

KiDS-1000 and DES-Y1 combined: Cosmology from peak count statistics

Joachim Harnois-Déraps^{1*}, Sven Heydenreich², Benjamin Giblin³, Nicolas Martinet⁴, Tillman Tröster⁵, Marika Asgari⁶, Pierre Burger^{7,8,9}, Tiago Castro^{10,11,12,13}, Klaus Dolag¹⁴, Catherine Heymans^{3,15}, Hendrik Hildebrandt¹⁵, Benjamin Joachimi¹⁶ & Angus Wright¹⁵

¹*School of Mathematics, Statistics and Physics, Newcastle University, Herschel Building, NE1 7RU, Newcastle-upon-Tyne, UK*

²*Department of Astronomy and Astrophysics, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064 USA*

³*Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh, EH9 3HJ, UK*

⁴*Université d'Aix-Marseille, CNRS, CNES, LAM, Marseille, France*

⁵*Institute for Particle Physics and Astrophysics, ETH Zürich, Wolfgang-Pauli-Strasse 27, 8093 Zürich, Switzerland*

⁶*Excellence for Data Science, AI, and Modelling (DAIM), University of Hull, Cottingham Road, Kingston-upon-Hull, HU6 7RX*

⁷*Waterloo Centre for Astrophysics, University of Waterloo, Waterloo, ON N2L 3G1, Canada*

⁸*Department of Physics and Astronomy, University of Waterloo, Waterloo, ON N2L 3G1, Canada*

⁹*Argelander-Institut für Astronomie, Auf dem Hügel 71, 53121 Bonn, Germany*

¹⁰*INAF - Osservatorio Astronomico di Trieste, via Tiepolo 11, I-34131 Trieste, Italy*

¹¹*INFN - Sezione di Trieste, Via Valerio 2, 34100 Trieste TS, Italy*

¹²*IFPU - Institute for Fundamental Physics of the Universe, via Beirut 2, 34151, Trieste, Italy*

¹³*ICSC - Centro Nazionale di Ricerca in High Performance Computing, Big Data e Quantum Computing, Via Magnanelli 2, Bologna, Italy*

¹⁴*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild Strasse 1, 85748 Garching, Germany*

¹⁵*Ruhr University Bochum, Faculty of Physics and Astronomy, Astronomical Institute (AIRUB), German Centre for Cosmological Lensing, 44780 Bochum, Germany*

¹⁶*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

Accepted XYZ, 2023. Received XYZ, 2023; in original form XYZ, 2023

ABSTRACT

We analyse the fourth data release of the Kilo Degree Survey (KiDS-1000) and extract cosmological **parameter** constraints based on the cosmic shear peak count statistics. Peaks are identified in aperture mass maps in which the filter is maximally sensitive to angular scales in the range 2–4 arcmin, probing deep into the non-linear regime of structure formation. We interpret our results with a simulation-based inference pipeline, sampling over a broad w CDM prior volume and marginalising over uncertainties on shape calibration, photometric redshift distribution, intrinsic alignment and baryonic feedback. Our measurements constrain the structure growth parameter and the amplitude of the non-linear intrinsic alignment model to $S_8 \equiv \sigma_8 \sqrt{\Omega_m}/0.3 = 0.732^{+0.032}_{-0.032}$ and $A_{IA} = 0.71^{+0.42}_{-0.42}$, respectively, in agreement with previous KiDS-1000 results based on two-point shear statistics. These results are robust against modelling of the non-linear physics, different scale cuts and selections of tomographic bins. The inferred cosmology is also consistent with the Dark Energy Survey Year-1 peak count analysis presented in Harnois-Déraps et al. (2021), and hence we jointly analyse both surveys with a common pipeline, obtaining $S_8^{\text{joint}} = 0.732^{+0.020}_{-0.020}$, one of the tightest constraints to date on this quantity, in 4σ tension with *Planck*. **Residual biases in the photometric redshifts of the DES-Y1 data could lower this tension, however it is robust against other systematics.** We also report from the joint-survey analysis the tightest constraint on the dark energy equation-of-state from non-Gaussian cosmic shear statistics, achieving $w_0^{\text{joint}} = -1.12^{+0.42}_{-0.31}$, in agreement with the concordance Λ CDM scenario. [Comments GoogleDoc: <https://docs.google.com/document/d/1i8yOL0mWnMXPDtxCCiybKDUVHDkJ-dMnjomVBNTOKpE/edit?usp=sharing>]

Key words: Gravitational lensing: weak – Methods: data analysis, numerical – Cosmology: dark matter, dark energy & cosmological parameters

1 INTRODUCTION

Cosmic shear cosmology has entered an era of high precision, with recent measurements from the Kilo Degree Survey¹ (KiDS), the Dark Energy Survey² (DES) and the Hyper Suprime Camera Survey³ (HSC) reaching a precision of a few percent on parameters central to the standard model of cosmology (e.g. Asgari et al. 2021; van den Busch et al. 2022; Li et al. 2023a; Secco et al. 2022b; Amon et al. 2022; Dalal et al. 2023; Li et al. 2023b). Based on the detection of weak correlations between the observed shapes of galaxies imparted by the foreground large scale structure, cosmic shear is mostly sensitive to the structure growth parameter $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$, a combination of the matter density **parameter** Ω_m and of the amplitude of the linear matter power spectrum smoothed on spheres of $8h^{-1}\text{Mpc}$, labelled as σ_8 (for lensing reviews, see e.g. Kilbinger 2015; Mandelbaum 2018). These Stage-III lensing surveys have been steadily improving the data quality and the analysis methods, in preparation for the next generation of cosmic shear experiments such as the **Rubin observatory**⁴ (Ivezić et al. 2019), *Euclid*⁵ (Laureijs et al. 2011) and the **Nancy Grace Roman space telescope**⁶ (Akeson et al. 2019).

Despite the large effort that is being invested by international collaborations in constructing accurate lensing catalogues of hundreds of millions of galaxies, it is not entirely clear how to best analyse these vast data, striking an optimal compromise between accuracy and precision. To date the shear two-point (2pt) **functions** are still regarded as the baseline summary statistics, having been tested for over a decade and achieving an unmatched level of understanding and control in all aspects of the analysis, including measurements tools (e.g. TREECORR and NAMASTER, see Jarvis et al. 2004; Alonso et al. 2019), theoretical predictions (e.g. Kilbinger et al. 2017) and the impact of systematics (see e.g. Mandelbaum 2018). The main drawback from these statistics is that they completely disregard the non-Gaussian information that is stored in the non-linear matter field, more precisely in the coupling between the phases of distinct Fourier modes, without which the cosmic web would look like a Gaussian random field. This is obviously sub-optimal, and this waste of information **will be aggravated** in the upcoming cosmic shear experiments. Accessing this non-Gaussian information is an active field of research: an array of novel weak lensing statistics are being developed specifically to utilise this complementary small-scale information. These new methods are reaching a level of maturity that makes them competitive at analysing existing cosmic shear data, carefully balancing the precision vs accuracy metric. Recent progress is largely due to the radically improved modelling of the signal, thanks to the increased accuracy of cosmological N -body codes and the availability of super-computers (see Angulo & Hahn 2022, for a recent review on N -body codes). Recent examples of these ‘beyond-2pt’ cosmic shear data analyses include the three-point function (Fu et al. 2014; Secco et al. 2022a; Burger et al. 2023), peak count statistics (Kacprzak et al. 2016; Martinet et al. 2018; Shan et al. 2018; Harnois-Déraps et al. 2021; Zürcher et al. 2022; Liu et al. 2023; Gatti et al. 2023a; Marques et al. 2023), density split statistics (Brouwer et al. 2018; Gruen et al. 2018; Burger et al. 2022), shear clipping (Giblin et al. 2018), persistent homology (Heydenreich et al. 2022), moments of convergence maps (van Waerbeke et al. 2013; Gatti et al. 2020), cumulative distribution functions (Anbajagane et al. 2023), **likelihood-free inference** (Jeffrey et al. 2021; Lin et al. 2023; Gatti et al. 2023b) or convolutional neural network inference (Fluri et al. 2019, 2022).

At the moment, these alternative methods exhibit a constraining power that is similar to that of two-point functions, which is not surprising given the noise levels of current lensing data, which make difficult the extraction of information stored in the noisy higher-order moments. The situation will change drastically with the upcoming surveys, where the cosmic web itself will be detectable with lensing, at which point the non-Gaussian information **will take on a larger proportion of the signal**.

All forecasts are clear about this: joint cosmic shear analyses that combine two-point functions and any complementary probe improve the constraints on cosmological parameters even in presence of systematic uncertainties (e.g. Li et al. 2019; Schneider et al. 2019; Zürcher et al. 2020; Pyne & Joachimi 2021; Harnois-Déraps et al. 2022; Giblin et al. 2023; Euclid Collaboration: Ajani et al. 2023). The main difficulty in many of these methods comes from their accrued dependence on numerical simulations, which adds a significant computational overhead to the data analysis compared to those for which an analytical model exists. Typically, simulations are needed for modelling the cosmological signal, for modelling some of the systematics **such as baryonic feedback or intrinsic alignments of galaxies**, and for the estimation of the covariance matrix (although this is **not always necessary, as demonstrated by the recent likelihood-free inference analyses mentioned above**).

This paper contributes an important step to this effort: we carry out a cosmological analysis based on lensing peak statistics measured from the fourth data release of the Kilo Degree Survey (KiDS-1000 hereafter). We use the exact same data as those used in the two-point function analyses of Asgari et al. (2021, A21 hereafter), while ignoring for now other non-lensing KiDS galaxy catalogues designed for galaxy clustering analyses (Vakili et al. 2020; Bilicki et al. 2021). Our method finds peaks in aperture mass maps with an aperture filter designed for the extraction of small-scale structure, with maximal sensitivity to scales of less than 4 arcmin, as in Martinet et al. (2018, hereafter M18) and Harnois-Déraps et al. (2021, HD21). This contrasts with the recent peak count analysis of Zürcher et al. (2022), in which peaks are extracted from convergence maps with pixel resolution of about 7 arcmin. Both methods have their advantages and downsides, ours strongly focuses on small, non-linear scales, which, as demonstrated in HD21 and Martinet et al. (2020), have a higher potential for complementarity with two-point functions. Finding a posterior that is statistically consistent with that from HD21, we combine both likelihoods and carry out a joint KiDS-1000 + DES DR1 data (DES-Y1 hereafter) peak count analysis, finding the tightest constraints on S_8 to date from peaks alone.

After describing the data and simulations in Sec. 2, we detail our measurement techniques and analysis pipeline in Sec. 3, and we present our mitigation strategy for the key systematic uncertainties in Sec. 4. We finally show our results in Sec. 5 and discuss our findings afterwards. Supplementary material is provided in the Appendices, including a thorough discussion of B -modes (in Appendix A), supplementary pipeline validation tests (in Appendix B) and a detailed discussion on goodness-of-fit for noisy covariance matrices (in Appendix C).

2 DATA AND SIMULATIONS

We present in this section the survey data and the various simulation suites that are used for the cosmological analysis.

2.1 KiDS-1000 data

The Kilo Degree Survey (Kuijken et al. 2015) is a multi-band photometric galaxy survey explicitly designed for weak lensing cosmology. Carried out at the European Southern Observatory by the VST-

¹ KiDS: kids.strw.leidenuniv.nl

² DES: www.darkenergysurvey.org

³ HSC: www.naoj.org/Projects/HSC

⁴ LSST: www.lsst.org

⁵ Euclid: www.euclid-ec.org

⁶ WFIRST: roman.gsfc.nasa.gov

OmegaCAM, we analyse here the public⁷ fourth data release (Kuijken et al. 2019). The observation conditions are of exceptional quality, with a mean seeing of 0.7 arcsec in the r -band, used for shape measurements. The photometric redshifts are obtained from a combination of nine optical and infrared bands ($ugriZYZJHK_s$, see Wright et al. 2020), thanks to the observations of the companion VIKING survey (VISTA Kilo-degree INfrared Galaxy, Edge et al. 2013). The galaxies selected in this analysis exactly match those used in the cosmic shear two-point function analyses of A21 and van den Busch et al. (2022), covering an effective area of 777.4 deg².

The KiDS DR4 data are reduced with the THELI (Erben et al. 2013) and Astro-WISE (Begeman et al. 2013) pipelines, following which the shear is inferred from *lensfit* (Miller et al. 2013; Fenech Conti et al. 2016). Shear additive and multiplicative biases (c - and m -corrections) are measured to a high accuracy (Giblin et al. 2021), where it is shown via a series of null tests that known residual systematics in the shear measurement could lead to no more than a 0.1σ shift in the structure growth parameter $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$, the composite quantity that is best measured by cosmic shear. Note that strictly speaking, the results from the tests carried out in Giblin et al. (2021) are only shown to hold for two-point cosmic shear statistics.

Following A21, we split the full DR4 galaxies in five tomographic bins according to their individual best-fit redshift z_B as measured by BPZ (Benítez 2000), with bin edges set to [0.1, 0.3, 0.5, 0.7, 0.9 and 1.2]. The tomographic redshift distributions, $n^a(z)$, are estimated via self-organising maps (SOM, see Wright et al. 2020), which group galaxies based on their nine-band photometric properties and assign redshifts based on similar studies made on spectroscopic samples; galaxies for which no match is found are rejected. We further reject galaxies for which the SOM redshift catastrophically differs from the initial z_B , resulting in the so-called ‘Gold Sample’ introduced in Hildebrandt et al. (2021) and used in the subsequent KiDS-1000 cosmic shear analyses mentioned above. As detailed in A21, the means and the error of the SOM redshift distributions are calibrated on KiDS-like mock data constructed from the MICE2 simulations (Fosalba et al. 2013; van den Busch et al. 2020) and accounted for during the inference stage of our analysis. The redshift accuracy is excellent due to the nine-band photometry, which helps breaking degeneracies in the galaxy spectral energy distributions: at worst, the difference between the mean redshift and that estimated from the matched spectroscopic sample is $z_{\text{est}} - z_{\text{true}} = 0.013 \pm 0.0118$, making this a sub-dominant source of uncertainty in our measurement. Note that Hildebrandt et al. (2021) further show that the SOM redshift distributions are fully consistent with independent estimates based on clustering cross-correlations with spectroscopic reference samples, providing extra robustness to the method. Fig. 1 shows the redshift distributions estimated in the five tomographic bins, along with the variations on these distributions allowed within our photometric uncertainty.

The SOM selection and the shear inference pipelines are both repeated on KiDS-like image simulations (Kannawadi et al. 2019), from which a relation between apparent size, magnitude and the observed galaxy shape is used to calibrate the inferred *lensfit* shear⁸. Whereas previous cosmological analyses use a single m -calibration factor per tomographic bin, the aperture mass map statistics exploited in this paper are subject to local variations in the noise levels and seeing conditions, and we therefore use the above-mentioned relation to extract a shear calibration *per object*, m_a . This is not necessary, but allows to capture




Figure 1. Tomographic redshift distribution of the KiDS-1000 data. The thinner lines represent the effect of photometric uncertainty on these distributions, characterised by $n^a(z) \rightarrow n^a(z + \Delta z_a)$, with Δz_a sampled 10 times from Gaussian distributions with widths listed in Table 1. All shifted $n^a(z)$ are then rebinned with the same z bins.

Table 1. Main properties of the KiDS-1000 data used in the current work. The Gold Sample redshift selection based on z_B is identical to that presented in Hildebrandt et al. (2021). The effective number densities are listed in the second column, in gal/arcmin². The shape noise (per component) listed in the third column reflects the dispersion measured in the observed galaxy shapes, as documented in Giblin et al. (2021), while the fifth column shows the mean shape calibration coefficients. **The redshift bias and errors** listed in the fourth column are estimated from the SOM method in Hildebrandt et al. (2021), while the last column shows the additive $c_{1/2}$ terms, **which has an uncertainty of 0.23×10^{-3}** (Giblin et al. 2021).

tomo	n_{eff}	σ_ϵ	$z_{\text{est}} - z_{\text{true}}$	m	$(c_1, c_2) \times 10^3$
bin1	0.62	0.27	0.000 ± 0.0106	-0.009 ± 0.019	(0.295, 0.156)
bin2	1.18	0.26	0.002 ± 0.0113	-0.011 ± 0.020	(0.004, 0.621)
bin3	1.85	0.27	0.013 ± 0.0118	-0.015 ± 0.017	(0.052, 0.728)
bin4	1.26	0.25	0.011 ± 0.0087	0.002 ± 0.012	(-0.360, 0.948)
bin5	1.31	0.27	-0.006 ± 0.0097	0.007 ± 0.010	(-1.363, 1.155)

possible correlations between the m -correction and the *lensfit* weights. These are inevitably noisier than the average over the full tomographic bins, but a large fraction of this noise cancels within our aperture mass map calculations as well, while providing optimal estimates of the local noise contribution (M18). Let us recall that this calibration corrects for known residual biases such as shape detection biases (Fenech Conti et al. 2016; Kannawadi et al. 2019) or blending of the images of galaxies (Hoekstra et al. 2015). While we apply the m -correction per object, the averaged multiplicative biases per redshift bin used in A21 enter our analysis at the inference level in the form of nuisance parameters over which we marginalise. Table 1 summarises the survey properties relevant to our analysis.

2.2 DES-Y1 data

The DES-Y1 measurements is based on the public year-1 data release from the Dark Energy Survey Collaboration (Abbott et al. 2018), with source galaxy selections that exactly follow the main cosmic shear results described in Troxel et al. (2018). The lensing catalogue consists of 26 million galaxies covering a footprint of 1320 deg² with a galaxy density of 5.07 gal arcmin⁻². The per-galaxy shear signal is inferred with the METACALIBRATION method (Sheldon & Huff 2017). Every galaxy is assigned to one of the four tomographic bins based on the photometric redshift posteriors estimated from the $griz$ flux measurements, as detailed in Hoyle et al. (2018). Following Troxel et al. (2018), the uncertainty on the shear multiplicative calibration is given by $m_a = 0.012 \pm 0.023$.

Whereas the original DES-Y1 results estimated the tomographic

⁷ KiDS-1000 data: <http://kids.strw.leidenuniv.nl/DR4>

⁸ A KiDS-1000 re-analysis has been presented in Li et al. (2023a) after correcting for an anisotropic error in the *lensfit* likelihood sampler. This error has not been corrected here, but their study shows the correction has a negligible impact on the inferred cosmology.

$n^a(z)$ from a Bayesian photometric redshift analysis calibrated on the COSMOS2015 field (Laigle et al. 2016), the HD21 reanalysis instead opted for $n^a(z)$ estimates based on a direct reweighted calibration of matched spectroscopic data (Lima et al. 2008, DIR hereafter), following the DES-Y1 reanalyses of Joudaki et al. (2020) and Asgari et al. (2020). The uncertainty on the DIR mean redshift distributions is $\Delta z_a = [0.008, 0.014, 0.011 \text{ and } 0.009]$ for redshift bins $a = 1...4$, respectively. **Both methods have their pros and cons. The calibration with COSMOS is by design based on a complete sample but suffers from imperfect redshifts (see e.g. Alarcon et al. 2020). In contrast, the spec- z samples used for the DIR method have (close to) perfect redshifts but are incomplete and not representative of the source sample, which is alleviated by the reweighting, but often cannot be fully eliminated (see Gruen & Brimiouille 2017). Importantly, the DIR $n^a(z)$ favours S_8 values that are smaller by $\Delta S_8 = 0.03$ compared to the COSMOS-calibrated $n(z)$, which is a 0.8σ shift (Joudaki et al. 2020).**

2.3 Simulations

As mentioned in the introduction, the accuracy of simulation-based inference pipelines fully depends on the quality of the numerical simulations it is calibrated on. The same way 2pt analyses must carefully understand the scales, cosmologies and redshifts that are well captured by their model, it is critical for our peak count analysis to identify the range of validity of our training simulations. The additional complexity here is that no simulation suite serves all purposes, and therefore we must carefully investigate, for all of them separately, the accuracy and limits of the measurements and how these impact the peak count statistics. The simulations used in this work are in many aspects identical to those presented in HD21, which we refer to for further details. Specifically:

(i) the cosmological dependence of the peak count statistics is calibrated on the w CDM *cosmo*-SLICS N -body simulations introduced in Harnois-Déraps et al. (2019). They sample a wide volume in S_8 , Ω_m , w_0 and h with 25 points arranged in a Latin hypercube (plus one Λ CDM point), each evolved with a pair of N -body simulations designed to suppress sample variance in 2pt functions, then ray-traced in ten light-cones of 100 deg^2 (10000 deg^2 in total area). These form our *Cosmology Training Set*, and resolve the non-linear physics to better than 2% up to k -modes of $2.0 h^{-1} \text{ Mpc}$, when compared to the Cosmic Emulator (Heitmann et al. 2014). Smaller scales gradually lose precision, affecting mainly their ability to resolve substructure in most massive objects. The exact impact of this loss on weak lensing peak counts is investigated in HD21 with a separate set of simulations ran with a much higher force resolution, where it is found that this leads to at most a 1% loss of the highest peaks, which is largely sub-dominant compared to both baryonic physics and statistical errors. We revisit this in Sec. 4 (see also point iv);

(ii) the covariance matrix that captures the sample variance is estimated from 124 fully independent SLICS N -body simulations described in Harnois-Déraps & van Waerbeke (2015). These are evolved from independent initial conditions at a fixed cosmology, and make our *Covariance Training Set*. They resolve the same non-linear physics as the *cosmo*-SLICS, and are shown in Harnois-Déraps et al. (2019) and HD21 to produce marginalised errors on cosmological parameters that are fully consistent with those obtained with an analytical calculation, when analysing 2pt statistics. Burger et al. (2022) further show in the context of density-split statistics that a covariance matrix estimated from the SLICS or from a much larger number of log-normal FLASK mocks (Xavier et al. 2016) produce fully consistent results, as expected for these mildly non-linear statistics. We further increase the effective number of covariance mocks by randomly rotating 10 times

the shape noise components. This works particularly well given that the peak statistics is currently shape-noise dominated: while the expectation value of standard 2pt statistics does not depend on the noise (only their variance does), shape noise affects both the signal and covariance of map-based statistics (see Appendix D of Heydenreich et al. 2021);

(iii) for the KiDS-1000 analysis, the impact of galaxy intrinsic alignments is measured from the IA-infused lensing simulations described in Harnois-Déraps et al. (2022). These are also constructed from the *cosmo*-SLICS and therefore resolve the same physical scales. This *IA Training Set* assumes a linear coupling between the projected **non-linear** tidal field and the intrinsic ellipticity of every galaxy, and is therefore physically modelling the non-linear **linear** alignment model of Bridle & King (2007) without explicit redshift nor luminosity dependence. It is expected that this effective IA model does not fully capture the alignment signal, and that a more physical model such as the Tidal Alignment And Torquing model (Blazek et al. 2019) or the **halo-model of Fortuna et al. (2020)** would provide a more accurate description, however current cosmic shear surveys do not have the statistical power to constrain parameters beyond the simpler NLA model (Secco et al. 2022b), which is therefore deemed sufficient for the current analysis. The IA infusion process has been shown in Harnois-Déraps et al. (2022) to accurately reproduce the NLA predictions for the 2pt correlation function down to scales of a few arcmin, beyond which the NLA is expected to fail in a manner that is undetectable in the current data. **Burger et al. (2023) further show that these same simulations agree with the IA modelling of three-point shear statistics.** The model fails at scales that correspond to high over-densities in our simulations, which contribute to lensing peaks that are excluded from our analysis. We infuse different levels of IA and marginalise over these choices in the end, as described in Sec. 4. for the DES-Y1 analysis, IA are included with a non-linear halo-based model, see HD21 for details;

(iv) limits in the force resolution of the *cosmo*-SLICS are bound to impact the weak lensing statistics in a manner that is not always predictable. We assess this with the SLICS-HR suite (Harnois-Déraps & van Waerbeke 2015), a high-resolution version of the SLICS light-cones recently used in a combined lensing-clustering cosmological analysis (Duncan et al. 2022). The SLICS-HR consist of ten independent $10 \times 10 \text{ deg}^2$ catalogues that are run at the same cosmology and with the same particle count and volume as the SLICS, but the N -body force accuracy has been increased such as to resolve k -modes up to $10 h^{-1} \text{ Mpc}$. We use these to validate the full inference pipeline in Sec. 4.7, acting as our *Validation Set*;

(v) the impact of baryon feedback is estimated with the *Magneticum* hydrodynamical simulations⁹, forming our *Baryons Training Set*. These have been shown to reproduce a number of key observations relevant to weak lensing studies (Castro et al. 2018), and notably the feedback on the matter distribution closely matches that of the BAHAMAS (McCarthy et al. 2017), another suite of hydro simulations with independent prescriptions for their sub-grid physics. The training set consists of ten $10 \times 10 \text{ deg}^2$ *pseudo*-independent light-cones extracted from full hydrodynamical simulations, and another 10 light-cones extracted from dark matter-only sister simulations, evolved from the same initial conditions (more details on the used simulations can be found in Martinet et al. 2021). There is a large uncertainty on the exact impact of baryonic physics on the matter distribution (and therefore on our lensing statistics), which we account for by linearly scaling the relative baryonic bias with a nuisance parameter, b_{bary} , which we marginalise over at the inference stage¹⁰.

⁹ Magneticum simulations: www.magneticum.org

¹⁰ Note that this parameter is not to be confused with A_{bary} used in A21 (see their Table 2), which specifically relates to one of the free parameters entering their HMCode halo model.




Figure 2. Tiling strategy adopted to pave the full KiDS-1000 data (black) with flat-sky $10 \times 10 \text{ deg}^2$ simulations (red squares). Some of the tiles slightly overlap due to the sky curvature, in which case the data is split at the mean Dec in overlapping regions.

2.3.1 Assembling mock surveys

Most of these simulations have been introduced in HD21, in Heydenreich et al. (2022), in Burger et al. (2022) and in the references listed in the previous section; we encourage the interested reader to consult these for a more complete technical description. To summarise some of the key properties, all of the above-mentioned simulations are organised in light-cones of 100 deg^2 each, populated with galaxy samples that match the tomographic $n^a(z)$ distributions, number densities and shape noise levels of the KiDS-1000 Gold Sample and DES-Y1 data. Except for the IA-infused simulations, the galaxy positions, the amplitude of their ellipticities $|\epsilon_{\text{data}}|$ and their multiplicative shear calibration factors m_a are exactly reproduced in each of the mock survey realisations (i.e. in the *Cosmology*, *Covariance*, *Validation* and *Baryon Training Sets*). To achieve this, the KiDS-1000 data are split into 18 tiles that each fit within 100 deg^2 regions, as depicted in Fig. 2, and the shear and convergence from every simulation is repeated across them, interpolated at the local galaxy positions. These tiles are analysed separately and combined only at the level of the summary statistics, ensuring that cross-tile correlations that exist in the data but not in the simulations are explicitly ignored. This effect is minor for localised non-Gaussian probes such as peak statistics, but is critical for e.g. shear 2pt functions. The shear and convergence are interpolated from the underlying simulations at the position of every galaxy, then combined with the (randomly rotated) observed ellipticity following:

$$\epsilon_{\text{mock}} = \frac{\epsilon_{\text{data}}^{\text{rand}} + \mathbf{g}}{1 + \epsilon_{\text{data}}^{\text{rand}} \mathbf{g}^*}. \quad (1)$$

In the above expressions, bold-font symbols are spin-2 complex quantities and $\mathbf{g}$ is the simulated reduced shear. As described in HD21, this involves rotating each tile at the equator, which preserves the relative positions of galaxies but modifies their ellipticities, defined with respect to the North pole.

We repeat this construction for all light-cones of the *Cosmology Training Set*, the *Covariance Training Set*, the *Baryon Training Set* and the *Validation Set*. Additionally, the uncertainty in the photometric redshifts is forward-modelled with a further 10 full survey realisations computed at the fiducial cosmology, in which the $n(z)$ is shifted by small amounts (details provided in Sec. 4). In total, this results in 414 simulated mosaic surveys that we analyse in preparation for the inference stage, with the majority (260) contributing to the *Cosmology Training Set*. Each mock further contains 10 random rotations of ϵ_{data} , to improve convergence of the signal.

The *IA Training Set* are treated slightly differently, since for these the positions of the mock galaxies must be sampled from the simulated over-density maps or halo catalogues, which do not correlate with the positions in the data (see Harnois-Déraps et al. 2022, for more details).

The mosaic survey tiling is therefore not possible, so we use instead 5 light-cones per IA model and explore four alignment strengths in KiDS, and one model in DES. Although these represent a lesser total area than the real data, their sole purpose is to capture the relative impact of IA on the signal, computed from ratios in which the sample variance cancels by design. We use $4 \times 5 \times 100 \text{ deg}^2$ of training data, which is enough to capture this.

The simulations are free of additive biases by construction, however Giblin et al. (2021) measures residual additive terms $c_{1/2}$ in the KiDS-1000 cosmic shear catalogues, caused by the shape measurement method itself. These are reported in Table 1 and subtracted from the observed ellipticities when analysing the real data. **We follow Troxel et al. (2018) by not accounting for any low-level additive terms in the DES Y1 catalogue.** The multiplicative biases are not easily removed from the data, hence we instead infuse the mocks with the m_a terms per object, and treat thereafter data and simulations on equal footings.

3 METHODS

3.1 Aperture mass map statistics

There exists a number of methods for identifying and counting lensing peaks, including finding maxima on convergence maps (Li et al. 2019), on wavelet-transformed maps (Ajani et al. 2020) or on aperture mass maps (Schneider 1996). We here opted for the aperture mass maps for the following reasons: as argued in M18, this statistics is immune to masking-induced biases and strong B -mode leakage common to methods based on reconstruction of convergence maps, plus it benefits from a local estimation of the shape and Poisson noise, yielding more accurate signal-to-noise maps.

Specifically, we cover each of the 18 tiles with a two-dimensional grid with a pixel size of 0.59 arcmin. We next reconstruct the mass inside an aperture filter Q centred on each pixel, at position θ on the sky, from the sum of all tangential ellipticities $\epsilon_{a,t}$ contained therein as:

$$M_{\text{ap}}(\theta) = \frac{1}{n_{\text{gal}}(\theta) \sum_a w_a (1 + m_a)} \sum_a w_a \epsilon_{a,t}(\theta, \theta_a) Q(|\theta - \theta_a|, \theta_{\text{ap}}, x_c). \quad (2)$$

The tangential ellipticity about θ is computed as $\epsilon_{a,t}(\theta, \theta_a) = -[\epsilon_1(\theta_a) \cos(2\phi(\theta, \theta_a)) + \epsilon_2(\theta_a) \sin(2\phi(\theta, \theta_a))]$, where θ_a is the position of galaxy a and $\phi(\theta, \theta_a)$ is the angle between both coordinates. The sum runs over all galaxies in the aperture, and $n_{\text{gal}}(\theta)$ is the local galaxy density in the filter when centred at θ . As in M18 and HD21, our filter $Q(\theta, \theta_{\text{ap}}, x_c)$, abridged to $Q(\theta)$, matches that of Schirmer et al. (2007), which is optimised for efficiently detecting NFW haloes:

$$Q(x) = \frac{\tanh(x/x_c)}{x/x_c} [1 + \exp(6 - 150x) + \exp(-47 + 50x)]^{-1}. \quad (3)$$

In the above expression, we use the standard value of $x_c = 0.15$, while $x = \theta/\theta_{\text{ap}}$, with θ the distance to the filter centre. We additionally use the same filter size, set to $\theta_{\text{ap}} = 12.5 \text{ arcmin}$, which is shown in M18 to better detect the cosmological signal over other filter sizes in KiDS data. We compute Eq. (2) at every pixel location to construct our signal map. The variance about this map is computed at every pixel location with:

$$\sigma_{\text{ap}}^2(\theta) = \frac{1}{2n_{\text{gal}}^2(\theta) [\sum_a w_a]^2} \sum_a w_a^2 |\epsilon_a|^2 Q^2(|\theta - \theta_a|), \quad (4)$$

where again the sum runs over all galaxies in the filter. The m -calibration estimated from the image simulations of Kannawadi et al. (2019) is meant to correct the inferred shear, not the ellipticity, which explains why it appears in the denominator of Eq. (2) but not in that of Eq. (4), which describes the noise map. Finally, we take the ratio between Eq. (2) and the square root of Eq. (4) at every pixel location to

construct our signal-to-noise maps, $S/N(\theta) \equiv M_{\text{ap}}(\theta)/\sqrt{\sigma_{\text{ap}}^2(\theta)}$, from which we identify peaks as local maxima and record their S/N -values. We repeat this process for the 10 realisations of random rotations and report the average, except for the *Covariance Training Set*, for which we do not take the average; instead, each noise realisation leads to an estimate of the covariance matrix, of which we take the average in the end.

As detailed in HD21, masking is dealt with naturally in aperture mass statistics, and no special treatment needs to be enforced as long as data and simulations are masked and analysed the same way. This is achieved by fixing galaxy positions in the simulations to that of the observed data, which ensures the impact of the mask is identical. In our case, we decided nevertheless to act upon masked pixels. These are identified from the galaxy catalogues as regions with an aperture galaxy density that is either critically low or null, then removed from the final $S/N(\theta)$ maps.

It has been shown that some additional information can be extracted by combining the peak count statistics measured from multiple filter sizes (e.g. Zürcher et al. 2022; Giblin et al. 2023), however M18 shows that this gain is mild for Stage III surveys. We therefore opted for a single-scale analysis here, but intend to revisit this in the future.

3.2 Tomography and selection

Tomographic decomposition of the lensing data allows us to probe the redshift evolution of the large scale structures, which is largely driven by Ω_m and w_0 via their impact on the growth of perturbations. A direct consequence of the improved sensitivity to these is a gain in precision in S_8 , arising from degeneracy breaking. This decomposition is different for the KiDS and DES surveys, which we detail here.

3.2.1 KiDS-1000

From the five KiDS tomographic bins, we include both the auto- and the cross-redshift measurements, as first defined in Martinet et al. (2020). To be specific, peaks are identified from the individual tomographic galaxy catalogues (the ‘auto’ redshift bins 1, 2, 3, 4, 5), from every possible combination of bin pairs (1U2, 1U3, 1U4, 1U5... 4U5), triplets (1U2U3, 1U2U4, 1U2U5...), quadruplets (1U2U3U4, 1U2U4U5...) and quintets (i.e. no tomography). As shown in HD21, Zürcher et al. (2022) and Heydenreich et al. (2022), these ‘cross-tomographic’ catalogues contain a significant amount of additional information that is not contained within the ‘auto’ case. The tomographic peak function is presented in Fig. 3, showing in the different panels the 30 different redshift bin combinations. For each case we overlay the predictions from the *Cosmology Training Set* in colour with the data measurements in black; the error bars are obtained from the *Covariance Training Set*. A similar measurement is presented in Fig B1, where the data is replaced by the mean over our *Baryons Training Set*; it was shown in HD21 that using instead the mean over the *Covariance Training Set* effectively validates that the ensemble average agrees with the predictions, a result we verified still holds with the KiDS-1000 mocks. In these figures, we have subtracted the peak function **measured from pure shape noise fields**, $N_{\text{peaks}}^{\text{noise}}$, to better highlight the cosmological dependence of the signal.

In all cases, we measure the peak function in S/N bins of width 0.5 in the range $[-2.5, 4.0]$, for a total of 13 bins per sub-panel and 390 elements in total. The motivation behind this initial choice of range is driven by a number of requirements, notably that of having a large number of peaks per bin to ensure the data is Gaussian-distributed (with our selection, every bin has at least 200 objects, while bins outside this range have far fewer objects). Additionally, our analysis has strict requirements on the modelling precision and on the level of contamination by residual systematic effects, resulting in this bin selection being

Table 2. Priors used in the KiDS likelihood sampling. The ranges for the four cosmological parameters are determined by the cosmo-SLICS simulations, while the prescription from sampling the nuisance parameters describing the photometric redshifts Δz_a and intrinsic alignments A_{IA} are taken from Joachimi et al. (2020). In particular, the redshift parameters are correlated and drawn from a **multivariate Gaussian distribution with means μ taken from Table 1 (fourth column) and a covariance matrix C_z described in Sec. 4.3**. The shear calibration parameters Δm_a are sampled from Gaussian priors characterised by a mean and a standard deviation (μ, σ) estimated in Giblin et al. (2021). The baryonic feedback parameter b_{bary} is used to scale the effect measured in the *Baryon Training Set*.

Parameter	range	prior
Cosmology		
Ω_m	[0.1, 0.55]	Flat
S_8	[0.6, 0.9]	Flat
h	[0.6, 0.82]	Flat
w_0	[-2.0, -0.5]	Flat
Nuisance		
$\Delta z_a \times 10^2$	[-10, 10]	$\mathcal{G}(\mu, C_z)$
$\Delta m_1 \times 10^2$	[-10, 10]	$\mathcal{G}(-0.9, 1.9)$
$\Delta m_2 \times 10^2$	[-10, 10]	$\mathcal{G}(-1.1, 2.0)$
$\Delta m_3 \times 10^2$	[-10, 10]	$\mathcal{G}(-1.5, 1.7)$
$\Delta m_4 \times 10^2$	[-10, 10]	$\mathcal{G}(0.2, 1.2)$
$\Delta m_5 \times 10^2$	[-10, 10]	$\mathcal{G}(0.7, 1.0)$
Astrophysics		
A_{IA}	[-5, 5]	Flat
b_{bary}	[0, 2]	Flat

in fact “aggressive”. We expand on this in Sec. 4, where we argue that instead the range $[-1.0 < S/N < 3.0]$ is a better choice with lower modelling errors, forming a ‘clean’ data vector of $7 \times 30 = 210$ elements in total that is used for the main cosmological analysis.

3.2.2 DES-Y1

Following HD21, our DES peak count analysis includes the auto- and cross-redshift measurements up to pairs of tomographic bins, for a total of 10 bin combinations. The peak function is measured in 12 S/N bins in the range $[0.0 < S/N < 4.0]$, forming a data vector with 120 elements. Although these details differ compared to the KiDS-1000 case described above, it is shown in HD21 to be accurate and competitive.

3.3 Analysis pipelines

Our cosmological inference pipeline heavily builds on the methods presented in HD21 and Heydenreich et al. (2022), which we briefly overview here. First, we model the peak function by training a Gaussian Process Regression¹¹ emulator (GPR) on the measurements from the *Cosmology Training Set*. The GPR can subsequently produce $N_{\text{peaks}}^{\kappa}$ predictions within a fraction of a second everywhere inside the parameter volume covered by the cosmo-SLICS. This therefore determines the prior ranges over Ω_m, S_8, w_0 and h , which we report in Table 2.

Second, we must estimate the covariance matrix, which captures the correlation between the elements of our data vector, central to the error propagation. As mentioned before, the *Covariance Training Set* consists of 124 full survey realisations, each replicated with 10 shape noise realisation, producing 1240 *pseudo-independent* data

¹¹ We use the GPR toolkit provided by SciKitLEARN (Pedregosa et al. 2011).




Figure 3. Tomographic weak lensing peak function $N^k_{\text{peaks}}(S/N)$ measured in the KiDS-1000 data (black squares) and in the *Cosmology Training Set* simulations, colour-coded by their S_8 value. The pure noise signal $N^{\text{noise}}_{\text{peaks}}$ has been removed to better highlight the variations with respect to cosmology. The panels show the results from different combinations of tomographic bins, in which the red dashed lines represents the best-fit model inferred from our fiducial analysis, see Sec. 5.

Table 3. Priors used for sampling the nuisance parameters in the DES-Y1 peak statistics analysis. The sampling of photometric redshifts Δz_a and shear bias Δm_a nuisance parameters follows the original cosmic shear paper by Troxel et al. (2018). The baryonic feedback parameter b_{bary} is the same as in the KiDS-1000 likelihood, however there is no IA parameter here.

Parameter	range	prior
$\Delta z_1 \times 10^2$	[-10, 10]	$\mathcal{G}(0.1, 1.6)$
$\Delta z_2 \times 10^2$	[-10, 10]	$\mathcal{G}(1.9, 1.3)$
$\Delta z_3 \times 10^2$	[-10, 10]	$\mathcal{G}(0.9, 1.1)$
$\Delta z_4 \times 10^2$	[-10, 10]	$\mathcal{G}(1.8, 2.2)$
$\Delta m_a \times 10^2$	[-10, 10]	$\mathcal{G}(1.2, 2.3)$
Astrophysics		
b_{bary}	[0, 2]	Flat

vectors from which our covariance matrix C is extracted. We show in Fig. 4 the cross-correlation coefficient matrix, defined as $C_{ij}/\sqrt{C_{ii}C_{jj}}$, which better highlights the correlations between the negative and pos-

itive peaks in each of the tomographic block. Also visible is the significant amount of correlation (and anti-correlation) present in the off-diagonal component. This matrix contains at most 390^2 elements and is thus invertible (since $390 < 1240$, see Hartlap et al. 2007), a criteria that is also naturally satisfied by the ‘clean’ KiDS-1000 data vector, which contains only 210 entries, and by the DES-Y1 data vector, which contains 120.

Having our model and covariance matrix, we are now in a position to evaluate the likelihood $\mathcal{L}$ of the model $\mathbf{x}(\boldsymbol{\pi})$ with parameters $\boldsymbol{\pi}$, given the data vector $\mathbf{d}$. We use the Sellentin & Heavens (2016) t -distribution likelihood, which is well suited for nearly Gaussian data vectors with simulation-based covariance matrices. It is constructed as:

$$\log \mathcal{L}(\boldsymbol{\pi}|\mathbf{d}) = \frac{N_{\text{sims}}}{2} \ln \left[1 + \chi^2 / (N_{\text{sims}} - 1) \right] + \text{const}, \quad \text{with} \quad (5)$$

$$\chi^2 = [\mathbf{x}(\boldsymbol{\pi}) - \mathbf{d}]^T C^{-1} [\mathbf{x}(\boldsymbol{\pi}) - \mathbf{d}]. \quad (6)$$

In the above, $N_{\text{sims}} = 1240$ is the number of realisations used to evaluate the covariance matrix C . The model depends on the four cosmological parameters Ω_m, S_8, w_0 and h , and on a set of 12 (9) astrophysical and




Figure 4. This figure highlights the correlations between the different elements of the KiDS-1000 data vector. From left to right, the 30 blocks show the correlation coefficients for the different redshift bin combinations, starting with singlets (i.e. auto-bins), pairs, triplets, quadruplet and the no-tomographic case, with redshift increasing towards the right and the top of the figure.

nuisance parameters for KiDS (DES), which characterise the dependence of our signal on the systematic effects mentioned previously. This is an excellent approximation to the more general likelihood suggested by Percival et al. (2022) in our case. Finally, the posteriors are sampled both by the nested sampling algorithm MULTINEST (Feroz et al. 2009) and by NAUTILUS (Lange 2023), implemented within CosmoSIS (Zuntz et al. 2015). While the latter sampler is more robust (Lange 2023), the former has been more widely used in the literature and is therefore useful to make fair comparisons with previous analyses.

Since our likelihood function differs from the widely used multivariate Gaussian, the goodness-of-fit evaluation must be adapted accordingly. For Gaussian likelihoods, the $\chi^2_{\text{best-fit}}$, estimated at the best-fit parameters, is to a very good approximation sampling an underlying χ^2_ν probability distribution, which depends only on the number of degrees of freedom ν – this is only an approximation however, because of informative priors, non-linear modelling and correlated error bars (see e.g. Joachimi et al. 2020). A good fit will have a $\chi^2_{\text{best-fit}}$ close to the maximum of the χ^2_ν probability distribution, while a bad fit will land far in the tail, leading to a probability to exceed (PTE) that is smaller than our acceptance threshold, set to 0.01. For our Student- t distribution likelihood, we still assess the goodness-of-fit with PTE values, however the χ^2_ν curve needs to be modified (see Appendix C for details).

A few differences exist between the KiDS-1000 and DES-Y1 likelihoods which are worth highlighting here, as these influence the construction of our joint pipeline. First, the original DES-Y1 peak count analysis samples σ_8 instead of S_8 ; the latter is a better option as it exactly covers the training volume and is therefore adopted for our DES-Y1 re-analysis. Second, the treatment of intrinsic alignments are simpler in the DES-Y1 analysis: the IA contribution is estimated from the alignments of dark matter haloes, which are assumed to fully correlate with the alignment of central galaxies. This non-linear prescription provides a single IA model that is then added to the predictions, without marginalisation. As discussed in Sec. 5.2, not marginalising over the IA in the KiDS analysis slightly underestimates the total error. This is

likely less important in the DES-Y1 likelihood since the statistical error is larger.

This also connects with the third difference, which is that in the baseline DES-Y1 measurement, only the auto-tomographic redshift bins are included, in an attempt to avoid possible residual contamination from unmodelled IA in the cross-redshift bins. This turns out to be an over-conservative data cut. Indeed, the recent **DES-Y1 persistent homology cosmic shear analysis** from Heydenreich et al. (2022) reveals that the constraints on S_8 are negligibly affected by these IA terms: they show that a full tomographic analysis including all cross-tomographic combinations shift the parameter by at most 0.3σ towards higher S_8 values, even when the inferred A_{IA} is as large as unity. Although their analysis is based on the different statistics (they use persistent homology instead of peak count), their results should hold here too, given that peaks are a subset of their data vectors. Therefore residual IA cannot play an important role in the DES-Y1 peak count analysis, justifying our choice to include the cross-redshift bins here (up-to-pairs, but not the triplets nor the quadruplets since these are not fully modelled yet for the DES-Y1 data).

A fourth difference in the likelihood concerns the treatment of the baryonic feedback: in HD21 the peak statistics are measured in the *Magneticum* simulations to ensure that the selected elements from the data vectors are immune to unmodelled baryonic mechanisms, but no marginalisation is included. This can potentially lead to a slightly over-optimistic precision on the DES likelihood compared to the KiDS-1000 likelihood, which includes marginalisation over the b_{bary} parameter. We therefore decided to include in our joint analysis the same marginalisation machinery for both the KiDS-1000 and the DES-Y1 pipelines. Moreover, we use a unique b_{bary} parameter to infuse baryonic feedback into both surveys, since these physical processes describe physics that affect the foreground matter distribution independently of survey-specific source selection. In total, the combined-survey analysis marginalises over nine redshift bias parameters, nine shear bias parameters, one IA and one baryon parameter. The sampling strategy of the DES-related parameters are listed in Table 3. Finally, given the absence of overlap between two survey footprints and the compatibility of the priors, the two likelihoods can be directly added at each evaluation point, without needing to consider cross-survey covariance.

4 SYSTEMATIC UNCERTAINTIES

As the amount of high quality lensing data keeps increasing, the statistical precision reaches unprecedented high levels, and consequently understanding and controlling the residual systematics in every segment of the data analyses has become one of the primary objectives and focus of development in the field of beyond-2pt statistics. We investigate here a number of such systematic effects that have been identified in the literature and mentioned earlier, including residual uncertainties related to interpolation in the modelling, shear calibration, photometric redshifts, astrophysics (intrinsic alignments and baryonic feedback), simulation-based covariance matrix, non-linear physics, source-lens coupling and likelihood sampling strategies. Some of the systematics that are ignored in the current work are those related to the effect of source blending, depth variations, PSF leakage, or the cosmology dependence of the IA signal. These will likely become important in the future, but can be safely omitted in current Stage-III lensing surveys (see HD21, Zürcher et al. 2022). Amongst those that we investigate here, many are shown to be sub-dominant or heavily suppressed by our range of S/N , while others are forward-modelled with nuisance parameters that are marginalised over in the likelihood analysis.




Figure 5. Accuracy of the KiDS-1000 GPR emulator, computed with a leave-one-out cross-validation test. The results are colour-coded with the S_8 value of the removed training point, and compared with the statistical precision on the measurement of the peak function (shown with the black dashed lines). The black solid line indicates the Λ CDM node, and the different panels show the auto- and cross-redshift (up to pairs) measurements; the other 15 tomographic combinations show a similar precision and hence are not shown. The outliers seen in a few panels are of extreme S_8 values and as such required the test emulator to extrapolate; this does not occur with the full emulator, and should therefore not be considered when estimating the interpolation error.

4.1 Modelling

The accuracy of the peak function modelling has two aspects to be considered: we must understand both how well the cosmology scaling is captured by the emulator, and whether any elements of our data vector are affected by either resolution limits of our simulations or the choice of gravity solver. Regarding the first aspect, the cosmo-SLICS have been shown to match in precision the **commonly used** CosmicEmu (Heitmann et al. 2014), and to even outperform the latter in terms of range, benefiting from more training nodes (Harnois-Déraps et al. 2019). The GPR interpolation uncertainty is **fully propagated in the likelihood**, and is quantified with a leave-one-out cross-validation test, in which the emulator is trained on all but one node, producing predictions at that node that are then compared with the measurement. As shown in Fig. 5, we cycle over all 26 nodes in this way and estimate an upper limit on the error (since the actual GPR has all the nodes). The accuracy degrades towards large positive or negative S/N values, but most of the cross-validation lines lie well within the statistical precision on the data, shown with the dashed black lines. There are a handful of exceptions with poorer accuracy, attributed to removing extreme values of S_8 from the training set and therefore effectively demanding the GPR to extrapolate. With these edge nodes included, the full emulator has no such outliers. In fact, as argued in HD21, the most reliable estimate of the emulator’s precision is evaluated by removing the fiducial cosmology and training on the others, which is shown as the thick black lines in the figure and always well within the statistical precision. For the KiDS modelling, the interpolation error is mostly at the 1% error over the range $-1.0 < S/N < 3.0$ (our ‘clean’ range), and is otherwise always under 10%. Similarly, the DES interpolation error is everywhere under 2% (see HD21), **the difference coming from the choice of S/N cuts**. This will likely be a limiting factor for future data analysis with sub-percent accuracy requirement on the modelling, and will likely be

addressed by increasing the number of nodes in the next generation of the *Cosmology Training Set*. At the moment however, the interpolation error is low enough for our analysis. We nevertheless include it in our error budget by averaging over the square of the residuals (after the outliers have been removed):

$$\text{Cov}_{\text{interp}} = \text{diag} \left\langle \left(N_{\text{peaks}}^{\text{GPR}} - N_{\text{peaks}}^{\text{sim}} \right)^2 \right\rangle, \quad (7)$$

and adding this to our statistical covariance. We could have instead used the errors directly provided by the Gaussian Process Emulator, however Heydenreich et al. (2021) have shown that the two methods yield posteriors with negligible differences. This contribution, although small, helps with the goodness-of-fit in the data analysis.

4.2 Shape calibration

Table 1 shows the average multiplicative correction factors m_a that must be applied to the observed galaxy shapes in order to correct for a combination of residual PSF leakage, blending and measurement noise, as assessed from Giblin et al. (2021). While in A21 the uncertainty on the shape calibration is absorbed directly in the analytical covariance matrix, our simulation-based method works instead at the level of the data vector, as for all other nuisance parameters. The impact of mis-calibrated shape measurements is estimated by infusing a non-corrected global term $m_a \rightarrow m_a + \Delta m_a$ directly in the simulations and measuring the effect on the different elements of the peak function N_{peaks}^k . As we show later, this systematic effect is completely subdominant compared to the others due to the tight priors on Δm_a (reported in Table 2), and hence it is sufficient to model its impact with a reduced accuracy. In HD21 the estimation is based on a linear regression (i.e. $\partial N_{\text{peaks}} / \partial \Delta m_a$ per data element) that is fit through 10 values of Δm_a . We use here only two points, at $\pm 1\sigma$, which is sufficient given the small values of Δm_a . The measured $\partial N_{\text{peaks}} / \partial \Delta m_a$ is further discussed in Appendix B, and is used to modify the data vector for any value of Δm_a (see Eq. 8) sampled in the likelihood. For cross-redshift tomographic bins, we use the mean shift, e.g. $\Delta m_a^{1/2} = (\Delta m_a^1 + \Delta m_a^2)/2$, which is consistent with what is currently done for all shear two-point function analyses. We could instead use an n_{gal} -weighted mean to compute the Δm_a shift in cross-redshift tomographic bin, however this should have a negligible effect given the tight priors on these parameters, and we therefore leave this for the future.

4.3 Photometric redshifts

The KiDS-1000 uncertainty on the redshift distributions has been fully quantified in Hildebrandt et al. (2021), where it is shown that the **mean of the $n(z)$ is captured to a high accuracy**, varying by no more than 0.014 at the 1σ level¹². As for the shear bias, the posteriors on the mean of the redshift distributions are used as priors on nuisance parameters in the current work, summarised in Table 2. In this case however, the five redshift bias parameters Δz_a must be drawn from a correlated distribution. This is achieved in a two-step operation where we first draw five uncorrelated numbers from the priors, then rotate into the correlated space using a Cholesky decomposition of the redshift covariance

¹² **The full shape of the $n(z)$ is less accurate than its mean, and which consequences we leave for future work.**

matrix:

$$C_z \times 10^5 = \begin{bmatrix} 11.20 & 2.600 & 1.562 & 0.056 & 0.622 \\ 2.600 & 12.78 & 4.081 & -1.692 & -0.2140 \\ 1.562 & 4.081 & 13.81 & -1.139 & 0.525 \\ 0.056 & -1.692 & -1.139 & 7.551 & 3.054 \\ 0.622 & -0.2140 & 0.525 & 3.054 & 9.496 \end{bmatrix},$$

which results in a correlated sampling of these five nuisance parameters (see A21, Hildebrandt et al. 2021, for more details). We produced a dedicated set of *Redshift Training Set* simulations in which the $n(z)$ are shifted, but which are otherwise identical to the *Cosmology Training Set* at the fiducial cosmology. Following HD21, we measure the peak function on full mock surveys with 10 shifts, each with a slightly different value of Δz_a sampled from the prior, then extract a linear fit per data element and estimate $\partial N_{\text{peaks}} / \partial \Delta z_a$. This derivative is used to forward model redshift uncertainties on our data vector for arbitrary Δz_a values. Again, we use the mean shift when considering cross-redshift bins, and the DES-Y1 $\partial N_{\text{peaks}} / \partial \Delta z_a$ measurements from HD21.

4.4 Astrophysics

Cosmic shear measurements are strongly affected by IA and baryon feedback. Using the *IA* and the *Baryons Training Sets* described in Sec. 2.3, we estimate in a similar way $\partial N_{\text{peaks}} / \partial A_{\text{IA}}$ and $\partial N_{\text{peaks}} / \partial b_{\text{bary}}$, where A_{IA} and b_{bary} are free parameters that control the levels of IA and baryon contaminations, respectively. The IA derivative is obtained by linear fitting the peak function's response to changes in A_{IA} , measured from the *IA Training Set* infused with $A_{\text{IA}} = 2.0, 1.0, 0.0, -1.0$ and -2.0 . Since IA is currently not well constrained and the NLA parameterisation is an effective model, we adopt a wide top-hat prior over the range $[-5.0; 5.0]$, as argued in Joachimi et al. (2020). **This extrapolates our fit to larger A_{IA} values, which can in principle become inaccurate, however in the end the 3σ region of our posterior is fully contained within the training range (see Sec. 5).** Similarly, the baryon derivative is measured from the *Baryons Training Set*, which we use to infuse a baryonic correction whose strength is controlled by the parameter b_{bary} . The case $b_{\text{bary}} = 0.0$ corresponds to a dark matter-only universe, while $b_{\text{bary}} = 1.0$ corresponds to the case where the feedback processes is exactly described by the *Magneticum* physics. There is a large uncertainty on the amplitude of this baryon correction, hence we scale the measured baryonic correction with a free parameter b_{bary} . Since the *Magneticum* suites are already a strong model (see Martinet et al. 2021, for a comparison with other hydrodynamical simulations), we sample the range $b_{\text{bary}} \in [0.0, 2.0]$, thereby spanning a variety of realistic models (albeit imposing a fixed shape for the relative signal). **As seen later, low b_{bary} values are not well constrained by the data while larger values are strongly disfavoured**, hence we do not extend the prior limit beyond 2.0.

4.5 Implementation of forward-modelled systematics

Four sources of systematics are forward-modelled in our pipeline. Following Heydenreich et al. (2022), we construct systematics-infused data vector as:

$$N_{\text{peaks}}^{\text{syst}}(\boldsymbol{\pi}, \Delta m_a, \Delta z_a, A_{\text{IA}}, b_{\text{bary}}) = N_{\text{peaks}}^{\text{GPR}}(\boldsymbol{\pi}) + \left[\partial N_{\text{peaks}} / \partial \Delta m_a \right] \Delta m_a + \left[\partial N_{\text{peaks}} / \partial \Delta z_a \right] \Delta z_a \dots + \left[\partial N_{\text{peaks}} / \partial A_{\text{IA}} \right] A_{\text{IA}} + \left[\partial N_{\text{peaks}} / \partial b_{\text{bary}} \right] b_{\text{bary}}, \quad (8)$$

where the **twelve** parameters $(\Delta m_a, \Delta z_a, A_{\text{IA}}, b_{\text{bary}})$ are sampled from the priors described in Table 2. We marginalise over these nuisance parameters when inferring the values of the cosmological parameters. Equation

(8) assumes that these different systematics are independent of cosmology and from each other, which we know is not entirely true. It has been shown that the cosmology dependence of the baryon feedback is a second order effect (McCarthy et al. 2017), supporting our simplified approach, however the intrinsic alignments couple to the tidal field that is in itself cosmology dependent. The shear calibration and redshift errors are independent of cosmology *a priori*, however the derivatives of the peak function with respect to Δm_a and Δz_a are not (see HD21), a secondary effect we neglect here. Moreover, it has been shown that the photometric and shape calibration errors are sometimes correlated (MacCrann et al. 2022). Although these approximation will become important in Stage-IV surveys, the current level of statistical precision allows us to relax the modelling of these effects without hurting our results. We illustrate this point in Section 5.2 by running inference MCMC chains in which the modelling of some or all of these systematic effects are switched off: the minor impact this has on the inference validates this approach. We also assume here that these systematic effects have a linear dependence on the nuisance parameter, which is probably not entirely true, but has been shown to be good enough for Stage-III lensing data in Heydenreich et al. (2022, see their figure 7).

4.6 Other sources of systematics

In addition to the main systematic effects described in the last section, we consider here other known sources of errors that could potentially impact our results.

N-body resolution

Being completely simulation-based, our analysis relies on the quality of the underlying training samples. As mentioned already in Sec. 2.3, the *Cosmology Training Set* has been shown to closely reproduce the non-linear clustering of the Cosmic Emulator (Heitmann et al. 2014), which is based on a completely independent *N-body* code. This agreement between different gravity solvers is key to assert the accuracy of the non-linear solution to structure formation (see, e.g. Euclid Collaboration: Knabenhans et al. 2019, for a comparison between different *N-body* solvers), and the convergence of the solution must be assessed via a comparison with calculations carried out with a higher force/mass resolution simulations. As shown in HD21, known limits in the mass resolution of the *cosmo-SLICS* used for the peak function emulation mainly affect high peaks. More precisely, $N_{\text{peaks}}^{\kappa}(S/N > 4.0)$ is systematically under-predicted by tens of percent, while the $S/N = 4.0$ count is affected by no more than 5%. This is in fact one of the main justification for our initial choice of upper S/N limit.

The KiDS-1000 data are deeper than DES-Y1, and hence the sensitivity to such non-linear effects could be accrued here. We verify this by running our cosmological inference on the peak count statistics measured from the SLICS-HR, in which the increased force resolution results in a slightly larger number of large positive and negative peaks. Details are presented in Appendix B, but in short our data selection and marginalisation scheme almost completely protects us against this, yielding no noticeable shifts on Ω_m nor S_8 . As in HD21, we nevertheless compute a multiplicative factor from the ratio between the SLICS-HR and the mean of the SLICS and apply it on our model predictions during the likelihood sampling. The overall effect is smaller than the baryon and IA corrections, hence marginalising over these latter two significantly washes out the impact of inaccurately-modelled non-linear physics under question here. In the future we intend to look into multi-fidelity emulators as in Ho et al. (2022).

Ray-tracing approximations

Our ray-tracing method in itself contains approximations and algorithmic components that are bound to affect to some level the lensing statistics. Most importantly, the finite thickness of the mass sheets and the randomisation process between them destroys correlations along the line of sight; in particular it can slice large galaxy clusters in two, and no structures larger than $257.5 h^{-1}\text{Mpc}$ can exist in any of our light-cones. This suppresses some of the large-scale power, as documented in Takahashi et al. (2017, see their Appendix B). However, smaller structures, such as those probed by the peak statistics, are left completely unaffected by this, which is why no forward modelling is needed here. This has been measured specifically for peak statistics in Zorrilla Matilla et al. (2020) where it was found that this plays a subdominant role even for Stage-IV surveys. Of course, full on-the-fly light-cones such as the ‘Onion Universe’ methods (Fosalba et al. 2008) avoid these problems, which we will consider for future analyses.

Another source of error comes from the fact that our simulations assume the Born approximation in the flat-sky limit, which introduces small inaccuracies at high- ℓ and low redshift, respectively (Hilbert et al. 2020). However, these are affecting the signal at a level much smaller than the statistical accuracy of our lensing data, and are not expected to matter here.

Covariance matrix

Estimation of the covariance matrix is one of the main computational challenges for non-Gaussian weak lensing probes, as it requires a large number of simulations with a resolution that is high enough to capture the non-linear physics being measured. Resorting to approximate methods such as FLASK (Xavier et al. 2016) and ICE-COLA (Izard et al. 2018) can significantly lower the computational cost of creating such mocks, but at the price of a reduced precision on the physics under investigation. We instead opted for mocks produced by a full N -body suite, our *Covariance Training Set*, and are therefore only limited by the number of mocks and their box size. To test the convergence of our covariance matrix with respect to N_{sims} , we run an inference analysis in which we increase the number of pseudo-independent realisations to 2120 (and adjusted the likelihood N_{sims} parameter accordingly), and find an excellent match to the posterior, with only the tail of the distribution being slightly modified. We could also have opted for a data compression such as in Zürcher et al. (2022) but that is not necessary given our results have converged, and our choice of likelihood accounts for the noise in the covariance matrix.

The simulation box size could also affect our results, however it has been shown in Harnois-Déraps et al. (2019) that the SLICS contains about 75% of the “super-sample covariance” term (SSC), when applied to 2pt statistics, yielding constraints on cosmological parameters that are highly accurate. Although this has not been demonstrated to date, peak count statistics are thought to be even less affected by the SSC, given that the covariance is close to being Poissonian, not Gaussian. In addition, as mentioned earlier, Burger et al. (2022) finds for the density-split statistics an excellent agreement between the SLICS covariance and that from full sky log-normal FLASK mocks (**which contain an incomplete contribution from the trispectrum term but the full SSC**), supporting our claim that the partly missing SSC must have a minimal influence on our error budget. This is also consistent with the recent findings from Linke et al. (2023) according to which the SSC term affects only the Fourier space estimators, whereas covariance matrix measured from intra-survey real-space statistics such as the M_{ap} are unbiased.

Source-lens coupling and blending

An important difference between real and mock galaxies is that those in the data are clustered, which leads to a number of effects that are systematically absent from the calibration sample. For example, the quality of the shape measurements is lowered in regions of high density due to blending and obscuration. More importantly, the uncertainty in photometric redshifts is particularly severe in such areas, which often results in cluster members being wrongly assigned a higher redshift. This subsequently creates a small population of apparently high-redshift outliers that carry an unexpectedly weak shear component, thus diluting the overall lensing signal. Correcting for this can be partially achieved with ‘boost factors’, however it was shown in HD21 and Zürcher et al. (2022) that even though the excess clustering around high peaks is indeed measured in the data, the impact this has on the inferred cosmology can be safely ignored. It was also shown in Gatti et al. (2023a) that source clustering had a minimal effect on the peak count statistics, supporting our choice to neglect this here.

Sampling the likelihood

Our likelihood sampling strategy, described in Sec. 3.3, assumes a flat prior for the four main cosmological parameters and the two astrophysical parameters (A_{IA} and b_{bary}), and Gaussian priors for the parameters associated with photometric redshifts and shape calibration. This is not strictly speaking a non-informative approach, **however the prior edges about the key measured parameters are sufficiently broad to have negligible impact on the posterior**. Since it is found in Lemos et al. (2022) that MULTINEST tends to yield slightly over-precise constraints, we use the NAUTILUS sampler for our fiducial results, but report both.

We also note that our cosmology sampling strategy is different from the other KiDS-1000 cosmic shear analyses, mainly due to the volume where our emulator is valid. For example, A21 sample uniformly the parameters S_8 , $\omega_c \equiv \Omega_c h^2$, $\omega_b \equiv \Omega_b h^2$, h and n_s . This choice is designed to avoid regions of parameter space that are strongly disfavoured by external data, and it was shown in Joachimi et al. (2020) that while it disfavoured high Ω_m values already in prior space, the resulting S_8 prior space is highly uninformative. We could have taken a similar approach, however our emulator is much quicker, and hence it is more natural to sample the full training space, ensuring a wide sampling of Ω_m , S_8 and w_0 .

Another aspect that currently limits our sampling strategy is the fact that we hold the value of many parameters fixed, notably Ω_b and n_s . In contrast, the DES-Y3 peak count analysis of Zürcher et al. (2022) use derivatives to marginalise over variation in these parameters, following the approach we adopt for IA and baryons. Neglecting to account for these has a small effect on current data sets (the DES-Y3 joint peaks+power spectrum analysis finds to be of about 0.13σ), which are thus ignored here.

$M_{\times}$ modes

The observed weak lensing signal can generally be decomposed into a combination of E - and B -modes, the latter of which can be estimated for any measurement by rotating all galaxies by 45 degrees; therefore, for the aperture mass map statistics, it is often referred to as $M_{\times}(\theta)$. The cosmic shear signal being a pure E -mode generator to first order, measurements of B -modes are therefore routinely used to assess the presence of residual systematics in lensing data (see, e.g. Zürcher et al. 2022, for a recent application to peak statistics). Whereas the two-point function B -mode signal is zero in absence of systematics, the construction of aperture mass maps on a grid inevitably injects non-zero $M_{\times}$ -modes due to the missing contribution from sub-pixel scales (Kilbinger et al. 2006). This can be important: for a small-angle cut-off scale of




Figure 6. Full inference analysis on the *Validation Set* (SLICS-HR) peak count data (blue) with MULTINEST, optionally infused with intrinsic alignments (red) or baryon feedback (grey). We marginalise over these two effects plus shape calibration and photometric uncertainty. The priors are shown by the dashed grey lines at the edge of the panels, while the cross-hairs show the input truths.

10 arcsec and an aperture of $\theta=2.0$ arcmin, B -modes measured this way can reach about 10 percent the size of the E -mode M_{ap}^2 signal. This effect is accentuated for larger cut-off scales and smaller opening angles θ . Given our pixel scale of 35 arcsec, we do expect non-zero $M_{\times}$ -modes to be introduced by our aperture map making, which we fully quantify in Appendix A. We show therein that the level of contamination is consistent with noise, that there is no evidence for residual systematics in the data from this measurement, hence that our cosmological analysis is clean of B -modes.

4.7 Peak count to cosmology pipeline validation

We test our KiDS-1000 cosmology inference pipeline by analysing simulated data of known cosmology, infused with a controlled amount of residual systematics. In order to avoid confirmation bias, these tests are carried out with the *Validation Set*, which have not been used in the cosmology training nor for the covariance estimation, with an N -body force resolution that is higher than the other simulations used in this work¹³. In addition, we use the forward-modelling approach presented in Sec. 4 to infuse the simulated data vectors with either intrinsic alignments (assuming $A_{\text{IA}} = 1.0$) or baryonic feedback (with $b_{\text{bary}} = 1.0$). Fig. 6 shows the results for these three analysis cases. The maxima of the projected posterior distributions are all centred on the input truth, except for the b_{bary} parameter, which are away from zero even in the no-baryon cases. This is a **projection effect similar to those discussed in Joachimi et al. (2020), Chintalapati et al. (2022) and Dark Energy Survey and Kilo-Degree Survey Collaboration (2023)**, and we have verified that reducing the lower prior limit to $b_{\text{bary}} = -2.0$ pushed both the red and blue maxima towards the ground truth.

¹³ We have further verified that the Λ CDM cosmology is correctly inferred when analysing data from the *Cosmology Training Set* but these tests are easier to satisfy since the data is used for training the emulator. We discuss these in greater details in Appendix B.




Figure 7. KiDS-1000 constraints on the two best-measured parameters from peaks count statistics, for different selection of redshift bins. Tomographic analyses all break the $S_8 - A_{\text{IA}}$ degeneracy.

In this test, there is a small secondary solution for $\Omega_m \sim 0.4$ that is slightly surprising. As detailed in Appendix A, this feature persists when analysing data from the *Cosmology Training Set* at the fiducial cosmology, but vanishes at other cosmologies. This is caused by the poor sensitivity of the current lensing data to Ω_m , as also seen in the large Ω_m scatter reported in A21 between different two-point functions. The posteriors on w_0 and b_{bary} are wide and significantly overlap with the prior limits, preventing us from achieving meaningful constraints on these parameters. We observe a degeneracy in the $[S_8 - w_0]$ plane here, however we show in Appendix B that it is not seen when analysing other cosmologies, making it impossible to draw physically meaningful conclusions about this. Only the $[S_8 - A_{\text{IA}}]$ plane is well constrained with the current KiDS-1000 peak count analysis: we achieve a 4.4% precision measurement on S_8 , with $S_8^{\text{SLICS-HR}} = 0.816^{+0.039}_{-0.033}$ (truth is 0.813), and a precision of $\sigma_{A_{\text{IA}}} = 0.45$ on A_{IA} , sampling the likelihood with MULTINEST.

The DES-Y1 pipeline validation is presented in HD21, while that for the joint KiDS-DES is presented in Appendix B, showing again an excellent agreement between the inferred cosmology and the input truth.

5 RESULTS: KIDS-1000

We present in this section the results from our cosmological inference analyses, beginning with the fiducial KiDS-1000 pipeline, then reporting on the importance of various selection cuts and systematic effects. For reasons explained in Sec. 4.7, we report only the constraints on S_8 and A_{IA} ; results are summarised in Table 4 and further condensed in Fig. 10.

From our fiducial full tomographic KiDS-1000 analysis of the measurements presented in Fig. 3, we obtain:

$$S_8^{\text{KiDS}} = 0.733^{+0.032}_{-0.032}, \quad A_{\text{IA}} = 0.71^{+0.49}_{-0.49}, \quad (9)$$

after marginalising over three cosmological parameters (Ω_m , w_0 and h) and 11 nuisance parameters ($5 \times \Delta m_a$, $5 \times \Delta z_a$, b_{bary}). Unless explicitly

mentioned, all quoted parameter constraints correspond to the **mean** $\pm 1\sigma$ region of the marginalised posterior, not to be confused with the point of maximum likelihood in the higher-dimensional space. This is therefore a 4.4% measurement of the structure growth parameter. The best-fit model is shown with the red line in Fig. 3. The joint constraints on these two parameters are shown in Fig. 7, along with results from different selections of tomographic bins. Importantly, the strong $S_8 - A_{IA}$ degeneracy seen in the no-tomographic case (the tilted dark purple contour) is lifted by tomographic decompositions, which capture the different redshift dependence of the cosmological and IA signals. This is an important verification of our IA modelling, which, under the NLA framework, compensates large S_8 values with large tidal alignments in the no-tomographic case only. Notably, all tomographic additions contribute to further tightening the constraints, once again demonstrating the power of using cross-redshift bins in non-Gaussian statistics. We also observe that all cases shown in Fig. 7 are consistent, providing statistical robustness to our measurement.

At our best-fit parameters the measurement yields a χ^2 of 250, which reduces to $\chi^2_{\text{red}} = 1.22$ after dividing by $\nu = (220 - 4.5) = 205.5$ degrees of freedom. Note that although we use six **unconstrained** parameters¹⁴ in our likelihood evaluation (the four cosmological parameters plus A_{IA} and b_{bary}), it was shown in Joachimi et al. (2020) that an effective number of $\nu = 4.5$ free parameters better describes the weak lensing data given the existing correlations and degeneracies, results which we have used here¹⁵. Our p -value for this measurement is 0.43, which is well above our threshold of 0.01. It is worth noting that the KiDS-1000 shear two-point correlation functions and band power analyses had a lower goodness-of-fit, with $p=0.034$ and 0.013, respectively.

In many previous analyses, sampling and marginalisation over w_0 is often excluded, being considered an extension to the vanilla Λ CDM scenario. In the present case, fixing w_0 to -1.0 when sampling the likelihood¹⁶ results in minor changes to the reported (S_8, A_{IA}) constraints, leading to $0.729^{+0.029}_{-0.029}$ and $0.73^{+0.48}_{-0.48}$. Interestingly, we find that the impact of opening up the w_0 dimension is far lower than for the two-point statistics, where Tröster et al. (2021) finds a degradation by a factor of a few on the S_8 constraints (compare their figures 1 and 6). Different degeneracy-breaking directions are likely causing this difference, which is promising for upcoming measurements of w_0 with alternative statistics (see Martinet et al. 2021, for a Stage-IV lensing forecast on the dark energy parameter with peak statistics).

One of the key questions to be explored by beyond-2pt statistics concerns the exact origin of the non-Gaussian cosmological information. Large peaks are often associated with massive galaxy clusters, which are known to be highly sensitive to the dark energy equation-of-state parameters for instance, however the wide projection effect and the fact that baryons, IA and non-linear physics maximally affect these large S/N peaks (Martinet et al. 2021; Harnois-Déraps et al. 2022) complicate the picture. To (partly) answer this question, we investigate the constraining power contained in the highest ($S/N > 3$) and lowest ($S/N < 0$) bins by removing these sequentially from the ‘aggressive’ data vector ($-2.5 \leq S/N \leq 4.0$) in the likelihood. The results are shown in Fig. 8, where we observe that the negative S/N peaks significantly help break the $[S_8 - A_{IA}]$ degeneracy, while the highest peaks help in tightening the S_8 constraints. In an analysis that ignored the role




Figure 8. Effect of S/N cuts on the KiDS-1000 constraints. Large peaks ($S/N > 3$) **slightly increase the statistical precision on S_8 , as seen by comparing the red and blue contours (see values in Table 4).** Negative peaks ($S/N < -1$), included in the red but excluded from the grey contours, help in breaking the $[S_8 - A_{IA}]$ degeneracy. The grey and blue contours correspond to the ‘clean’ and ‘aggressive’ cases, respectively.

of IA, M20 found that the amount of information about S_8 that is contained in negative peaks is quite small, however here we find that they actually play a key role once IA are forward-modelled.

5.1 Internal consistency

It has been found in previous cosmic shear analyses (e.g. A21, Hamana et al. 2020; Amon et al. 2022) that internal consistency tests can help differentiate residual systematics from statistical fluctuations. We therefore stress-test our results by removing data from tomographic bins one at a time before proceeding to the inference. For example, we consider results obtained from an analysis where exactly no data from bin1 (i.e. 1, 1U2...1U2, 1U2U3 ... 1U2U3U4U5) is used, then no data from bin2, and so on. The results are shown in Fig. 9, where we observe that all cases are self-consistent, in agreement with the full selection. Note that the S_8 shifts per-bin are not expected to match exactly those measured with other lensing probes due to different responses of the cosmic shear estimators to noise in the data. For example, A21 found that removing the fifth tomographic bin maximally degrades the precision on S_8 , confirming the large amount of information on this parameter carried by high redshift bins in shear two-point functions. In contrast, we find here that removing the third redshift bin has the worst impact on the precision. The third bin has the greatest number density of galaxies, hence better captures the information in peak statistics, whose mean value is affected by the noise level. The constraints on A_{IA} fluctuate about the fiducial results by less than 2σ , while those on S_8 agree within 1σ , as expected.

5.2 Impact of systematics

We present in this section additional variations with respect to the fiducial analysis, designed to better understand our results and assess their robustness to residual systematics. We first investigate the impact of

¹⁴ We do not count as free parameters those nuisance parameters for which we impose a tight prior.

¹⁵ It is not guaranteed that the exact same effective number of degrees of freedom applies here, given that the likelihood is not sampled over the same volume. We have checked that our goodness-of-fit is robust over choices for this quantity, with p -value varying between 0.48 and 0.37 over the range $2 < \nu < 7$.

¹⁶ This still uses the same w CDM emulator, but only varying the other three cosmological parameters.

Table 4. Summary of our cosmological inference analyses. Posteriors on Ω_m , h and w_0 are prior-limited, so their constraints are not reported here. Unless explicitly specified in the first column, the KiDS-1000 measurements are based on the ‘clean’ data vector, i.e. $-1.0 < S/N \leq 3.0$. Validation of the inference pipelines on mock data are presented in Appendix B.

Peak count analysis of KiDS-1000				
	NAUTILUS		MULTINEST	
	S_8	A_{IA}	S_8	A_{IA}
Fiducial	$0.733^{+0.032}_{-0.032}$	$0.71^{+0.49}_{-0.49}$	$0.733^{+0.021}_{-0.027}$	$0.74^{+0.43}_{-0.43}$
Λ CDM	$0.733^{+0.029}_{-0.029}$	$0.73^{+0.48}_{-0.48}$	$0.729^{+0.026}_{-0.026}$	$0.73^{+0.43}_{-0.43}$
Auto-only	$0.734^{+0.050}_{-0.050}$	$-0.2^{+1.0}_{-1.0}$	$0.732^{+0.039}_{-0.048}$	$-0.2^{+1.0}_{-1.0}$
Up to pairs	$0.750^{+0.036}_{-0.036}$	$0.10^{+0.69}_{-0.69}$	$0.742^{+0.032}_{-0.032}$	$0.11^{+0.63}_{-0.63}$
Up to triplets	$0.740^{+0.035}_{-0.035}$	$0.35^{+0.53}_{-0.53}$	$0.740^{+0.029}_{-0.029}$	$0.37^{+0.48}_{-0.48}$
No tomo	$0.695^{+0.033}_{-0.087}$	$-0.6^{+1.7}_{-2.2}$	$0.690^{+0.038}_{-0.068}$	$-0.7^{+1.5}_{-2.0}$
$-2.5 < S/N \leq 4.0$	$0.720^{+0.036}_{-0.026}$	$0.08^{+0.30}_{-0.30}$	$0.717^{+0.027}_{-0.022}$	$0.07^{+0.27}_{-0.27}$
$-2.5 < S/N \leq 3.0$	$0.717^{+0.031}_{-0.031}$	$0.14^{+0.31}_{-0.31}$	$0.713^{+0.023}_{-0.023}$	$0.22^{+0.27}_{-0.27}$
$0.0 < S/N \leq 4.0$	$0.739^{+0.031}_{-0.026}$	$0.77^{+0.47}_{-0.47}$	$0.744^{+0.023}_{-0.023}$	$0.80^{+0.38}_{-0.38}$
No IA	$0.726^{+0.024}_{-0.042}$	–	$0.720^{+0.021}_{-0.031}$	–
No baryons	$0.732^{+0.032}_{-0.032}$	$0.71^{+0.49}_{-0.49}$	$0.725^{+0.022}_{-0.027}$	$0.70^{+0.43}_{-0.43}$
No syst	$0.729^{+0.024}_{-0.056}$	–	$0.723^{+0.022}_{-0.048}$	–
No bin1	$0.734^{+0.043}_{-0.037}$	$0.10^{+0.74}_{-0.74}$	$0.735^{+0.035}_{-0.035}$	$0.10^{+0.70}_{-0.10}$
No bin2	$0.740^{+0.042}_{-0.048}$	$-0.78^{+0.72}_{-0.72}$	$0.740^{+0.040}_{-0.040}$	$-0.76^{+0.70}_{-0.70}$
No bin3	$0.775^{+0.049}_{-0.055}$	$0.97^{+0.63}_{-0.63}$	$0.777^{+0.048}_{-0.048}$	$0.95^{+0.60}_{-0.60}$
No bin4	$0.701^{+0.037}_{-0.037}$	$0.41^{+0.68}_{-0.68}$	$0.702^{+0.029}_{-0.034}$	$0.45^{+0.59}_{-0.59}$
No bin5	$0.720^{+0.036}_{-0.028}$	$0.53^{+0.61}_{-0.61}$	$0.723^{+0.030}_{-0.025}$	$0.53^{+0.57}_{-0.57}$
Peak count analysis of DES-Y1				
DH21	–	–	$0.737^{+0.027}_{-0.031}$	–
This work	$0.743^{+0.036}_{-0.036}$	–	$0.742^{+0.030}_{-0.034}$	–
Joint peak count analysis				
	NAUTILUS		MULTINEST	
	S_8	A_{IA}	S_8	A_{IA}
Fiducial	$0.733^{+0.020}_{-0.020}$	$0.82^{+0.47}_{-0.47}$	$0.732^{+0.012}_{-0.010}$	$0.82^{+0.33}_{-0.33}$
Λ CDM	$0.736^{+0.016}_{-0.018}$	$0.81^{+0.46}_{-0.46}$	$0.736^{+0.012}_{-0.015}$	$0.79^{+0.40}_{-0.40}$
No baryons	$0.728^{+0.020}_{-0.016}$	$0.82^{+0.46}_{-0.46}$	$0.725^{+0.018}_{-0.014}$	$0.83^{+0.39}_{-0.39}$
No IA	$0.726^{+0.020}_{-0.016}$	–	$0.729^{+0.015}_{-0.015}$	–

IA on the uncertainty by fixing A_{IA} to 0.72, the best-fit value in the fiducial analysis. Doing so, the error bars on S_8 shrink by less than 10%, while the mean value is not affected, by construction. Setting instead the IA parameter to 0.0, we can estimate the bias on the inferred cosmology if **IA are completely neglected**. We measure in this case $S_8^{\text{no-IA}} = 0.725^{+0.024}_{-0.042}$, a 0.22σ shift from the fiducial results. Intrinsic alignments are therefore a **modest** part of the error budget, suggesting that peak count analyses where IA are not modelled or held fixed (e.g. M18, HD21, Marques et al. 2023) likely yield both biased low and slightly optimistic constraints for S_8 .

We next carry out a similar study this time removing the modelling of baryons, fixing the associated nuisance parameter to $b_{\text{bary}} = 0.0$. As reported in Table 4, the measurements are mostly unchanged. As shown in M21, any non-zero residual feedback tends to lower the number of high S/N peaks in all tomographic bins, which, when confronted to fixed data, must be compensated with an increased value of inferred S_8 . Therefore, removing the baryon modelling goes the other way and reduces the inferred S_8 . This is not clearly seen with the NAUTILUS chains, but the MULTINEST runs shows this shift with 0.2σ significance.

Then, removing modelling of all systematics (photo- z , shape calibration, IA and baryons) results in S_8 values half way between the




Figure 9. Internal consistency: effect of removing tomographic data from the KiDS-1000 analysis.




Figure 10. Summary of S_8 constraints from this work, from recent cosmic shear data analyses and from *Planck*. This figure shows the projected 1σ errors.

no-baryon and no-IA cases, but the error bars are the **larger**. This suggests that marginalisation over these systematics helps in finding the true maximal likelihood. This is supported by the fact that the error on S_8 becomes smaller than the fiducial case if A_{IA} and b_{bary} are fixed to their best-fit value (0.72 and 0.5, respectively) instead of zero, **leading to $S_8^{\text{syst-fixed}} = 0.728 \pm 0.030$, a slight reduction in error compared to the fiducial error.**

Finally, we compared our fiducial results with those from the MULTINEST nested sampler and recover negligible biases in the inferred parameters, but with smaller error bars ($S_8 = 0.732 \pm 0.032$ versus

0.730 ± 0.026 for MULTINEST). This is consistent with previous findings (Lemos et al. 2022) and justifies our choice of NAUTILUS as our main sampler. We nevertheless report MULTINEST results as well for fairer comparison with previous results.

5.3 Comparison with previous KiDS-1000 results

The S_8 measurement presented here is not the first carried out from KiDS-1000. Previous analyses include the measurements of A21 and van den Busch et al. (2022), the latter of which used an upgraded photometric calibration compared to the former, **followed by that of Li et al. (2023a) based on upgraded shear measurements**. Loureiro et al. (2022) carried out a *pseudo- C_ℓ* analysis, Fluri et al. (2022) used instead a convolutional neural network, while Longley et al. (2023) re-analysed the data within the LSST-DESC pipeline. We report these results as the purple symbols in Fig. 10, where we see that they all seem to prefer slightly higher values of S_8 compared to our own measurements, albeit not by a significant amount. Given the important differences in the analysis pipelines between these efforts, it is reassuring to recover $< 1\sigma$ agreements. The constraints from the w CDM **band power analysis from Tröster et al. (2021) are reported in Fig. 11 and are broadly consistent with our peak statistics constraints, even though peaks are clearly more constraining on S_8 (0.732 ± 0.032 for peaks vs 0.742 ± 0.047 for band power), due to the reduced degeneracy in the $[S_8 - w_0]$ plane.** It is worth noting that both statistics provide similar constraints on the A_{IA} parameter ($\sigma_{A_{IA}} = 0.42$ for peaks, compared to $\sigma_{A_{IA}} = 0.36$ for band power), which is reassuring given that both use the same NLA approach. This error is significantly reduced ($\sigma_{A_{IA}} = 0.30$) when considering the more aggressive data selection ($-2.5 \leq S/N \leq 4.0$), but since the associated goodness-of-fit is poor, the results are not straight-forward to interpret. We nevertheless expect tighter constraints on A_{IA} to be achievable coming from non-Gaussian probes.

6 JOINT ANALYSIS WITH DES-Y1

The posterior obtained from the KiDS-1000 peak count analysis is fully consistent with that from the peak count analysis of the Dark Energy Year 1 (DES-Y1) presented in HD21. In particular, the latter finds $S_8^{\text{HD21}} = 0.737^{+0.027}_{-0.031}$, which significantly overlaps with our $S_8^{\text{KiDS}} 1\sigma$ results. Other parameters less well measured such as Ω_m and w_0 are also largely overlapping at the 1σ level (see the lower part of Fig. 11), which means the intersection between the two likelihood hyper-volumes must be large enough to safely combine the two data sets. Furthermore, both measurements are based on the similar analysis pipeline and, in particular, exploit the same simulations to model the cosmology dependence, thereby suppressing the risk of mis-interpreting the joint data due to non-uniform modelling of the signal.

6.1 Results: DES-Y1 re-analysis

As detailed in Sec. 3.3, there are differences between our DES-Y1 pipeline and that presented in HD21, including the S_8 sampling, the treatment of baryons, **the inclusion of the emulator uncertainty on the covariance** and the choice of sampler. The results from these re-analyses are presented in the lower panel of Fig. 11 (in grey and blue). The difference induced on these contours are small, but the goodness-of-fit improvement is important, with a p -value of 0.53 (using the same p -value estimator as HD21, we obtain 0.25, which is still a massive improvement compared to their p -value=0.005).

We remark that our joint pipeline contains a slight inconsistency: we include IA with the NLA model in the KiDS-1000 data (with




Figure 11. Comparison with the previous cosmic shear results. The top part shows a comparison with the KiDS-1000 band power analysis from Tröster et al. (2021, based on MULTINEST), while the bottom part presents the results from the peak count analysis of the DES-Y1 from HD21 (blue), along with the DES-Y1 re-analysis presented in this work (red) and detailed in Sec. 6. Note that the posteriors obtained from the NAUTILUS sampler are typically noisier than those from MULTINEST. The excellent agreement seen in this figure warrants the joint survey analysis.

marginalisation over A_{IA}) and with the non-linear halo-based IA model for the DES-Y1 data (without marginalisation, but with an on/off switch instead). **We verify the impact of this feature by analysing the likelihood with the DES IA model turned on and off and report on the difference, which is sub-dominant ($\Delta S_8 = 0.002$).** We also compare the results from turning off the modelling of baryons, and from replacing the w CDM by a Λ CDM analysis, finding in all cases results consistent with the fiducial analysis. The re-analysis presented in this work has slightly larger error bars compared to that of HD21, due to the marginalisation over baryons, and additionally NAUTILUS yields con-




Figure 12. Joint peak count analysis of the KiDS-1000 and DES-Y1 data. In the upper panel, the green bands indicate the 1σ and 2σ confidence intervals from the DES-Y1 re-analysis presented in this paper. The A_{IA} parameter shown here describes only the IA signal in the KiDS likelihood, since the IA is fixed to a non-linear model in the DES likelihood (see main text for details).

straints slightly larger compared to MULTINEST, as summarised in Table 4.

6.2 Results: joint KiDS + DES

We present in this section the results from our joint KiDS-1000 + DES-Y1 peak count analysis. Sampling the joint likelihood with our fiducial setup, we achieve improved constraints on S_8 with:

$$S_8^{\text{joint}, \Lambda\text{CDM}} = 0.732^{+0.020}_{-0.020} \quad (10)$$

and

$$S_8^{\text{joint}, \Lambda\text{CDM}} = 0.735^{+0.016}_{-0.018} \quad (11)$$

the tightest results obtained from non-Gaussian cosmic shear statistics to date, comparable to the recent joint ΛCDM analysis of the KiDS-1000 and DES-Y3 data (Dark Energy Survey and Kilo-Degree Survey Collaboration 2023), which measured $S_8^{\text{NLA}} = 0.792^{+0.016}_{-0.013}$. The two-dimensional posterior is shown in Fig. 12 (in blue) and compared to the fiducial KiDS-1000 (red) and DES-Y1 (green) peak statistics constraints. Recall that the A_{IA} parameter affects only the KiDS likelihood since, as explained in the previous section, the DES likelihood assumes instead a fixed halo-based IA model with no free parameter. We should therefore use caution when interpreting this parameter. The reported value is close to the point of maximum likelihood ($S_8^{\text{ML}} = 0.728$), and the size of the error bars on S_8 is consistent with our expectation: for example, we read from Table 4 that the ΛCDM KiDS-1000 analysis has a mean error of $\sigma_{S_8} = 0.029$. Scaling this precision by the square root of the area, we naively predict a joint survey error of around 0.018, and obtain 0.017. The error would be slightly larger had we included as well a marginalisation of the IA in the DES-Y1 part of this analysis, possibly explaining this slight difference. At the joint best-fit cosmology, the p -values for the KiDS and DES pipelines are basically unchanged, while the joint analysis has a $\chi^2_{\text{red}} = 1.15$ and a p -value of 0.96, all satisfying our goodness-of-fits criteria.

If we restrict the joint analysis to $w_0 = -1.0$, the S_8 values are minimally affected while the uncertainty is reduced, as expected from lowering the dimensionality of the likelihood. Alternatively, turning on the IA modelling in the DES likelihood only yields a 0.2σ downward shift, also expected whenever IA modelling is added. The smallness of this shift is once again showing that the intrinsic alignment do not significantly impact the peak count statistics as measured in the DES-Y1 data. In comparison, setting to zero the IA model in both KiDS and DES results in $S_8^{\text{joint, no-IA}} = 0.725^{+0.020}_{-0.016}$. Holding fixed the baryonic feedback parameter to $b_{\text{bary}} = 0.0$ has similar consequences on this joint analysis, shifting the best fit value to $S_8^{\text{joint, no-bary}} = 0.727^{+0.020}_{-0.016}$, a 0.25σ shift compared to the fiducial case. All these values are summarised in Table 4 and in Fig. 10 (with the brown symbols).

The dark energy equation-of-state is constrained from this joint analysis, with

$$w_0^{\text{joint}} = -1.12^{+0.42}_{-0.31}, \quad (12)$$

which is the first measurement of this quantity from peak statistics, and arguably one of the best from cosmic shear-only data analyses. The upper limit is close to the prior edge on w_0 that might lead to a slight underestimation of the error on this side. However, this measurement is robust against the choice of sampler ($w_0 = -1.09^{+0.29}_{-0.29}$ for MULTINEST), baryon modelling ($w_0 = -1.05^{+0.51}_{-0.22}$ setting $b_{\text{bary}} = 0.0$), IA ($w_0 = -1.13^{+0.44}_{-0.33}$ setting $A_{\text{IA}} = 0.0$) and scale cuts ($w_0 = -0.958^{+0.45}_{-0.093}$ when including the aggressive S/N cut in the KiDS-1000 data vector). As shown in Martinet et al. (2020), aperture-mass maps statistics are highly sensitive to dark energy and these results seem to be showing exactly that. Previously, the shear two-point function measurement from Troxel et al. (2018) on DES-Y1 achieved $w_0 = -0.77^{+0.30}_{-0.37}$ when varying the baryonic feedback model, using the MULTINEST sampler. The GCNN analysis of Fluri et al. (2022) was also able to set constraints on dark energy, with $w_0 = -0.93^{+0.32}_{-0.29}$, although they recognise that their results are affected by the prior boundary on the low side, **just like ours is on the high side**¹⁷. Similarly, HD21 found $w_0 > -1.5$, also prior-dominated on one side. Other cosmic shear measurements of w_0 involve additional data

¹⁷ We review the A21 definition that constraints are uninformed by the prior when the posterior drops below 0.135 of its maximum at the edges of a uniform prior volume. In the case of w_0 we find that the posterior is slightly above this threshold (0.156) at the upper edge. As this is at the borderline of the A21 criteria, we therefore caution that the error on this side might be slightly under-estimated.

(Abbott et al. 2023; Tröster et al. 2021), making this an unfair comparison.

6.3 Tension with *Planck*

The S_8 tension between recent CMB anisotropy and weak lensing data analyses is drawing a lot of attention, as it could point towards new physics or hidden systematics (see, e.g. Abdalla et al. 2022a, for a review). The *Planck* mission reports $S_8^{Planck} = 0.830 \pm 0.013$ (Planck Collaboration et al. 2020), **which is higher than many lensing results (see Amon et al. 2023, and references therein)**. The tension τ can be evaluated with a number of metrics, and we use here a relatively simple one used in A21, which compares the difference in the mean with the combined variances, $\text{var}[S_8]^{Planck}, \text{var}[S_8]^{peaks}$:

$$\tau = \frac{S_8^{Planck} - S_8^{peaks}}{\sqrt{\text{var}[S_8]^{Planck} + \text{var}[S_8]^{peaks}}} \quad (13)$$

With the fiducial setup shown in Table 4, and using the above definition, our results from the KiDS-1000 peak count analysis are in $\tau = 2.8\sigma$ with the *Planck* nominal constraints on S_8 . A21 finds a similar tension with their COSEBIs analysis, both using this simple tension metric and using a more sophisticated complementary method based on the full shape of the likelihoods. Similarly, we evaluate our joint KiDS-DES analysis to be in $\tau = 4.1\sigma$ tension with *Planck*, an increase that is driven by the decrease in error bars. Note that this tension is not only seen in weak lensing, but also in other late-time probes, e.g. data involving galaxy clustering, as recently summarised in Alonso et al. (2023, see their figure 7) and reviewed with greater details in Abdalla et al. (2022b). The current work exacerbates the existing trend, without providing an obvious solution. Again, large unaccounted contributions from IA and baryons could push the inferred S_8 value towards *Planck*, **but our analysis prefers lower values: in particular, we measure $b_{\text{bary}} < 1.05$ at 95% CL in the KiDS-1000 analysis, and < 0.82 in the joint analysis, excluding baryonic feedback models that are stronger than the *Magneticum*. Note that the redshift estimation methods used to analyse the DES-Y1 data are suboptimal compared to recent developments (see e.g. Hildebrandt et al. 2021), potentially causing biases of up to 0.03 in the inferred S_8 value. The observed tension with *Planck* would be reduced if that bias was real.**

7 CONCLUSIONS

We report in this paper a 4.4% measurement of S_8 from the tomographic peak count statistics measured from the KiDS-1000 data. Our simulation-based inference method exploits the non-linear features extracted from aperture mass statistics, sensitive to scales as small as 2.0 arcmin. We model the cosmological dependence with simulated w CDM weak lensing light-cones, we estimate the covariance matrix numerically, and forward model the effect of intrinsic alignments, baryonic feedback, photometric redshift error and galaxy shape mis-calibrations. We find a value of $S_8^{KiDS} = 0.732^{+0.032}_{-0.032}$, which aligns well with previous KiDS-1000 measurements. We show that our results are robust to residual systematics and that, of these, intrinsic alignment of galaxies plays the most important role, shifting the best-fit S_8 value by 0.22σ if left unmodelled.

The inferred posterior distribution is consistent with the peak count measurement carried out on the DES-Y1 data using a similar analysis pipeline (HD21), allowing us to jointly analyse the two data sets, which yields $S_8^{\text{joint}} = 0.732^{+0.020}_{-0.020}$, one of the tightest constraints on this parameter from lensing data alone. The combined data sets have enough statistical precision to allow the first measurement of the

dark energy equation-of-state parameter from non-Gaussian statistics: $w_0^{\text{joint}} = -1.12^{+0.42}_{-0.31}$, in agreement with the Λ CDM scenario, and robust to variations in the analysis choices.

Our best-fit S_8^{joint} is also in statistical agreement with all previous KiDS-1000 analyses and with the HSC-Y3 and DES-Y3 γ -2PCF results, but lower than the DES-Y3 measurements from peaks and moments, and in 4σ tension with *Planck*. There is no obvious solution to relax this tension, which therefore must either be explained from new physics or unaccounted systematics either in the CMB or lensing data, or in their treatment. Our pipeline has been thoroughly tested, however we recognise it is incomplete. As detailed in Sec. 4, we hold fixed a number of cosmological parameters, which likely affect our results, including Ω_b, n_s and m_ν . We also consider a single baryonic feedback model (although we allow its amplitude to vary), knowing that other hydrodynamical simulations would provide slightly different responses. Furthermore, we model IA with the redshift-independent NLA model, which we know is an incomplete effective model, and we neglect source clustering, as it was shown to be completely subdominant. Addressing these will therefore be the object of future work. It will also be informative to compare our results to other non-Gaussian probes of cosmic shear, and possibly combine the methods to further reduce the uncertainty on S_8 and w_0 .

ACKNOWLEDGEMENTS

We would like to thank Dr. Eric Tittley from the Royal Observatory of Edinburgh for his continuous help and support with the Cuillin cluster on which the bulk of the computations were carried on, as well as Elena Sellentin for discussions on the use of Hotelling's T^2 statistics.

JHD acknowledges support from an STFC Ernest Rutherford Fellowship (project reference ST/S004858/1). SH is supported by the U.D Department of Energy, Office of Science, Office of High Energy Physics under Award Number DE-SC0019301. BG acknowledges support from a UK STFC Consolidated Grant and from the Royal Society of Edinburgh through the Saltire Early Career Fellowship (ref. number 1914). TT acknowledges funding from the Swiss National Science Foundation under the Ambizione project PZ00P2_193352. KD is supported by the grant agreements ANR-21-CE31-0019 / 490702358 from the French Agence Nationale de la Recherche / Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for the LOCALIZATION project. TC is supported by the INFN INDARK PD51 grant, by the Fondazione ICSC National Recovery and Resilience Plan (PNRR) Project ID CN-00000013 “Italian Research Center on High-Performance Computing, Big Data and Quantum Computing” funded by MUR Missione 4 Componente 2 Investimento 1.4: “Potenziamento strutture di ricerca e creazione di “campioni nazionali” di R&S (M4C2-19)” – Next Generation EU (NGEU), and by the FARE MIUR grant ‘ClustersXEuclid’. HH is supported by a DFG Heisenberg grant (Hi 1495/5-1), the DFG Collaborative Research Center SFB1491, as well as an ERC Consolidator Grant (No. 770935). BJ acknowledges support by STFC Consolidated Grant ST/V000780/1. AHW is supported by the Deutsches Zentrum für Luft- und Raumfahrt (DLR), made possible by the Bundesministerium für Wirtschaft und Klimaschutz, and acknowledges funding from the German Science Foundation DFG, via the Collaborative Research Center SFB1491 “Cosmic Interacting Matters - From Source to Signal”.

The KiDS-1000 results in this paper are based on data products from observations made with ESO Telescopes at the La Silla Paranal Observatory under programme IDs 177.A-3016, 177.A-3017 and 177.A-3018, and on data products produced by Target/OmegaCEN, INAF-OACN, INAF-OAPD and the KiDS production team, on behalf of the KiDS consortium. This project also used public archival data

from the Dark Energy Survey (DES). Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministério da Ciência, Tecnologia e Inovação, the Deutsche Forschungsgemeinschaft, and the Collaborating Institutions in the Dark Energy Survey. The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zürich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l'Espai (IEEC/CSIC), the Institut de Física d'Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, the National Optical Astronomy Observatory, the University of Nottingham, The Ohio State University, the OzDES Membership Consortium, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, and Texas A&M University. Based in part on observations at Cerro Tololo Inter-American Observatory, National Optical Astronomy Observatory, which is operated by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

The chains presented in this paper are plotted with GETDIST (Lewis 2019), available at <https://getdist.readthedocs.io>.

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

DATA AVAILABILITY

The SLICS numerical simulations can be found at <http://slics.roe.ac.uk/>, while the SLICS-HR, the cosmo-SLICS and the *Magneticum* can be made available upon request. This work also uses public KiDS-1000 and DES-Y1 data, which can be found at <https://kids.strw.leidenuniv.nl/DR4/index.php> and <https://des.ncsa.illinois.edu/releases/y1a1>, respectively.

REFERENCES

- Abbott, T. M. C. et al. 2018, *ApJS*, 239, 18, 1801.03181
 —. 2022, *Phys. Rev. D*, 105, 023520, 2105.13549
 —. 2023, *Phys. Rev. D*, 107, 083504, 2207.05766
 Abdalla, E. et al. 2022a, *Journal of High Energy Astrophysics*, 34, 49, 2203.06142
 —. 2022b, *Journal of High Energy Astrophysics*, 34, 49, 2203.06142
 Ajani, V., Peel, A., Pettorino, V., Starck, J.-L., Li, Z., & Liu, J. 2020, *Phys. Rev. D*, 102, 103531, 2001.10993
 Akesson, R. et al. 2019, arXiv e-prints, arXiv:1902.05569, 1902.05569
 Alarcon, A. et al. 2020, arXiv e-prints, arXiv:2007.11132, 2007.11132
 Alonso, D., Fabbian, G., Storey-Fisher, K., Eilers, A.-C., García-García, C., Hogg, D. W., & Rix, H.-W. 2023, arXiv e-prints, arXiv:2306.17748, 2306.17748
 Alonso, D., Sanchez, J., Slosar, A., & LSST Dark Energy Science Collaboration. 2019, *MNRAS*, 484, 4127, 1809.09603
 Amon, A. et al. 2022, *Phys. Rev. D*, 105, 023514, 2105.13543
 —. 2023, *MNRAS*, 518, 477, 2202.07440
 Anbajagane, D. et al. 2023, arXiv e-prints, arXiv:2308.03863, 2308.03863
 Angulo, R. E., & Hahn, O. 2022, *Living Reviews in Computational Astrophysics*, 8, 1, 2112.05165
 Asgari, M. et al. 2021, *A&A*, 645, A104, 2007.15633
 —. 2020, *A&A*, 634, A127, 1910.05336
 Begeman, K., Belikov, A. N., Boxhoorn, D. R., & Valentijn, E. A. 2013, *Experimental Astronomy*, 35, 1, 1208.0447
 Benítez, N. 2000, *ApJ*, 536, 571
 Bilicki, M. et al. 2021, *A&A*, 653, A82, 2101.06010
 Blazek, J. A., MacCrann, N., Troxel, M. A., & Fang, X. 2019, *Phys. Rev. D*, 100, 103506, 1708.09247
 Bridle, S., & King, L. 2007, *New Journal of Physics*, 9, 444, 0705.0166
 Brouwer, M. M. et al. 2018, ArXiv e-prints, 1805.00562
 Burger, P. A. et al. 2022, arXiv e-prints, arXiv:2208.02171, 2208.02171
 —. 2023, arXiv e-prints, arXiv:2309.08602, 2309.08602
 Castro, T., Quartin, M., Giocoli, C., Borgani, S., & Dolag, K. 2018, *MNRAS*, 478, 1305, 1711.10017
 Chintalapati, P. R. V., Gutierrez, G., & Wang, M. H. L. S. 2022, *Phys. Rev. D*, 105, 043515, 2108.11518
 Dalal, R. et al. 2023, arXiv e-prints, arXiv:2304.00701, 2304.00701
 Dark Energy Survey and Kilo-Degree Survey Collaboration. 2023, *The Open Journal of Astrophysics*, 6, 36, 2305.17173
 Duncan, C. A. J., Harnois-Déraps, J., Miller, L., & Langedijk, A. 2022, *MNRAS*, 515, 1130, 2111.09867
 Edge, A., Sutherland, W., Kuijken, K., Driver, S., McMahon, R., Eales, S., & Emerson, J. P. 2013, *The Messenger*, 154, 32
 Erben, T. et al. 2013, *MNRAS*, 433, 2545, 1210.8156
 Euclid Collaboration: Ajani, V. et al. 2023, *A&A*, 675, A120, 2301.12890
 Euclid Collaboration: Knabenhans, M. et al. 2019, *MNRAS*, 484, 5509, 1809.04695
 Fenech Conti, I., Herbonnet, R., Hoekstra, H., Merten, J., Miller, L., & Viola, M. 2016, ArXiv e-prints, 1606.05337
 Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, 398, 1601, 0809.3437
 Fluri, J., Kacprzak, T., Lucchi, A., Refregier, A., Amara, A., Hofmann, T., & Schneider, A. 2019, *Phys. Rev. D*, 100, 063514, 1906.03156
 Fluri, J., Kacprzak, T., Lucchi, A., Schneider, A., Refregier, A., & Hofmann, T. 2022, *Phys. Rev. D*, 105, 083518, 2201.07771
 Fortuna, M. C., Hoekstra, H., Joachimi, B., Johnston, H., Chisari, N. E., Georgiou, C., & Mahony, C. 2020, arXiv e-prints, arXiv:2003.02700, 2003.02700
 Fosalba, P., Crocce, M., Gaztanaga, E., & Castander, F. J. 2013, ArXiv e-prints, 1312.1707
 Fosalba, P., Gaztanaga, E., Castander, F. J., & Manera, M. 2008, *MNRAS*, 391, 435, 0711.1540
 Fu, L. et al. 2014, *MNRAS*, 441, 2725, 1404.5469

- Gatti, M. et al. 2020, MNRAS, 498, 4060, 1911.05568
- . 2023a, arXiv e-prints, arXiv:2307.13860, 2307.13860
- . 2023b, arXiv e-prints, arXiv:2310.17557, 2310.17557
- Giblin, B., Cai, Y.-C., & Harnois-Déraps, J. 2023, MNRAS, 520, 1721, 2211.05708
- Giblin, B. et al. 2021, A&A, 645, 105
- . 2018, MNRAS, 480, 5529, 1805.12084
- Gruen, D., & Brimiouille, F. 2017, MNRAS, 468, 769, 1610.01160
- Gruen, D. et al. 2018, Phys. Rev. D, 98, 023507, 1710.05045
- Gupta, A., & Nagar, D. 1999, Matrix Variate Distributions, Monographs and Surveys in Pure and Applied Mathematics (Taylor & Francis)
- Hamana, T. et al. 2020, PASJ, 72, 16, 1906.06041
- Harnois-Déraps, J., Giblin, B., & Joachimi, B. 2019, A&A, 631, A160, 1905.06454
- Harnois-Déraps, J., Martinet, N., Castro, T., Dolag, K., Giblin, B., Heymans, C., Hildebrandt, H., & Xia, Q. 2021, MNRAS, 506, 1623, 2012.02777
- Harnois-Déraps, J., Martinet, N., & Reischke, R. 2022, MNRAS, 509, 3868, 2107.08041
- Harnois-Déraps, J., & van Waerbeke, L. 2015, MNRAS, 450, 2857, 1406.0543
- Hartlap, J., Simon, P., & Schneider, P. 2007, A&A, 464, 399, arXiv:astro-ph/0608064
- Heitmann, K., Lawrence, E., Kwan, J., Habib, S., & Higdon, D. 2014, ApJ, 780, 111, 1304.7849
- Heydenreich, S., Brück, B., Burger, P., Harnois-Déraps, J., Unruh, S., Castro, T., Dolag, K., & Martinet, N. 2022, arXiv e-prints, arXiv:2204.11831, 2204.11831
- Heydenreich, S., Brück, B., & Harnois-Déraps, J. 2021, A&A, 648, A74, 2007.13724
- Hilbert, S. et al. 2020, MNRAS, 493, 305, 1910.10625
- Hildebrandt, H. et al. 2021, A&A, 647, A124, 2007.15635
- Ho, M.-F., Bird, S., & Shelton, C. R. 2022, MNRAS, 509, 2551, 2105.01081
- Hoekstra, H., Herbonnet, R., Muzzin, A., Babul, A., Mahdavi, A., Viola, M., & Cacciato, M. 2015, MNRAS, 449, 685, 1502.01883
- Hotelling, H. 1931, The Annals of Mathematical Statistics, 2, 360
- Hoyle, B. et al. 2018, MNRAS, 478, 592, 1708.01532
- Ivezić, Ž. et al. 2019, ApJ, 873, 111, 0805.2366
- Izard, A., Fosalba, P., & Crocce, M. 2018, MNRAS, 473, 3051
- Jarvis, M., Bernstein, G., & Jain, B. 2004, MNRAS, 352, 338
- Jeffrey, N., Alsing, J., & Lanusse, F. 2021, MNRAS, 501, 954, 2009.08459
- Joachimi, B. et al. 2020, arXiv e-prints, arXiv:2007.01844, 2007.01844
- Joudaki, S. et al. 2020, A&A, 638, L1, 1906.09262
- Kacprzak, T. et al. 2016, MNRAS, 463, 3653
- Kannawadi, A. et al. 2019, A&A, 624, A92, 1812.03983
- Kilbinger, M. 2015, Reports on Progress in Physics, 78, 086901
- Kilbinger, M. et al. 2017, MNRAS, 472, 2126, 1702.05301
- Kilbinger, M., Schneider, P., & Eifler, T. 2006, A&A, 457, 15, astro-ph/0604520
- Kuijken, K. et al. 2019, A&A, 625, A2, 1902.11265
- . 2015, MNRAS, 454, 3500, 1507.00738
- Laigle, C. et al. 2016, ApJS, 224, 24, 1604.02350
- Lange, J. U. 2023, MNRAS, 525, 3181, 2306.16923
- Laureijs, R. et al. 2011, arXiv e-prints, arXiv:1110.3193, 1110.3193
- Lemos, P. et al. 2022, arXiv e-prints, arXiv:2202.08233, 2202.08233
- Lewis, A. 2019, arXiv e-prints, arXiv:1910.13970, 1910.13970
- Li, S.-S. et al. 2023a, arXiv e-prints, arXiv:2306.11124, 2306.11124
- Li, X. et al. 2023b, arXiv e-prints, arXiv:2304.00702, 2304.00702
- Li, Z., Liu, J., Zorrilla Matilla, J. M., & Coulton, W. R. 2019, Phys. Rev. D, 99, 063527, 1810.01781
- Lima, M., Cunha, C. E., Oyaizu, H., Frieman, J., Lin, H., & Sheldon, E. S. 2008, MNRAS, 390, 118, 0801.3822
- Lin, K., von Wietersheim-Kramsta, M., Joachimi, B., & Feeney, S. 2023, MNRAS, 524, 6167, 2212.04521
- Linke, L., Burger, P. A., Heydenreich, S., Porth, L., & Schneider, P. 2023, arXiv e-prints, arXiv:2302.12277, 2302.12277
- Liu, X., Yuan, S., Pan, C., Zhang, T., Wang, Q., & Fan, Z. 2023, MNRAS, 519, 594, 2210.07853
- Longley, E. P. et al. 2023, MNRAS, 520, 5016, 2208.07179
- Loureiro, A. et al. 2022, A&A, 665, A56, 2110.06947
- MacCrann, N. et al. 2022, MNRAS, 509, 3371, 2012.08567
- Mandelbaum, R. 2018, ARA&A, 56, 393, 1710.03235
- Marques, G. A. et al. 2023, arXiv e-prints, arXiv:2308.10866, 2308.10866
- Martin, N., Castro, T., Harnois-Déraps, J., Jullo, E., Giocoli, C., & Dolag, K. 2021, A&A, 648, A115, 2012.09614
- Martin, N., Harnois-Déraps, J., Jullo, E., & Schneider, P. 2020, arXiv e-prints, arXiv:2010.07376, 2010.07376
- Martin, N. et al. 2018, MNRAS, 474, 712, 1709.07678
- McCarthy, I. G., Schaye, J., Bird, S., & Le Brun, A. M. C. 2017, MNRAS, 465, 2936, 1603.02702
- Miller, L. et al. 2013, MNRAS, 429, 2858
- Pedregosa, F. et al. 2011, Journal of Machine Learning Research, 12, 2825
- Percival, W. J., Friedrich, O., Sellentin, E., & Heavens, A. 2022, MNRAS, 510, 3207, 2108.10402
- Planck Collaboration et al. 2020, A&A, 641, A6, 1807.06209
- Pyne, S., & Joachimi, B. 2021, MNRAS, 503, 2300, 2010.00614
- Schirmer, M., Erben, T., Hettterscheidt, M., & Schneider, P. 2007, A&A, 462, 875, astro-ph/0607022
- Schneider, A., Teyssier, R., Stadel, J., Chisari, N. E., Le Brun, A. M. C., Amara, A., & Refregier, A. 2019, JCAP, 2019, 020, 1810.08629
- Schneider, P. 1996, MNRAS, 283, 837, astro-ph/9601039
- Secco, L. F. et al. 2022a, Phys. Rev. D, 105, 103537, 2201.05227
- . 2022b, Phys. Rev. D, 105, 023515, 2105.13544
- Sellentin, E., & Heavens, A. F. 2016, MNRAS, 456, L132
- Shan, H. et al. 2018, MNRAS, 474, 1116, 1709.07651
- Sheldon, E. S., & Huff, E. M. 2017, ApJ, 841, 24
- Takahashi, R., Hamana, T., Shirasaki, M., Namikawa, T., Nishimichi, T., Osato, K., & Shiroshima, K. 2017, ApJ, 850, 24
- Tröster, T. et al. 2021, A&A, 649, A88, 2010.16416
- Troxel, M. A. et al. 2018, Phys. Rev. D, 98, 043528, 1708.01538
- Vakili, M. et al. 2020, arXiv e-prints, arXiv:2008.13154, 2008.13154
- van den Busch, J. L. et al. 2020, A&A, 642, A200, 2007.01846
- . 2022, A&A, 664, A170, 2204.02396
- van Waerbeke, L., et al. 2013, MNRAS, 433, 3373
- Wright, A. H., Hildebrandt, H., van den Busch, J. L., Heymans, C., Joachimi, B., Kannawadi, A., & Kuijken, K. 2020, A&A, 640, L14, 2005.04207
- Xavier, H. S., Abdalla, F. B., & Joachimi, B. 2016, FLASK: Full-sky Lognormal Astro-fields Simulation Kit, Astrophysics Source Code Library, record ascl:1606.015, 1606.015
- Zorrilla Matilla, J. M., Waterval, S., & Haiman, Z. 2020, AJ, 159, 284, 1909.12345
- Zuntz, J. et al. 2015, Astronomy and Computing, 12, 45
- Zürcher, D. et al. 2022, MNRAS, 511, 2075, 2110.10135
- Zürcher, D., Fluri, J., Sgier, R., Kacprzak, T., & Refregier, A. 2020, arXiv e-prints, arXiv:2006.12506, 2006.12506




Figure A1. Noise-subtracted peak count statistics measured from B -modes data (black squares) for a representative subset of the 30 tomographic bins, compared with the E -modes cosmological predictions.

APPENDIX A: B -MODES

To leading order, B -modes are not produced by gravitational lensing, hence their detection in cosmic shear data is generally regarded as an indication of residual systematics. As mentioned in Sec. 4.6, the aperture mass map statistics constructed on a grid inevitably induces B -modes from the missing sub-pixel contributions, resulting in a non-zero $M_\times$ signal. This section presents a careful investigation of the amplitude, origins and consequences of these induced B -modes. In particular, and we find that finite sampling of the shear field itself is also a source of B -modes in aperture mass maps, on top of pixelisation.

We first quantify here the strength of these effects by measuring the peak function from $M_\times(\theta)$ in our data, i.e. aperture mass maps in which the galaxies are rotated by 45 deg. The (noise-subtracted) signal N^B_{peaks} is shown in Fig. A1 for a representative subset of the tomographic bins. We observe that the residual signal is much flatter than what we would expect from a cosmological signal consistent with pure noise with a p -value of $p = 0.12$, above the threshold of $p = 0.01$ (the same threshold is used in the main text, and in the DES-Y3 results for this type of hypothesis testing, see Appendices G and D of Zürcher et al. 2022; Abbott et al. 2022, respectively). This agrees with A21, namely that there is no evidence of residual B -modes in the KiDS-1000 data. It is therefore safe to keep all data entries in our inference, but investigate further the source of the $M_\times(\theta)$ signal seen by eye in Fig. A1 to confirm it is not problematic.

We carried out peak count measurements of $M_\times(\theta)$ on 20 full survey realisations from the *Covariance Training Set* (again, these are pure E -mode mocks rotated by 45 degrees for this exercise), expecting to find large p -values in all of these trials. Instead, this test revealed that p -values range from 10^{-10} to 0.1. Some of these trials seem to rule out completely the null hypothesis (that the B -modes are consistent with pure noise), even though no B -mode exists at the catalogue level. The observed $M_\times$ signal must therefore come from the aperture map method itself, and is consequently a poor test for residual observational systematics.

Interestingly, the measured N^B_{peaks} averaged over 20 noise realisations has a p -value of 1.0, namely $\langle N^B_{\text{peaks}} \rangle = N^{\text{noise}}_{\text{peaks}}$, suggesting that these B -modes contain mostly noise, even though on a case-by-case some realisations see strong deviations. We hypothesise that this stems from E -modes leaking into B -modes due to an incomplete knowledge of the shear field: assuming a noiseless, pure E -mode shear field, the average cross-shear $\gamma_\times$ on a circle around every point in the field vanishes by definition, and thus $M_\times(\theta) \equiv 0$ holds everywhere. However, that is no longer guaranteed once the shear field is only known at a discrete set of




Figure B1. Tomographic weak lensing peak function in the *Baryons Training Set*. The coloured lines are obtained by scaling the GPR predictions (at the *Magneticum* cosmology) by the b_{bary} parameter, over the full prior range, demonstrating that peak statistics are fairly insensitive to changes in baryon feedback.

positions, as the average $\gamma_\times$ on a circle no longer necessarily vanishes. We investigate this by varying the number of source galaxies in our simulations. We achieve this by running our measurements on Stage-IV mocks created with a number density of 30.0 arcmin, introduced in Heydenreich et al. (2021), without any tomographic split. We measure on these the $M_\times$ signal from maps sampled a) at every pixel location, b) at all galaxy positions, and c) at galaxy positions downsampled to match the KiDS-1000 number density. In the first case, we find that the B -mode field $M_\times(\theta)$ vanishes completely (up to numerical precision). The second case induces B -modes of approximately 0.5% of the E -mode signal, whereas the third case (KiDS-1000-like number density) yields B -modes of approximately 4% of the E -mode signal. We note that this is likely to be exacerbated by the splitting the galaxies into different tomographic bins, which further decreases the number density per aperture. We further note that these tests were performed in the absence of shape noise to better isolate this effect.

More importantly, since these non-zero N^B_{peaks} are caused by finite sampling of the shear field, and that this sampling is exactly the same for the data and the *Cosmology Training Set*, the same amount of leakage should occur on average. In particular this should be fully converged in simulation-based model once averaged over the 50 mock survey $\times$ 10 noise realisations per cosmology (20 was shown to be enough in the discussion above). Therefore our inference must be immune to these by construction.

APPENDIX B: VALIDATION OF THE COSMOLOGY INFERENCE PIPELINE

In this section we present a series of validation tests we performed on our cosmology inference pipeline. First, we verified that the derivatives $\partial N_{\text{peaks}} / \partial \Delta m_a$, $\partial N_{\text{peaks}} / \partial \Delta z_a$, $\partial N_{\text{peaks}} / \partial b_{\text{bary}}$ and $\partial N_{\text{peaks}} / \partial A_{\text{IA}}$ are consistent with the results found in HD21 and Harnois-Déraps et al. (2022). The element-by-element values are different since these are survey-specific, but they agree qualitatively. We also verified that the impact of increased N -body force is of no consequence, consistent with HD21. This is achieved by carrying the inference with the *Validation Set* (high-




Figure B2. Cosmological constraints inferred from mock data vectors extracted from two of our *Cosmology Training Set* models.




Figure B3. Joint-survey mock analysis of the *Validation Training Set*.

resolution) instead of the mean of the *Covariance Training Set*, as done in HD21.

We also verified that the peak function measured from the *Baryons Training Set* is consistent with the *Cosmology Training Set* in Fig. B1. This is an important test, as the baryon mocks based on a completely independent N -body code. **This figure also shows the impact of varying b_{bary} on the data vector. Stronger feedback models (purple) tend to have fewer large peaks ($S/N > 2$), and more in the range $-1 < S/N < 1.5$. This is best seen in panel 4 $\cup$ 5 but a common feature to most panels.** Fig. B2 validates the KiDS-1000 cosmology inference pipeline by showing the excellent recovery of input parameters for two different cosmologies selected from the *Cosmology Training Set* (the fiducial Λ CDM model as well as w CDM model 12; we verified that other cosmologies work equally well).

We finally tested the joint-survey pipeline with the *Validation Set*

defined for both KiDS-1000 and DES-Y1 analyses, and report our results in Fig. B3. We observe that it recovers very well the input truth: the best-fit value is $S_8 = 0.818^{+0.030}_{-0.025}$, the maximum likelihood value is 0.831, while the truth is 0.813. These results are obtained from w CDM pipeline assuming the “clean” selection of S/N bin, marginalising over all nuisance parameters. All parameters are accurately recovered.

APPENDIX C: GOODNESS-OF-FIT FOR STUDENT- t LIKELIHOODS

Noisy numerical covariance matrices need to be treated carefully in likelihood analyses to avoid biases incurred during the inversion. A commonly used approach is to debias the inverse matrix with the Hartlap-Anderson coefficient (Hartlap et al. 2007), however this often leads to over-estimates in the contours. Instead, Sellentin & Heavens (2016) suggested to replace the Hartlap-corrected multivariate Gaussian likelihood by a Student- t distribution, which better accounts for the noise present when estimating the matrix from N_{sims} realisations of the data.

Once the likelihood has been sampled and the best-fit parameters found, one of the key subsequent steps is to estimate the goodness-of-fit. This is usually achieved by means of the p -value, which determines how likely it is that the difference between the best-fit model and the measured data is due to a random noise fluctuation. Given the number of degrees of freedom ν , best-fit χ^2 measurements from data that is well described by a multi-Gaussian likelihood will be sampled from a χ^2_ν distribution. Using this metric with noisy numerical covariance matrices will yield p -values that are biased towards low values if the inverse matrix is not Hartlap-corrected. Conversely, if corrected, the p -values are at risk to be on the high-side (Sellentin & Heavens 2016).

This is demonstrated by a toy model, which is created to follow our analysis: we generate a matrix A with 210^2 Gaussian random numbers (the same dimension as our KiDS-1000 analysis) and define a ‘true’ covariance matrix $\Sigma = A^T A$. We also define the ‘true’ data vector as the zero-vector.

Afterwards, the following procedure is repeated 10 000 times: we generate 1240 realizations of a multivariate normal distribution with mean 0 and covariance Σ , from which we estimate our sample covariance matrix C , mimicking the *Covariance Training Set*. We then also draw one additional realisation X of the same multivariate normal distribution, which constitutes our measurement. Finally, we calculate the p -value given X and C and a chosen p -value test, assuming that the degrees of freedom equal the number of elements in the data vector (since our toy model contains no free parameters).

This procedure yields 10 000 p -values which, if the chosen test is appropriate for our analysis, form a uniform distribution between 0 and 1. As can be seen in the lower panel of Fig. C1, the χ^2 -based p -value tests are heavily biased towards 0, as is expected. The Hartlap-corrected p -values are more uniformly distributed, but still show slight biases towards 0 and 1, and consequently a reduced probability towards central values. Although this effect is relatively weak for our setup, it becomes more prominent if the degrees of freedom increase. Nevertheless, this means that a Hartlap-corrected p -value test is more likely to favour extreme values, but it appears to be relatively robust.

An unbiased solution to this problem can be achieved by deriving the sampling distribution of our quadratic statistics¹⁸ defined in Eq. (6), specifically:

$$T_{\text{best-fit}}^2 = (d - x(\pi_{\text{best-fit}}))^T C^{-1} (d - x(\pi_{\text{best-fit}})), \quad (\text{C1})$$

¹⁸ The quadratic statistics described by Eq. 6 should not be labelled ‘ χ^2 ’ unless it is sampling a χ^2_ν distribution. We used the χ^2 notation in the main text only to align with the notation in the weak lensing literature.

where the data $\mathbf{d}$ has dimension p and is drawn from a normal distribution $\mathbf{d} \sim N(\boldsymbol{\mu}, \Sigma)$, for unknown mean $\boldsymbol{\mu}$ and unknown covariance Σ , and the C covariance is drawn from a Wishart distribution with $N_{\text{sims}} - 1$ degrees of freedom $(N_{\text{sims}} - 1)C \sim W_p(\Sigma, N_{\text{sims}} - 1)$.

We now define $LL^T = \Sigma^{-1}$ and $\mathbf{w} = L(\mathbf{d} - \mathbf{x}(\boldsymbol{\pi}_{\text{best-fit}}))$, such that $\text{Cov}[\mathbf{w}, \mathbf{w}] = \mathbb{I}_p$, the $p \times p$ identity matrix. With this, we can express Eq. (C1) as

$$T_{\text{best-fit}}^2 = (N_{\text{sims}} - 1)\mathbf{w}^T V^{-1} \mathbf{w}, \quad (\text{C2})$$

where we have defined $V = (N_{\text{sims}} - 1)LCL^T$ and note that $V \sim W_p(\mathbb{I}_p, N_{\text{sims}} - 1)$. We can now introduce an orthogonal matrix $M^T M = \mathbb{I}_p$ with the first row being $\frac{\mathbf{w}}{\|\mathbf{w}\|}$ and the others orthogonal to it, such that

$$M\mathbf{w} = \begin{pmatrix} \|\mathbf{w}\| \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \quad (\text{C3})$$

Conditional on M , we have that $Q = MVM^T \sim W_p(\mathbb{I}_p, N_{\text{sims}} - 1)$. Since this does not depend on M , $Q \sim W_p(\mathbb{I}_p, N_{\text{sims}} - 1)$ also holds in the unconditional case. With this transformation, Eq. (C2) can be written as

$$T_{\text{best-fit}}^2 = (N_{\text{sims}} - 1)\|\mathbf{w}\|^2 (Q^{-1})_{11}^{-1}, \quad (\text{C4})$$

with $(Q^{-1})_{11}^{-1}$ being the 1-1 entry of Q^{-1} . Writing out $\|\mathbf{w}\|^2$ as

$$\begin{aligned} \mathbf{w}^T \mathbf{w} &= (\mathbf{d} - \mathbf{x}(\boldsymbol{\pi}_{\text{best-fit}}))^T L^T L (\mathbf{d} - \mathbf{x}(\boldsymbol{\pi}_{\text{best-fit}})) \\ &= (\mathbf{d} - \mathbf{x}(\boldsymbol{\pi}_{\text{best-fit}}))^T \Sigma^{-1} (\mathbf{d} - \mathbf{x}(\boldsymbol{\pi}_{\text{best-fit}})), \end{aligned} \quad (\text{C5})$$

we recognise this as the usual χ^2 quantity where the true covariance Σ is assumed to be known. In other words, $\|\mathbf{w}\|^2 \sim \chi_\nu^2$, with $\nu = p - n_{\text{eff}}$, where n_{eff} is the effective number of free parameters that are being varied when finding $\boldsymbol{\pi}_{\text{best-fit}}$, **which accounts for the fact that the model may contain both constrained and unconstrained parameters.**

Returning to the last term in Eq. (C4), we have

$$(Q^{-1})_{11}^{-1} = \frac{1}{(Q^{-1})_{11}} = Q_{11} - Q_{12}Q_{22}^{-1}Q_{21}, \quad (\text{C6})$$

where Q_{12} and Q_{22} are the $1 \times (p - 1)$ and $(p - 1) \times (p - 1)$ submatrices of Q . From this follows (e.g. Gupta & Nagar 1999) that $\frac{1}{(Q^{-1})_{11}} \sim W_1(I_1, N_{\text{sims}} - p) = \chi_{N_{\text{sims}} - p}^2$. Putting things together, we therefore have that

$$T_{\text{best-fit}}^2 \sim (N_{\text{sims}} - 1) \frac{\chi_{p - n_{\text{eff}}}^2}{\chi_{N_{\text{sims}} - p}^2} = \frac{(N_{\text{sims}} - 1)(p - n_{\text{eff}})}{(N_{\text{sims}} - p)} F_{p - n_{\text{eff}}, N_{\text{sims}} - p}, \quad (\text{C7})$$

where $F_{p - n_{\text{eff}}, N_{\text{sims}} - p}$ is the F -distribution. For the case of no free parameters, $n_{\text{eff}} = 0$, this reduces to Hotelling's T^2 distribution (Hotelling 1931).

To calculate a p -value, one just has to evaluate the cumulative distribution function of the F -statistics at $\frac{N_{\text{sims}} - p}{(p - n_{\text{eff}})(N_{\text{sims}} - 1)} T^2$, replacing the $\chi^2(p - n_{\text{eff}})$ distribution. Clearly seen in Fig. C1, this constitutes an ideal solution for our toy model, so we use this statistics to estimate the p -values of our measurements.

When applied to our fiducial KiDS-1000 peak count data-vector, along with the best-fit model and our numerical covariance matrix, we obtain p -values of 0.43 with the (unbiased) Hotelling's statistics and for the (slightly biased) Hartlap-corrected χ^2 approach, and 0.02 for the (heavily biased) normal χ^2 statistics.

In this toy example, the difference between the Hotelling's and the Hartlap-corrected p -values distributions is quite small, however this is not always the case. The upper panel of Fig. C1 shows a second case where now the number of degrees of freedom has been increased to




Figure C1. Distribution of p -values extracted from our toy examples based on three commonly used goodness-of-fit statistics, for $\nu = 210$ (lower) and $\nu = 400$ (upper). Given a noisy numerical covariance matrix, only the Hotelling's T^2 distribution returns p -values evenly sampling the range $[0, 1]$; the χ^2 distribution (orange) is heavily skewed towards low p -values, while the Hartlap-corrected χ^2 (green) is slightly skewed towards extrema p -values.

$\nu = 400$, overshooting our joint-survey setup, but close to typical sizes of data vectors used in 2pt statistics. Keeping N_{sims} unchanged, in this case the Hartlap-corrected distribution shows a clear excess towards low and high p -values. The Hotelling's distribution is still flat however, showcasing the advantage of the F -statistics.

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