsu et al. 2009), which has slightly different parameter values from the Planck Collaboration (2016) cosmology used by the SURFS simulations. Nevertheless, the weak-lensing shear amplitude is dominated by the enclosed mass of the lens but has minor sensitivity to the concentration (e.g., Viola et al. 2015). Therefore, we ignored any potential WMAP5-to-Planck cosmology correction to Eq. (3.22).

Recognising the spherically symmetric feature of the NFW profile, we can derive the radial-dependent tangential shear as (Bartelmann 1996; Wright & Brainerd 2000):

$$\gamma_{\rm t}(x) = \frac{\rho_{\rm cr} \, \delta_c \, r_{\rm s}}{\Sigma_{\rm cr}} g(x) \,, \tag{3.23}$$

where $x \equiv R_{\rm FB}/r_{\rm s}$ is a dimensionless radial distance factor defined as the ratio of $R_{\rm FB}$, the projected radial separation between the lens and the source, to the scale radius of the lens. The critical surface mass density

$$\Sigma_{\rm cr} \equiv \frac{c^2}{4\pi G} \frac{D_{\rm a,B}}{D_{\rm a,F} D_{\rm a,FB}} \tag{3.24}$$

is a geometric term depending on the angular diameter distances to the source $D_{a,B}$, to the lens $D_{a,F}$ and between the lens and the source $D_{a,FB}$. The radial dependence of the shear is captured by the function g(x) as

$$g(x) = \frac{4}{x^2} \ln\left(\frac{x}{2}\right)$$

$$+\begin{cases} \frac{2}{1-x^2} + \frac{8-12x^2}{x^2(1-x^2)^{3/2}} \operatorname{arctanh}\sqrt{\frac{1-x}{1+x}} & (x < 1) \\ 10/3 & (x = 1) \end{cases}.$$

$$\frac{2}{1-x^2} + \frac{12x^2 - 8}{x^2(x^2 - 1)^{3/2}} \operatorname{arctanh}\sqrt{\frac{x-1}{1+x}} & (x > 1)$$

With all these ingredients in hand, we can now assign galaxy a specific shear value based on its redshift and neighbouring conditions. In summary, those identified as foreground galaxies only contain the redshift-dependent mean amplitude $\gamma_F(z_F)$, whilst the background galaxies combine the cosmic shear from Eq. (3.18) and the tangential shear from Eq. (3.23). This treatment accounts for not only the redshift-shear dependence but also the correlations between the blended objects.