

Dark Energy Survey Year 1 results: the impact of galaxy neighbours on weak lensing cosmology with IM3SHAPE

S. Samuroff,^{1★} S. L. Bridle,¹ J. Zuntz,² M. A. Troxel,^{3,4} D. Gruen,^{5,6†} R. P. Rollins,¹ G. M. Bernstein,⁷ T. F. Eifler,⁸ E. M. Huff,⁸ T. Kacprzak,⁹ E. Krause,⁵ N. MacCrann,^{3,4} F. B. Abdalla,^{10,11} S. Allam,¹² J. Annis,¹² K. Bechtol,¹³ A. Benoit-Lévy,^{14,10,15} E. Bertin,^{14,15} D. Brooks,¹⁰ E. Buckley-Geer,¹² A. Carnero Rosell,^{16,17} M. Carrasco Kind,^{18,19} J. Carretero,²⁰ M. Crocce,²¹ C. B. D’Andrea,⁷ L. N. da Costa,^{16,17} C. Davis,⁵ S. Desai,²² P. Doel,¹⁰ A. Fausti Neto,¹⁶ B. Flaugher,¹² P. Fosalba,²¹ J. Frieman,^{12,23} J. García-Bellido,²⁴ D. W. Gerdes,^{25,26} R. A. Gruendl,^{18,19} J. Gschwend,^{16,17} G. Gutierrez,¹² K. Honscheid,^{3,4} D. J. James,^{27,28} M. Jarvis,⁷ T. Jeltema,²⁹ D. Kirk,¹⁰ K. Kuehn,³⁰ S. Kuhlmann,³¹ T. S. Li,¹² M. Lima,^{32,16} M. A. G. Maia,^{16,17} M. March,⁷ J. L. Marshall,³³ P. Martini,^{3,34} P. Melchior,³⁵ F. Menanteau,^{18,19} R. Miquel,^{36,20} B. Nord,¹² R. L. C. Ogando,^{16,17} A. A. Plazas,⁸ A. Roodman,^{5,6} E. Sanchez,³⁷ V. Scarpine,¹² R. Schindler,⁶ M. Schubnell,²⁶ I. Sevilla-Noarbe,³⁷ E. Sheldon,³⁸ M. Smith,³⁹ M. Soares-Santos,¹² F. Sobreira,^{40,16} E. Suchyta,⁴¹ G. Tarle,²⁶ D. Thomas,⁴² D. L. Tucker¹² and (DES Collaboration)

Affiliations are listed at the end of the paper

Accepted 2017 December 16. Received 2017 November 27; in original form 2017 August 14

ABSTRACT

We use a suite of simulated images based on Year 1 of the Dark Energy Survey to explore the impact of galaxy neighbours on shape measurement and shear cosmology. The HOOPOE image simulations include realistic blending, galaxy positions, and spatial variations in depth and point spread function properties. Using the IM3SHAPE maximum-likelihood shape measurement code, we identify four mechanisms by which neighbours can have a non-negligible influence on shear estimation. These effects, if ignored, would contribute a net multiplicative bias of $m \sim 0.03\text{--}0.09$ in the Year One of the Dark Energy Survey (DES Y1) IM3SHAPE catalogue, though the precise impact will be dependent on both the measurement code and the selection cuts applied. This can be reduced to percentage level or less by removing objects with close neighbours, at a cost to the effective number density of galaxies n_{eff} of 30 per cent. We use the cosmological inference pipeline of DES Y1 to explore the cosmological implications of neighbour bias and show that omitting blending from the calibration simulation for DES Y1 would bias the inferred clustering amplitude $S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.5}$ by 2σ towards low values. Finally, we use the HOOPOE simulations to test the effect of neighbour-induced spatial correlations in the multiplicative bias. We find the impact on the recovered S_8 of ignoring such correlations to be subdominant to statistical error at the current level of precision.

Key words: gravitational lensing: weak – galaxies: statistics – cosmological parameters – cosmology: observations.

* E-mail: simon.samuroff@postgrad.manchester.ac.uk

† NASA Einstein Fellow.

1 INTRODUCTION

A standard and well-tested prediction of general relativity is that a concentration of mass will distort the spacetime around it, and thus produce a curious phenomenon called gravitational lensing. The most obvious manifestation is about massive galaxy clusters, where background galaxies can be elongated into crescent-shaped arcs. So-called strong lensing of galaxies was first observed in the late 1980s and has been confirmed many times since. A subtler, but from a cosmologist’s perspective more powerful, consequence of gravitational lensing is that background fluctuations in the density of dark matter will induce coherent distortions to photons’ paths. This effect is known as cosmic shear, and it was first detected by four groups at around the same time close to two decades ago (Bacon, Refregier & Ellis 2000; Kaiser, Wilson & Luppino 2000; Van Waerbeke et al. 2000; Wittman et al. 2000). Cosmic shear has the potential to be the single most powerful probe in the toolbox of modern cosmology. The spatial correlations due to lensing are a direct imprint of the large-scale mass distribution of the Universe. Thus it allows one to study the total mass of the Universe and the growth of structure within it (Maoli et al. 2001; Jarvis et al. 2006; Massey et al. 2007b; Heymans et al. 2013; Kilbinger et al. 2013; Dark Energy Survey Collaboration 2016; Jee et al. 2016; Hildebrandt et al. 2017; Köhlinger et al. 2017), or to map out the spatial distribution of dark matter on the sky (e.g. Kaiser 1994; Van Waerbeke et al. 2013; Chang et al. 2015). As a probe of both structure and geometry, cosmic shear is also attractive as a method for shedding light on the as yet poorly understood component of the Universe known as dark energy (Albrecht et al. 2006; Weinberg et al. 2013). Alternatively, lensing will allow us to place ever more stringent tests of our theories of gravity (Simpson et al. 2013; Harnois-Déraps et al. 2015; Brouwer et al. 2017). This is also theoretically very clean, responding directly to the power spectrum of dark matter, which is affected by baryonic physics only on small scales, and avoids recourse to poorly understood phenomenological rules. Indeed galaxy number density enters only at second order as a weighting of the observed shear due to the fact that one can only sample the shear field where there are real galaxies (Schmidt et al. 2009).

Though well modelled theoretically, cosmic shear is technically highly challenging to measure; as with all these probes it is not without its own sources of systematic error. It also cannot be reiterated too many times that the shear component of even the most distant galaxy’s shape is subdominant to noise by an order of magnitude. Indeed, the ambitions of the current generation of cosmology surveys will require sub-per cent level uncertainties (both systematic and statistical) on what is already a tiny cosmological ellipticity component $g \sim 0.01$.

It was realized early on how significant the task of translating photometric galaxy images into unbiased shear measurements would be. In response came a series of blind shear measurement challenges, designed to review, test, and compare the best methods available. The first of these, called STEP1 (Heymans et al. 2006) grew out of a discussion at the 225th IAU Symposium in 2004. The exercise was based around a set of simple SKYMAKER simulations (Bertin & Fouqué 2010), which were designed to mimic ground-based observations but with analytic galaxies and point spread functions (PSFs) and constant shear. The algorithms at this point represented a first wave of shear measurement codes and included several moment-based algorithms (Kaiser, Squires & Broadhurst 1995; Kuijken 1999; Rhodes, Refregier & Groth 2001), some early forward modelling methods (Bridle et al. 2002), as well as

a technique called shapelets, which models a light profile as a set of 2D basis functions (Bernstein & Jarvis 2002; Refregier & Bacon 2003).

The simulations and the codes themselves steadily grew in complexity. STEP2 was followed by series of GREAT challenges (Massey et al. 2007a; Bridle et al. 2009; Kitching et al. 2011; Mandelbaum et al. 2014), which focused on different aspects of shape measurement bias and have been essential in quantifying a number of significant effects. In recent years the drive to find ever more accurate ways to measure shear has intensified, with many novel approaches being suggested. For example Fenech Conti et al. (2017) use a form of self-calibration, which repeats the shape measurement on a test image based on the best-fitting model for each galaxy. A related approach, named metacalibration, involves deriving corrections to the galaxy shape measurements directly from the data, using modified copies of the image with additional shear (Huff & Mandelbaum 2017; Sheldon & Huff 2017). More advanced moment-based approaches include the BFD method (Bernstein & Armstrong 2014), which derives a prior on the ensemble ellipticity distribution using deeper fields, and (Herbonnet, Buddendiek & Kuijken 2017), a similar approach which builds ensemble shear estimates using shear nulling.

This paper is intended as a companion study to Zuntz et al. (2017) (Z17), where we present two shear catalogues derived from Year One of the Dark Energy Survey (DES Y1) data set. It is also presented alongside a raft of other papers, which use both catalogues and show them to be consistent in a number of different scientific contexts (Chang et al. 2017; Dark Energy Survey Collaboration 2017; Prat et al. 2017; Troxel et al. 2017) containing 22 million and 35 million galaxies, respectively, these catalogues are the product of two independent maximum likelihood codes. The first, called IM3SHAPE, implements simultaneous fits using multiple models and we calibrate externally using simulations. The second implements a Gaussian model fitting algorithm, supplemented by shear response corrections using METACALIBRATION. Whereas in Z17 we focus on the catalogues themselves, presenting a raft of calibration tests and a broad overview of the value-added data products, here we use the same resources to explore a narrower topic: the impact of image plane neighbours on shear measurement. Specifically we use the image simulations described in Z17, from which the Y1 IM3SHAPE calibration is derived, to explore the mechanisms for neighbour bias, and then propagate the results to mock shear two-point data to investigate the consequences for weak lensing cosmology. The results presented in this paper will be somewhat dependent on the choice of measurement algorithm, selection cuts, and the configuration of the object detection code. Unlike previous studies on this subject, however, we make use of a highly realistic simulation and measurement pipeline. Our choices on each of aspects are realistic, if not unique, for a leading-edge cosmology analysis.

It is worth remarking, however, that the tests described in this paper make use of IM3SHAPE only, and should not be assumed to apply generically to its sister Y1 METACALIBRATION catalogue. A complementary set of tests using METACALIBRATION are presented in section 4.5 of Z17.

This paper is structured as follows. In Section 2 we briefly review the formalism of lensing and the observables discussed in this work. In Section 3 we present a series of numerical calculations using a toy model to characterize neighbour bias. Section 4 describes the simulated DES Y1 data sets, generated using our HOOPOE simulator. We test the earlier predictions under more typical observing conditions in Section 5, and extend them into a

quantitative set of results using the more extensive Y1 HOPOE data set. Section 6 then presents a numerical analysis designed to test the cosmological implications of neighbour bias of the nature and magnitude found in our simulations. We conclude in Section 7.

2 THE SHEAR MEASUREMENT PROBLEM

The problem of shape measurement is far more intricate than it might first appear. Any cosmological analysis based on cosmic shear is reliant on a series of technical choices, which can have a non-trivial impact on measurement biases, precision, and cosmological sensitivity. Specifically we must choose (a) how to parametrize each galaxy's shape, and which measurement method to use to estimate it, (b) what selection criteria are needed to obtain data of sufficiently high quality for cosmology, and (c) how biased is the measurement and what correction is needed? These choices should be made on a case-by-case basis, since the optimal solutions are dependent on a number of survey-specific factors. We discuss each briefly in turn below.

2.1 Shape measurement with IM3SHAPE

The shape measurements upon which the following analyses are based make use of the maximum likelihood model fitting code IM3SHAPE¹ (Zuntz et al. 2013). It is a well-tested and understood algorithm, which has since been used in a range of lensing studies (Whittaker, Brown & Battye 2015; Dark Energy Survey Collaboration 2016; Kacprzak et al. 2016; Clampitt et al. 2017). It was also one of two codes used to produce shear catalogues in the Science Verification (SV) stage and Year 1 of the Dark Energy Survey. We refer the reader to Jarvis et al. (2016) (hereafter J16) and Z17 for the most recent modifications to the code.

We use the definition of the flux signal-to-noise ratio of Z17, J16, and Mandelbaum et al. (2015):

$$S/N \equiv \frac{\left(\sum_{i=1}^{N_{\text{pix}}} f_i^m f_i^{\text{im}} / \sigma_i^2 \right)}{\left(\sum_{i=1}^{N_{\text{pix}}} f_i^m f_i^m / \sigma_i^2 \right)^{\frac{1}{2}}}. \quad (1)$$

The indices $i = (1, 2, \dots, N_{\text{pix}})$ run over all pixels in a stack of image cutouts at the location of a galaxy detection. The model prediction and observed flux in pixel i are denoted by f_i^m and f_i^{im} , respectively, and σ_i is the RMS noise. This signal-to-noise measure is maximized when the differences between the model and the image pixel fluxes are small. Note that if the best-fitting model f^m is identical for two different postage stamps, S/N will favour the image with the greater total flux.

A useful size measure, referred to as R_{gp}/R_p is defined as the measured full width at half-maximum (FWHM) of the galaxy after PSF convolution, normalized to the PSF FWHM.

Real galaxy images are not perfectly symmetric (i.e. size is not independent of azimuthal angle about a galaxy's centroid), and single-number size estimates are obtained by circularizing (azimuthally averaging) the galaxy profile and computing the weighted quadrupole moments of the resulting image. For each galaxy we take the mean measured size across exposures.

2.2 Shear measurement bias

There are many ways bias can enter an ensemble shear estimate based on a population of galaxies. Although the list is not exhaustive, a handful of mechanisms are particularly prevalent and have been extensively discussed in the literature.

(i) **Noise Bias:** On addition of pixel noise to an image, the best-fitting parameters of a galaxy model will not scale linearly with the noise variance. This is as an estimator bias as much as a measurement bias, and results in an asymmetric, skewed likelihood surface (Hirata & Seljak 2003; Kacprzak et al. 2012; Refregier et al. 2012; Miller et al. 2013). Any code which uses the point statistics of the distribution (either mean or maximum likelihood) as a single-number estimates of the ellipticity results in a bias. This is true even in the idealized case where the galaxy we are fitting can be perfectly described by our analytic light profile. The bias is sensitive to the noise levels and also the size and flux of the galaxy, and thus is specific to the survey and galaxy sample in question. For likelihood-based estimates one solution would be to impose a prior on the ellipticity distribution and propagate the full posterior. However, the results can become dependent on the accuracy of that prior, and such codes require cautious testing using simulations (Bernstein & Armstrong 2014; Simon & Schneider 2017)

(ii) **Model Bias:** In reality galaxies are not analytic light profiles with clear symmetries. For the purposes of model-fitting, however, we are constrained to use models with a finite set of parameters. A model which does not allow sufficient flexibility to capture the range of morphological features seen in the images may produce biased shape measurements (Lewis 2009; Voigt & Bridle 2010; Kacprzak et al. 2014).

(iii) **Selection Bias:** Even if we were to devise an ideal shape measurement algorithm, capable of perfectly reconstructing the histogram of ellipticities in a certain population of galaxies, our attempts to estimate the cosmological shear could still be biased. If a measurement step prefers rounder objects or those with a particular orientation, the result would be a net alignment that could be mistaken as having cosmological origin. In practice selection bias can commonly arise if the object detection algorithm preferentially selects galaxies aligned with the PSF ellipticity (e.g. Kaiser et al. 2000; Bernstein & Jarvis 2002), or if it fails less frequently on rounder galaxies (Hirata & Seljak 2004). It is such effects that make post facto quality cuts on quantities such as signal-to-noise or size (both of which correlate with ellipticity) particularly delicate.

(iv) **Neighbour bias:** In practice, galaxies in photometric surveys like DES are not ideal isolated objects. Rather, they are extracted from a crowded image plane using imperfect deblending algorithms. The term 'neighbour bias' refers to any biases in the recovered shear arising from the interaction between galaxies in the image plane. This can include both the direct impact on the per-galaxy shapes (e.g. Hoekstra, Viola & Herbonnet 2017) and changes in the selection function (e.g. Hartlap et al. 2011). Neighbour bias is the subject of relatively few previous studies, and is the focus of this paper.

3 A SIMPLE TOY MODEL FOR NEIGHBOUR BIAS

To develop a picture of how image plane neighbours affect shear estimates with IM3SHAPE, we build a simplified toy model. Using

¹ <https://bitbucket.org/joezuntz/im3shape-git>




Figure 1. Postage stamp snapshots of the basic two-object toy model described in Section 3. The overlain ellipse shows the maximum likelihood fit to the image. The panels show three neighbour positions in the range $\theta = [0, \pi/2]$ rad. The best-fitting ellipticity and half-light radius are shown above each image. In all cases the input values are $e = (0, 0)$, $r = 0.5$ arcsec.

GALSIM² we generate a 48×48 pixel postage stamp containing a single exponential disc profile convolved with a tiny spherically symmetric PSF (though we confirm that our results are insensitive to the exact size of the PSF). We can then apply a small shear along one coordinate axis prior to convolution and use IM3SHAPE to fit the resulting image.

In the absence of noise or model bias the maximum of the likelihood of the measured parameters coincides exactly with the input values. The basic setup then has four adjustable parameters: the flux and size of the galaxy plus two ellipticity components, denoted by f_c , r_c , g_1^{tr} , and g_2^{tr} . Unless otherwise stated we fix these to the median values measured from the DES Y1 IM3SHAPE catalogue. We do not model miscentering error between the true galaxy centroid and the stamp centre.

It is worth noting that neither this basic model nor the more complex simulations that follow attempt to model spatial correlations in shear. Even at different redshifts, a real neighbour-central pair share some portion of their line of sight. These spatial correlations will amplify the impact of blending, and are worthy of future investigation. This is, however, likely a second-order effect of neighbours, and we postpone such study to a future date.

To explore the interaction in single neighbour-galaxy instances we introduce a round neighbour into the postage stamp, convolved with the same nominal PSF.

This simple two-galaxy set up is shown in Fig. 1, with the best-fitting elliptical profile overlain in black. As illustrated in Fig. 2, as the second object is moved outwards, the inferred shape of the blend is gradually elongated up to some cut-off; beyond this the two are sufficiently well separated that the neighbour effectively contributes a flat background level to the stamp.

Next, the calculation is repeated with a moderately sheared neighbour ($g_2 = 0.1$). The result is shown by the blue lines in Fig. 3. In cases where the objects share a large portion of their half-light radii, we are fitting a strongly blended pair with a single profile, and the neighbour/central distinction becomes difficult to define. The best-fitting ellipticity recovered from the blended image in this regime is not a pure measurement of either galaxy's shape, but rather a linear combination of the two.

While useful for understanding what follows, the impact of neighbours on individual galaxy instances is not particularly informative about the impact on cosmic shear measurements. Even significant bias in the per-object shapes could average away over many galaxies with no residual impact on the recovered shear. More important is the *collective* response to neighbours. To estimate the ensemble




Figure 2. Tangential shear measured using the numerical toy model described in Section 3 as a function of radial neighbour distance. The solid purple line shows the shape component aligned with the central-neighbour separation vector and the dot-dashed line is measured along axes rotated through 45° . Note that the latter is smaller than 10^{-6} at all points on this scale. The dashed and dotted black lines show the same ellipticity components when the neighbour is sheared in the e_2 direction by $g_2 = 0.1$.

effect, we measure a neighbour-central image at 70 positions on a ring of neighbour angles. Again, under zero shear $g^{\text{tr}} = 0$ the measured shape is constant in magnitude, and the two components oscillate about 0 with peaks of amplitude $|\tilde{g}(\theta|0, d_{\text{gn}})|$. This sinusoidal variation is shown by the dotted lines in Fig. 3 (right) at two values of the neighbour distance d_{gn} (7 and 8 pixels). By averaging over a (large) number of neighbours one is effectively marginalizing over the neighbour position angle θ , which results in an unbiased measurement of the shear $\langle \tilde{g}(\theta|g^{\text{tr}} = 0, d_{\text{gn}}) \rangle_\theta = g^{\text{tr}} = 0$. A non-zero shear $g^{\text{tr}} \neq 0$, however breaks the symmetry of the system. A galaxy sheared along one axis will not respond to a neighbour in the same way irrespective of θ , which can result in a net bias. To show this we fix $g^{\text{tr}} = -0.05$ and proceed as before. The solid lines in Fig. 3 (right) show the periodicity in the measured shear at two d_{gn} . The mean value averaged over θ is shifted incrementally away from the input shear, shown by the horizontal dot-dashed line. Specifically we should note that the peaks below g^{tr} at $\pi/2$ and $3\pi/2$ rad are deeper and narrower than those above it. Imagine one has a strongly elliptical central profile, with a neighbour placed at a position angle of π radians from its major axis. If one rotates the neighbour at constant d_{gn} , one is gradually shifting it out of the relatively flat low wings of the central galaxy's light profile closer

² <https://github.com/GalSim-developers/GalSim>




Figure 3. Best-fitting galaxy ellipticity as a function of neighbour position angle at fixed neighbour distance d_{gn} from the toy model described in the text. The two panels (left, right) show the same central-neighbour system ($g^{\text{tr}} = -0.05$), but with different d_{gn} (7 and 8 pixels) and biases m (shown atop each panel). The solid line in each case is the recovered galaxy shape at each θ , and the integrated mean along this range is shown by the horizontal dot-dashed line. The dotted lines show the zero-shear shape (i.e. the ellipticity that would be measured if the input shear were zero), but shifted downwards such that the mean is at -0.05 . Finally, to illustrate the (a)symmetry of the system we show the solid line flipped about $y = g_1^{\text{tr}}$ and shifted by $\pi/2$ rad as a dashed curve.

to the core. Perturbing an object about $\theta = \pi$ rad results in a much greater change in the local flux gradient, $\nabla f_c(x, y)$ than doing the same about $\theta = 0$ rad. This asymmetry results in the asymmetry in the width of the positive and negative peaks seen in Fig. 3. The depth of the peak can be explained qualitatively by similar arguments. At $\theta = 0$ a neighbour of given flux is closer to the centre of the light distribution and thus has a greater flux overlap with the central galaxy than at $\theta = \pi$. Naturally, then, one might expect the first neighbour to have less impact than second. Returning to Fig. 3, we can see that the two effects are in competition. Depending on the exact neighbour configuration, the simultaneous narrowing and deepening the negative peaks can result in a bias in the neighbour-averaged ellipticity towards large or small values.

The level of this effect will clearly correlate with the magnitude of the shear, and so induce a multiplicative bias. To illustrate this point the above exercise is repeated with a range of different input shears. At small shears ($|g| < 0.1$), the neighbour-induced bias $\tilde{g} - g^{\text{tr}}$ is well approximated as linear in g^{tr} . Though the above numerical exercise demonstrates that it is *possible* for significant multiplicative bias to arise as a result of neighbours, it does not make a clear prediction of the magnitude or even the sign. Indeed, our toy model is effectively marginalized over θ , but there is nothing to guarantee that fixing the other neighbour parameters to the median measured values is representative of the real level of neighbour bias in a survey like DES. Motivated by this observation we add a final layer of complexity to the model, as follows. A single neighbour-central realization is created as before, defined by a unique set of model parameters. Now, however, the values of those parameters $\mathbf{p} = (d_{\text{gn}}, f_n, r_n, f_c, r_c)$ are drawn randomly from the DES data. As these quantities will, in reality, be correlated we sample from the five-dimensional joint distribution rather than each 1D histogram individually. We then fit the model at 70 neighbour angles and two input shears $g_{\pm} = \pm 0.05$ (a total of 140 measurements), and estimate the multiplicative bias as a two-point finite-difference derivative:

$$m + 1 = \frac{\langle \tilde{g}(\theta|g_+) \rangle_{\theta} - \langle \tilde{g}(\theta|g_-) \rangle_{\theta}}{g_+ - g_-}. \quad (2)$$

This process is repeated to create 1.33M unique toy model realizations. Binning by neighbour distance we can then make a rough

prediction for the level of neighbour-induced bias and the angular scales over which it should act. The majority of cases yield a negative bias, particularly at low neighbour separation (where the broadening of the peak around $\theta = \pi$ dominates over the increased flux overlap at $\theta = 0$ rad). In the real data, of course, we apply a quality-based selection and übersege object masking (J16), both of which are neglected here. We can, however, test the impact of selecting on fitted quantities that respond to neighbour bias. Imposing a flat prior on the centroid offset $\Delta r_0 = (x_0^2 + y_0^2)^{1/2}$ (i.e. discarding randomly generated model realizations where the galaxy centroid is displaced from the stamp centre by more than a fixed number of pixels) changes the shape of this curve significantly, as illustrated by the thick purple line in Fig. 4.

One can understand the difference between the results with and without the centroid cut as a form of selection bias, whereby the cut preferentially removes toy model realizations in which the neighbour is bright relative to the central galaxy. At any given d_{gn} we are left with a relative overrepresentation of galaxies with $f_n/f_c \ll 1$. Faint neighbours, which in reality tend to be compact high-redshift objects, have little impact when they sit on the outskirts of the central profile (the regime which produces negative m). The same faint galaxy has a stronger impact if it is rotated to a position closer to the centre of the central’s flux profile. Thus one might expect a selection on Δr_0 to make the mean m in a particular bin less negative (or even positive) by preferentially removing brighter galaxies.

4 HOOPOE IMAGE SIMULATIONS

In this section we provide a brief overview of the simulation pipeline. The process is the same as that described in section 4 of Z17, and we refer the reader to that work for more detail. The end point of the pipeline is a cloned set of survey images with many of the observable characteristics of a chosen set of parent images, but for which we know the input noise properties and galaxy population perfectly. The simulated images inherit the pixel masking, PSF variation and noise maps measured from the progenitor data. Each simulated galaxy is then inserted into a subset of overlapping exposures and into the co-add at the position of a real detection in the DES Y1 data. Object detection is rerun on the new co-add images and galaxy cutouts and new segmentation masks are extracted and stored in the MEDS format described by J16. The mock survey footprint is shown in Fig. 5. In the lower panels we show an example of a simulated co-add (left) and the spatial variation in PSF orientation within the same image (right).

4.1 Parent data

We use reduced images from DES Y1 (Diehl et al. 2014) as input to the simulations discussed in this paper. The Dark Energy Survey is undertaking a 5 yr programme with the ultimate aim of observing $\sim 5000 \text{ deg}^2$ of the southern sky to ~ 24 th magnitude in five optical bands, *grizY*, covering $0.40\text{--}1.06 \mu\text{m}$. The data set is recorded using a 570 megapixel camera called DECam (Flaugher et al. 2015), which has a pixel size of 0.26 arcsec. In full it will consist of ~ 10 interwoven sets of exposures in the $g, r, i, z,$ and Y bands.

The Y1 data were collected between 2013 August and 2014 February, and cover a substantially larger footprint than the preliminary SV stage at 1500 deg^2 , albeit to a reduced depth. Details of the reduction and processing are presented in Z17. Our HOOPOE simulations use a selection of the total 3000 $0.75 \times 0.75 \text{ deg}$ co-added patches known as ‘tiles’.

4.2 Input galaxy selection

For populating the mock survey images a sample of real galaxy profiles from the *HST* COSMOS field, imaged at significantly lower noise and higher resolution than DES by the *Hubble Space Telescope* Advanced Camera for Surveys (*HST* ACS) (Scoville et al. 2007). The COSMOS catalogue extends significantly deeper than the Y1 detection limit of $M_{r,\text{lim}} = 24.1$, extending to roughly 27.9 mag in the SDSS *r* band. A main sample for our DES Y1 simulations is defined by imposing a cut at <24.1 mag.

Since the DES images do not cut off abruptly at 24th magnitude, in reality they contain a tail of fainter galaxies that contribute flux are not identifiable above the pixel noise. To assess the impact of these objects on shape measurements in Y1, we simulate a population of sub-detection galaxies in addition to the main sample. In brief we use the full histogram of COSMOS magnitudes to estimate the number of faint galaxies within a given tile. The required profiles are selected randomly from the faint end of the COSMOS distribution. Each undetected galaxy is paired with a detection, and inserted at a random location within the overlapping bounds of the same (subset of) single-exposure images. A more detailed description of this process can be found in Z17.

If these galaxies were present in the data they would enter the background flux calculation, and thus the subtraction applied would change due to their presence. Since the simulation pipeline produces images effectively in a post-background subtraction state this is not captured. To test this we rerun the SExtractor background calculation on a handful of tiles drawn with and without the faint galaxies. The impact was found to be well approximated as a uniform shift in the background correction. A flux correction equal to the pixel-averaged flux of the sub-detection galaxies over each image plane is, then, applied to postage stamps prior to shape measurement.

In reality the overdensity of sub-threshold galaxies will be coupled to the density of detectable objects, which is clearly not the case in our simulations. To gauge the impact of this we perform the following test. Each tile is divided into a 6×6 grid, and the mean multiplicative bias is calculated in each sub-patch. We bin sub-patches according to the ratio $f_{\text{faint}} \equiv N_{\text{faint}}/N_{\text{det}}$, or the total number of faint galaxies relative to the number of detectable ones. The impact is significant, but not leading order; excluding patches outside the range $0.9 < f_{\text{faint}} < 1.1$ induces a shift of $\Delta m \sim -0.005$.

An independent noise realization is generated for each exposure using the weight map from the parent data. We simulate the noise in each pixel by drawing from a Gaussian of corresponding width. The co-addition process is not rerun, but rather we compute an independent noise field by drawing the flux in each pixel from a zero-centred Gaussian of width determined by the measured variance in that pixel.

4.3 Neighbour-free resimulations

For the purpose of untangling the impact of image plane neighbours we use the simulated HOOPOE images to create a new spin-off data set. In a subset of a little over 500 tiles we store the (convolved) input profile for each object and the noise-only cutout, taken from the same position in the image plane prior to objects being drawn. By adding together these two components we can generate a suite of spin-off MEDS files, which are equivalent to the results of a simpler neighbour-free simulation (e.g. Miller et al. 2013; J16). The pixel noise realization, COSMOS selection, and input shears, however, are identical to the progenitor HOOPOE simulations.

We will call this process ‘resimulating’, and the basic concept is illustrated in Fig. 6. The 506-tile set of neighbour-free data are named the WAXWING resimulations. Finally the (now empty) segmentation masks corresponding to the subtracted neighbours are also removed. In subsequent IM3SHAPE runs on these data we ignore the SExtractor flags obtained from the main simulations.

5 QUANTIFYING NEIGHBOUR BIAS WITH HOOPOE

Equipped with qualitative predictions from Section 3, we now turn to the question of neighbour bias in the more complete simulations described in Section 4. The mock survey was designed to capture as much of the complexity of shape measurements on real photometric data as possible. We refer to Section 4 of this paper for a short overview and to section 5 of Z17 for a more detailed discussion of the simulation pipeline and validation tests. The simulated galaxy catalogue used in the following is identical to the one used to calibrate the DES Y1 IM3SHAPE catalogue, including quality cuts and selection masks.

5.1 Single-galaxy effects

The most straightforward way to assess the impact of neighbours on individual shape measurements in our simulations is to rotate the measured shapes into a frame defined by the central-neighbour separation vector. Whereas in the earlier toy model we had only one neighbour per galaxy, we now have a crowded image plane containing many objects simultaneously. For simplicity, in the earlier case we included no masking. For HOOPOE we wish to mimic the process of shape measurement on real data as closely as possible. We generate new segmentation maps by running SExtractor on the simulated images, and incorporate them into our shape measurements using the überseg algorithm (J16). Each simulated galaxy is allocated a nearest neighbour using a k -d tree matching algorithm constructed on the co-add pixel grid using every galaxy simulated at r -band magnitude $M_r < 24.1$. The quantities d_{gn} and θ are now redefined slightly as nearest-neighbour distance and angle. We define the tangential shear of a galaxy relative to its nearest neighbour using the standard convention,

$$e_+ = -[e_1 \cos(\theta) + e_2 \sin(\theta)], \quad (3)$$

and the cross shear

$$e_\times = -[e_2 \cos(\theta) - e_1 \sin(\theta)]. \quad (4)$$

Note that negative values of e_+ imply a net tangential alignment of the measured shapes towards neighbours. By analogy, we define $e_{1,n}$ and $e_{2,n}$, which are the measured ellipticity components, rotated into a reference frame defined by the major axis of the neighbour. Non-zero $e_{i,n}$ would indicate leakage of the neighbour’s shape into the measurement, which might conceivably be induced by inadequate deblending of very close neighbours or by extensive non-circular masking. We first divide the main simulated catalogue into bins according to d_{gn} , and measure the tangential shear about nearest neighbours in each bin. The result is shown by the purple curve in Fig. 7. Note that the statistical uncertainty is within the width of the line in all bins. The results here show qualitative agreement with the numerical predictions in Fig. 3. As we found earlier, the exact shape of this curve is sensitive to the properties of both the neighbour and the central galaxy. Despite small differences, the range of variation is comfortably within the scale of the postage stamp for the bulk of galaxies in DES Y1. Repeating the measurement, rotated into the




Figure 4. Multiplicative bias estimated using the Monte Carlo toy model described in the text. For each neighbour realization, defined by a particular distance, flux, and size we compute the average of the measured ellipticity components over 70 rotations on a ring of neighbour angles. To estimate the bias we perform this averaging twice at two non-zero shears, g_+ and g_- , and compute the finite-difference derivative using equation (2). The dashed thin blue line shows the result of using all measurements, while the bold purple line has a cut based on the offset between the centroid position of the best-fitting model and the stamp centre.

plane of the neighbour shape results in the dotted and dot-dashed lines in this figure. As noted above, there are not necessarily reliable ellipticity measurements for each neighbour, so we instead use the sheared input ellipticities. Both components of $e_{i,n}$ are seen to be negligible over all scales.

5.2 Neighbour ensemble biases

To explore the more practical question of how neighbours impact shear estimates we divide the catalogue into bins according to neighbour distance. Within each d_{gn} bin, the galaxies are further split into 12 bins of input shear, which are fitted to estimate the multiplicative and additive bias. We show the result as the purple points in Fig. 8, which can be compared with the earlier numerical model prediction in Fig. 4. The horizontal band on these axes shows the 1σ mean m measured using all galaxies in the HOOPOE catalogue, and sits at $m \sim -0.12$. We note a steeper decline than in the bold line (without the centroid cut), more akin to the case with the centroid cut ($\Delta r_0 < 1$ arcsec). This is not surprising given that the quality selection implemented by IM3SHAPE includes exactly this cut. We do not report a local peak at ~ 11 pixels, which we saw before in Fig. 4. We suggested previously that effect was the result of positive m in galaxies where the nearest neighbour is relatively faint and at middling distance. It is likely that many of these objects manifest themselves as large changes in other quantities to which IM3SHAPE’s INFO_FLAG (see Z17) is sensitive such as ellipticity magnitude and fit likelihood, or are flagged by the SExtractor deblending cuts.

When divided into broad bins according to the r -band magnitude of the nearest neighbour $M_{r,neigh}$ (the coloured stripes in Fig. 8), we find the surviving objects show relatively weak dependence on neighbour brightness, except at the neighbour distances, where bright neighbours have a slightly stronger (negative) impact than faint ones.

We measure the additive bias components in the same bins, but find no systematic variation with d_{gn} above noise.

Finally we show the analogous measurement in bins of galaxy magnitude in Fig. 9. The steep inflation of $|m|$ at the faint end of this plot has been seen elsewhere (e.g. Fenech Conti et al. 2017; Zuntz et al. 2017), and is easily understandable as the result of noise bias. We find that splitting by neighbour magnitude does not reveal any obvious trend here.

5.3 Untangling the knot of neighbour bias

A central plank of this analysis rests on a comparison of the main HOOPOE simulations with the neighbour-free WAXWING resimulations described in Section 4.3. The simplest comparison would be between multiplicative bias values, calculated using all galaxies in each catalogue after cuts. These values are shown by the two uppermost lines (purple) in Fig. 10. The difference is an indicator of the net impact of neighbours through any mechanism, which we find to be $\Delta m \sim -0.05$.

To untangle the various contributions to this shift, we construct a matched catalogue. Galaxies in the overlap between HOOPOE and WAXWING (12M galaxies over 183 deg²) are matched by ID; quality cuts are calculated for each set of measurements (see appendix E from Z17). Geometric masking from the DES Y1 GOLD catalogue (Drlica-Wagner et al. 2017) and SExtractor deblending flags are included for HOOPOE. Since the latter are irrelevant to WAXWING, we omit them from quality flags on that data set. For conciseness we will refer to the two measurements as ‘matched HOOPOE’ and ‘matched WAXWING’, and their cuts as ‘HOOPOE cuts’ and ‘WAXWING cuts’. Since the images are identical in all respects, but for the presence of neighbours, the statistical noise on the *change* in measured quantities should be smaller than the face-value uncertainties.

The appropriate cuts are first applied to each catalogue, then the results are divided into equal number signal-to-noise bins and fitted for the multiplicative bias in each. The result is shown by the points in the upper left-hand part of Fig. 11. The equivalent in bins of PSF-normalized size is shown on the right. The difference between the blue and the purple points gives an indication of the total effect of all neighbour-induced effects on m , indicated by the solid purple line in the lower panel. The generic shift attributed to ‘neighbour bias’ is in reality a collection of distinct effects. By comparing the matched catalogues we identify four main mechanisms: direct contamination, selection bias, S/N bin shifting and neighbour dilution. Each of these components that we describe is shown by one of the lines in Fig. 11. For a visual summary of the various tests designed to isolate them see Fig. 10.

5.3.1 Direct flux contamination

The most intuitive form of neighbour bias arises from the fact that, even after masking, neighbours contribute some flux to the cutout image of a galaxy. To gauge its impact we take the common sample of galaxies, which pass cuts in both data sets. The comparison is complicated somewhat by binning in measured S/N or R_{gp}/R_p ; for this test, we divide both sets of galaxies using the WAXWING-derived quantities. The resulting m measured using the HOOPOE galaxies is unrealistic in the sense that we are binning measurements made in the presence of neighbours by quantities derived from neighbour-subtracted images. This exercise does, however, isolate the impact of the neighbour flux on the measured ellipticity. The result is shown by the purple dotted and purple dot-dashed lines in the upper and




Figure 5. *Top:* The projected footprint of the simulated survey, visualized using the SKYMAPPER package (<https://github.com/pmelchior/skymapper>). The colour indicates the local raw number density in HEALPIX cells of $n_{\text{side}} = 1024$. The axes shown are right ascension and declination in units of degrees. The full simulation comprises 1824 0.73×0.73 deg tiles drawn randomly from the DES Y1 area. The solid blue line indicates the bounds of the planned area to be covered by the complete Y5 data set. *Bottom:* A random tile (DES0246-4123) selected from the HOOPOE area. The left-hand panel (a) shows a square sub-region of approximately 4×4 arcmin. The right-hand panel (b) shows a PSF whisker plot covering the full 0.73×0.73 tile. The length and orientation of each line represents the magnitude and position angle of the spin-2 PSF ellipticity at that position. Only galaxies which pass IM3SHAPE quality cuts are shown. The white patches show the spatial masking inherited from the GOLD catalogue, and correspond to the positions of bright stars in the parent data.

lower panels of Fig. 11. The effect scales significantly with signal-to-noise and size. Faint small galaxies are affected strongly by neighbour light, while larger brighter ones are relatively immune.

5.3.2 Neighbour-induced selection bias

To gauge the neighbour-induced selection effect, we take the WAXWING catalogue but now impose, in addition to its own quality cuts, the selection function derived from the with-neighbour HOOPOE data set. The double masking removes an additional 0.5M galaxies, which survive cuts in WAXWING but would be cut from the HOOPOE catalogue. The resulting change in m is shown by the dot-dashed blue lines in the upper panels of Fig. 11 (dashed green in the lower). The multiplicative bias arising from this cut is less than

1 per cent in all but the faintest and smallest galaxies, where it can reach up to $m \sim -0.02$.

5.3.3 Bin shifting

The above two tests encapsulate the impact on the measured ellipticities, and the selection flags from neighbour flux. An additional subtlety arises from the fact that the measured quantities used to bin galaxies (S/N and R_{gp}/R_p) are themselves affected by the presence of neighbours. To test this we recalculate m using the same galaxy selection as in Section 5.3.1 (i.e. passing both sets of cuts), but now binned by the appropriate *measured* S/N . For clarity, the bin edges are unchanged, defined to contain equal numbers of WAXWING galaxies. The result is shown by the dashed lines in Fig. 11. The




Figure 6. An example of an object in the main DES Y1 calibration simulation and the neighbour-free resimulation. The upper panels show the co-add cutout in the original simulated images (left, labelled HOOPOE) and in the neighbour-subtracted version (right, labelled WAXWING). The lower panels are the segmentation masks for the same galaxy. A number of neighbours, both masked (upper left and centre left) and unmasked (lower right) are visible within the stamp bounds.




Figure 7. Tangential shear around image plane neighbours in the full HOOPOE simulation. The purple solid line shows the mean component of the measured galaxy shapes radial to the nearest image plane neighbour. Dashed blue shows the component rotated by 45° , which we have no reason to expect should be non-zero. The dotted and dot-dashed lines show the measured ellipticity components when rotated into a coordinate frame defined by the major axis of the neighbour. The inset shows the same range in d_{gn} (the x -axis tick markers are the same), but with a magnified vertical axis.

difference compared with the case using fixed binning is purely the result of galaxies moving between bins. This shifting contributes multiplicative bias if one bins galaxies by observed quantities such as S/N, as we do in order to calibrate IM3SHAPE’s shear estimates. The amplitude of this is illustrated by the blue dotted line in the lower panels.




Figure 8. Multiplicative bias as a function of separation from the nearest image plane neighbour. The purple points show the bias calculated in bins of neighbour distance using the main HOOPOE simulated shape catalogue. The coloured bands show the same data set divided into four equal-number bins according to the r -band magnitude of the neighbour. As shown in the legend, the median values in the four bins are 21.5, 23.0, 23.5, and 24.0. The mean bias and its uncertainty across all distance bins is indicated by the horizontal band.

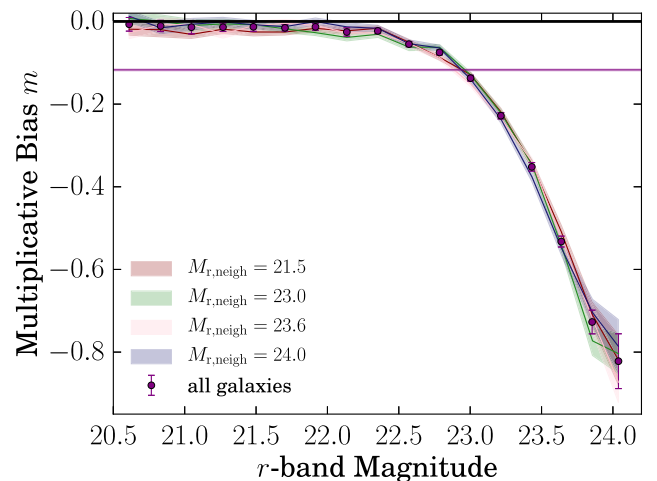


Figure 9. Multiplicative bias as a function of r -band magnitude. As in Fig. 8 the four coloured bands represent equal number bins of neighbour magnitude. Purple points show the full catalogue, with no magnitude binning. The mean bias and its uncertainty are shown by the purple horizontal band.

Such neighbour-induced shifting is noticeable if we plot out the S/N of objects in HOOPOE against the S/N of the same objects in WAXWING. Objects which are strongly shifted in S/N are more likely to scatter upwards than downwards. A similar skew can be seen in the R_{gp}/R_p plane; when galaxies are scattered in size they tend to be thrown further and more often upwards than downwards. Small galaxies (which we know already are more strongly affected by noise bias) are shifted strongly upwards by the presence of neighbour flux in the HOOPOE images. The result is a net negative m added to the upper R_{gp}/R_p bins, and a simultaneous positive shift in the lowest bins from which galaxies are lost. In the case of galaxy size we see the effects of bin scatter and direct neighbour bias almost negate each other, although the degree of cancellation




Figure 10. Graphical illustration of the measured multiplicative bias in the various scenarios considered in this paper. The upper two lines show the mean m in the main DES Y1 HOOPOE simulations and a spin-off neighbour-free resimulation named WAXWING, as described in Section 4.3. The middle section (green) shows results using only galaxies which appear in both the HOOPOE and WAXWING simulations. The matching process alone does not imply any quality-based selection function. The final three lines in red are from a similar matching between a smaller rerun of the simulation with and without sub-detection limit galaxies. See the text for details about each of these cases.

is likely dependent on the specifics of the measurement code and the data set.

5.3.4 Neighbour dilution

A final point can be gleaned from Fig. 11: that applying the WAXWING cuts to HOOPOE induces a shift in m . Naively one might expect the same galaxies as the WAXWING selection, plus some extra strongly blended galaxies. It is true that a sizeable number of galaxies are cut in the presence of neighbours, but would otherwise not be. There is also, however, a smaller population that survive cuts *because* they have image plane neighbours.

We can see this clearly from the fact that the purple points and the dashed purple lines in Fig. 11 are non-identical. We identify three separate (but partially overlapping) galaxy selections in this figure: (a) galaxies passing both sets of cuts, (b) galaxies passing cuts in the absence of neighbours, but cut by the HOOPOE selection, and (c) galaxies which pass cuts in the presence of neighbours, but cut by the WAXWING selection. We find that populations (b) and (c) have much smaller mean neighbour separation than the full population (the histograms of d_{gn} show a sharp peak at under 10 pixels). In contrast, both the full catalogue and population (a) objects a much broader distribution ($\bar{d}_{\text{gn}} \sim 24$ pixels).

Based on the toy model predictions in Section 3 we set out a working proposal: that population (c), objects cut out only when

neighbours are removed, are extreme blends dominated by a superposed neighbour. We will assume these objects are boosted considerably in size, S/N or both, such that what would otherwise be a small faint galaxy is now sufficiently bright to pass quality cuts. In these cases the measured shape of a simulated galaxy might be expected to be only weakly linked with the input ellipticity. To approximate this effect we take population (a) HOOPOE galaxies, subject to both sets of cuts, and add back some of the population (c) galaxies. Specifically, we include any objects shifted in S/N or $R_{\text{gp}}/R_{\text{p}}$ by more than 20 per cent. The true shears associated with these galaxies are now randomized to eliminate any correlation with the measured ellipticity. The result is shown as a pink dot-dashed line in Fig. 11. We can see that this effect, which we call neighbour dilution, to good approximation accounts for the residual difference between the population (a) and (c) samples. Particularly in the upper size bins of the right-hand panel the differences are not eliminated entirely. This is thought to be the result of residual (albeit weakened) covariance between the measured shapes of strongly blended objects and the input shears. Clearly the scenario in which a neighbour totally overrides the original shape of a galaxy is extreme, and there will be an indeterminate number of moderate blends which are boosted sufficiently to survive cuts but which retain some correlation with their original unblended shapes. Such cases are, however, extremely difficult to model accurately with the resources available for this investigation.




Figure 11. *Top half of each panel* Multiplicative bias as a function of signal-to-noise and size. The purple circles show the measured bias using the main HOOPOE simulation, and the blue diamonds are the resimulated neighbour-free version. The lines show permutations of the same measurements to highlight the neighbour-induced effects causing the two to differ. The dot-dashed blue and dashed purple lines show the impact of applying the HOOPOE selection mask to WAXWING and vice versa. The impact of bin shifting is shown by the purple dotted line, which is calculated from the same matched galaxies, using the HOOPOE shape measurements for the bias and WAXWING size and S/N for binning. The pink curve is the same as the dashed purple, but with a fraction of heavily blended galaxies added back with randomized shear (see Section 5.3.4). *Bottom half of each panel* The change in bias due to the effects described above. The green (dashed) line shows the impact of selection effects (the difference between the blue diamonds and the dashed line in the top panel). The direct neighbour bias due to light contamination is shown by the purple dash-dotted line (purple dotted minus blue dash-dotted top). The impact of shifting between bins is shown by the blue dotted (dashed minus purple dotted, top). The pink dot-dot-dashed line illustrates the impact of adding back randomized shears, as described. Finally the solid line represents the total neighbour bias, which includes all these effects (circles minus diamonds, top).

5.4 Redshift dependence of neighbour bias

The majority of shear measurement biases, including neighbour bias, are expected to vary systematically as a function of redshift. In simple terms, this is because the observed galaxy properties to which such biases evolve along the line of sight. Since the virtually all modern weak lensing analyses rely on tomographic binning as a way to boost the constraining power of late-time data sets, understanding such biases is particularly important.

The COSMOS galaxies used in the HOOPOE simulations are not altered relative to how they appear in the original *HST* images beyond a random rotation, shear, and PSF convolution. We thus expect the redshift dependence of morphology, shape, and brightness to be accurately represented in our simulations. Dividing galaxies into tomographic bins using their COSMOS photo- z estimates (e.g. Z17) should then capture the first-order impact of such redshift dependence.

Insofar as our simulations are realistic, we can measure such redshift dependence directly. To do this we first distribute the ~ 17 M HOOPOE galaxies amongst four redshift bins using the procedure described in section 5.3 of Z17, which is designed to roughly mimic the smearing in redshift due to photo- z error. The panels in Fig. 12 show the distributions of a number of salient galaxy properties, measured in each of the four DES Y1 redshift bins (left to right).

In each redshift bin we measure m first using HOOPOE, and then using the same galaxies in WAXWING. Since the populations are identical barring the presence of neighbour light, the difference between the two should capture the net impact of blending in each redshift slice.³ The result is a monotonic increase in m between ~ -0.035

and more than -0.06 in the upper redshift bin. This is shown by the solid purple points in Fig. 13. We can isolate the impact of neighbour-induced selection bias by comparing the m values computed from the neighbour-free WAXWING, under selection functions derived from the two different simulations. The difference is shown by the open blue circles in Fig. 13, and accounts for percentage-level biases with no strong redshift dependence.

One additional complication is that galaxy clustering also evolves as a function of redshift, and neighbour bias is a function of local number density. It is not inconceivable that this might introduce a redshift dependence bias term. For this to have a significant impact one must, of course, assume that in a given patch of sky the source density is dominated by localized clusters of galaxies, and not field galaxies at independent random redshifts. In the HOOPOE simulations we select an independent random COSMOS profile at each detection position, and so do not include such spatial correlations. If one compares the distributions of d_{gn} as measured from the simulations and the DES Y1 data (top row of Fig. 12) one finds low-level disagreement but no clear trend with redshift.

Finally, we can test the impact of such disagreement by reweighting the simulations to match the data and recomputing the bias. Unfortunately, a naive reweighting in one parameter (d_{gn}) risks inducing changes in the histograms of correlated properties such as ellipticity and size. We thus perform the following exercise. We repeat the fiducial shear calibration described in Z17 twice to obtain a multiplicative correction m_i for each HOOPOE galaxy. Each time we reweight the simulated galaxies within each grid cell of S/N and R_{gp}/R_p ; first reweighting such that the joint $p(e, d_{\text{gn}})$ distribution matches the real Y1 data, and then only for $p(e)$. We apply each of these weighted calibrations to the simulation and remeasure the residual bias in four tomographic bins. We find the shift in m due to additionally reweighting to account for the small differences in the histograms in the top row of Fig. 12 is negligible compared with statistical error at $\mathcal{O}(10^{-6})$.

³ Although this neglects the impact of neighbours on photometric redshift error. Though this topic is important, it is beyond the scope of the present study.




Figure 12. Measured histograms of observable properties in the main HOOPOE and neighbour-free WAXWING simulations. The rows show (top to bottom) nearest neighbour distance, ellipticity, signal-to-noise ratio, and PSF normalized size. Each column (left to right) shows the distribution of galaxies in one of the four tomographic bins used in the DES Y1 cosmic shear analysis.

5.5 Isolating the impact of sub-detection galaxies

A handful of previous studies have attempted to quantify the impact of galaxies below, or close to, a survey’s limiting magnitude. For example, Hoekstra et al. (2015, 2017) suggest they can induce a non-trivial multiplicative bias, which is dependent on the exact detection limit. They recommend using a shear calibration sample at least by 1.5 mag deeper than the survey in question (which ours does). Their findings, however, make exclusive use of the moment-based KSB algorithm (see Kaiser et al. 1995); such techniques are known to probe a galaxy’s ellipticity at larger radii than other methods, which could in principle make them more sensitive to nearby faint galaxies. It is thus a worthwhile exercise to gauge their impact in our case with IM3SHAPE.

5.5.1 Impact on multiplicative bias

We first compare our HOOPOE simulations with the neighbour-free WAXWING resimulations. Since WAXWING postage stamps consist of

only a single profile added to Gaussian pixel noise, they are unaffected by neighbours of any sort (faint or otherwise). We have seen that the impact of neighbours is strongly localized, with the excess m converging within a nearest neighbour distance d_{gn} of a dozen pixels or so. Thus selecting galaxies that are well separated from their nearest *visible* neighbour will isolate the impact of the *undetected* ones.

A further cut is thus imposed on $d_{gn} < 20$ pixels. Relative to the case with quality cuts only, the global multiplicative bias now shifts from $m \sim -0.119$ to $m = -0.064 \pm 0.006$ (the first and second lines in green in Fig. 10). This measurement is in mild tension with the value measured from WAXWING (again under its own cuts, with the selection on d_{gn}). This difference, which amounts to a negative shift in m of ~ 0.01 is, we suggest, the net effect of the sub-detection galaxies. From these numbers alone we cannot tell if this is a result of selection effects, flux contamination, bin shifting, or some combination thereof.

Interestingly we find that imposing both the HOOPOE and WAXWING selection functions in addition to the cut on d_{gn} brings m into

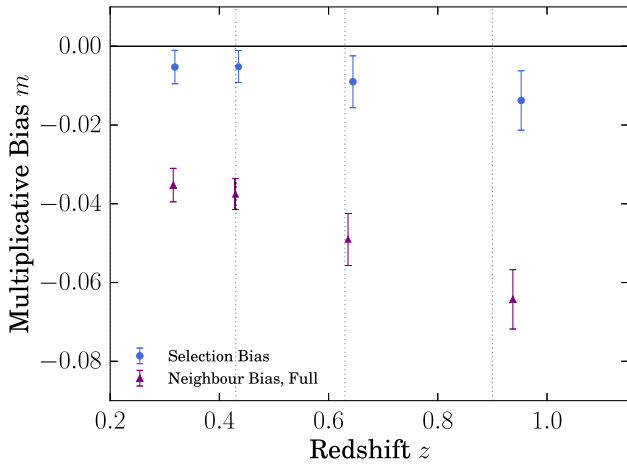


Figure 13. Measured neighbour-induced multiplicative bias in the four redshift bins of Troxel et al. (2017). Each simulated galaxies is allocated randomly to one of four bins, such that the distribution of COSMOS redshifts in each bin roughly matches the photometric redshift distribution as estimated from the real data. The solid purple triangles show the impact of neighbour bias in each bin, computed as the net difference between two simulations with and without blending. The open circles represent the selection bias component, as calculated by comparing results from the no-neighbours simulation using different selection functions.

agreement to well within the level of statistical precision (compare the final and penultimate lines in green in Fig. 10). That is, when restricted to a subset of galaxies that pass quality cuts in both simulations the flux contributed by the faint objects has little impact. Their main impact is rather that they allow a population of marginal faint galaxies which would otherwise be flagged and removed by quality cuts to pass into the final HOOPOE catalogue.

To test this idea further we rerun a subset of 100 random tiles from the simulated footprint, without the final step of adding sub-detection galaxies. To minimize the statistical noise in this comparison we enforce the same COSMOS profiles, shears, and rotations as well as the per-pixel noise realization as before. SExtractor source detection is applied and the blending flags are propagated into the post-processing cuts.

The raw m values calculated from the rerun and the main HOOPOE simulations, matched to the same tiles, are shown by the upper two red lines in Fig. 10. The downward shift of ~ 0.01 is consistent with the previous result based on the main simulation. This comparison should encapsulate the full effect of the faint objects (since there are no other differences between these data sets).

For each galaxy we next measure the distance to the nearest faint object d_{gf} , the distribution of which is shown under various selections in Fig. 14. Like in the comparison in Section 5.3.4, there is a population of galaxies that survive cuts *only* in the simulation with the sub-detection objects, and these galaxies tend to be ones with extremely close faint neighbours. Interestingly the inverse population surviving only when they are removed do not preferentially have small d_{gf} . This is intuitively understandable: a faint object might boost its neighbour’s apparent size or S/N if it were centred within a few pixels. Otherwise it would act as a source of background noise, which would reduce the quality of the fit.

Finally we find that if we apply both selection functions to the with-faint galaxies, the measured biases become roughly consistent. These findings, combined with the observations in the previous section, lead us to an interesting conclusion: the major effect of the faint galaxies in the DES Y1 IM3SHAPE catalogue is to allow




Figure 14. Histogram of radial distances between galaxies in our measured shape catalogues (the full HOOPOE simulations) and the nearest object below the DES detection limit. The dotted line includes all objects prior to quality cuts, while the solid line shows the impact of applying IM3SHAPE’s INFO_FLAG cuts (see J16). The dashed blue line shows the population of galaxies which survive cuts only in the presence of the faint galaxies.

a population of small faint galaxies to pass quality cuts, where otherwise they would have been removed. This is analogous to the neighbour dilution effect described above, but is subdominant to the influence of visible neighbours.

5.5.2 Impact on background flux subtraction

As a test of the robustness of this result we recompute our IM3SHAPE fits on the faint-free images, with and without the correction for the shift in the background flux that *would* have been applied if the sub-detection galaxies had been drawn. The mean per-tile correction is $\Delta f \sim 0.05$, against typical noise fluctuations $\sigma_n \sim 6.5$. Matching galaxies and examining the histograms of $\Delta S/R$ and $\Delta R_{gp}/R_p$ reveals weak downwards scatter in both quantities (i.e. the flux subtraction alone makes galaxies appear smaller and fainter). The magnitude of the shift is, however, tiny, peaking at ~ -0.1 and -0.005 , respectively. This is logical given the definition in equation (1). If the change is small enough such that the best-fitting model is stable, then an incremental reduction in flux will reduce the signal-to-noise of the measurement. Looking at the best-fitting shapes, we find a small shift towards high ellipticities, which can likewise be understood as a numerical effect; imposing a flat positive field of zero ellipticity will dilute the measured shear, producing a bias towards round $|e|$. The reverse logic applies with the flux correction, and subtracting a flat value from all pixels will make galaxies appear slightly *more* elliptical. In practice we find a sharp peak at $\Delta e \sim 0.001$.

5.6 Suppressing neighbour bias

There is no universal definition for the shape-weighted effective number density commonly used as proxy for cosmological constraining power in a shear catalogue. One which is particularly useful in the context of weak lensing, and which has been adopted in DES Y1 is the prescription of Chang et al. (2013), which is designed to account for shape noise and fitting error (see equation 7.5 in Zuntz et al. 2017). A second useful definition is set out by Heymans et al. (2012) in terms of the intrinsic shape dispersion and measurement error (see also Z17). We compute a neighbour

distance d_{gn} for every object in the real data, which allows us to cut on this quantity. Removing any galaxy with a neighbour detected within a radius of $d_{\text{gn}} = 20$ pixels reduces the effective number density of sources using either definition to about 70 per cent of its initial value, from $n_{\text{eff}}^{\text{H13}} = 5.48$ to $n_{\text{eff}}^{\text{H13}} = 3.68 \text{ arcmin}^{-2}$ using Heymans et al. (2012)’s definition. Using the prescription of Chang et al. (2013), the equivalent density drops from $n_{\text{eff}}^{\text{C13}} = 3.18$ prior to cuts and $n_{\text{eff}}^{\text{C13}} = 2.18 \text{ arcmin}^{-2}$ afterwards. This cut is stringent, as we have shown that beyond ~ 12 pixels the multiplicative bias becomes insensitive to further selection on d_{gn} . There are, however, a number of limitations in our analysis, including the fact that d_{gn} is defined using the true input positions, and indeed that we are using only the detected positions in DES to draw our simulated $M_r < 24.1$ galaxies. We thus judge that a level of conservatism is appropriate here. Relaxing the cut to $d_{\text{gn}} > 14$ pixels leaves n_{eff} at 84 per cent of its full value.

6 COSMOLOGICAL IMPLICATIONS

As we have shown in the previous sections, if ignored completely image plane neighbours can induce negative calibration biases in IM3SHAPE of a few per cent or more. The earlier part of the investigation focused on *when* and *how* neighbour bias can arise, first in the context of single-galaxies and then on ensemble shear estimates. We now turn to a more pressing question from the general cosmologist’s perspective: *how far should I be concerned about these effects in practice?* We present a set of numerical forecasts using the MULTINEST nested sampling algorithm (Feroz et al. 2013) to sample trial cosmologies. Each of the likelihood analyses presented in this paper has been repeated using a Markov Chain Monte Carlo sampler (EMCEE). Although we see the same small shift in contour size noted by Dark Energy Survey Collaboration (2017) (see their appendix A), which diminishes as the length of the MCMC chains increases, we find our conclusions are robust to the choice of sampler. Our basic methodology here follows previous numerical forecasts (e.g. Joachimi & Bridle 2010; Krause & Eifler 2017; Krause et al. 2017). We construct mock DES Y1 cosmic shear measurements using a matter power spectrum derived from the Boltzmann code CAMB ⁴ with late-time modifications from HALOFIT . The cosmic shear likelihood surface is sampled at trial cosmologies using COSMOSIS .⁵ The final data used for the likelihood calculation have the form of real-space ξ_{\pm} correlations. For the photometric redshift distributions we use the measured estimates in four tomographic bins, obtained from runs of the BPZ code on the Y1 IM3SHAPE catalogue, as described by Hoyle et al. (2017). Since this analysis was completed before the details of the photometric redshift calculation for DES Y1 had been finalized, these distributions differ marginally from (but are qualitatively the same as) the final version used in Troxel et al. (2017) and Dark Energy Survey Collaboration (2017). In all chains which follow we marginalize over two nuisance parameters (an amplitude and a power law in redshift) for intrinsic alignments, photo- z bias, and shear calibration bias. In total this gives 10 extra free parameters in addition to six for cosmology ($\Omega_m, \Omega_b, A_s, n_s, h, \Omega_v h^2$), which are also allowed to vary. Apart from the difference in redshift distributions remarked upon above, our analysis choices match the DES Y1 cosmic shear analysis of Troxel et al. (2017). We refer the reader to that paper for details of the priors and scale cuts, and their derivation. Finally, the following adopts shear–shear

covariance matrices derived from the analytic halo model calculations of Dark Energy Survey Collaboration (2017). We assume a fiducial ΛCDM cosmology $\sigma_8 = 0.83, \Omega_m = 0.295, \Omega_b = 0.047, n_s = 0.97, h = 0.688, \tau = 0.08$, with non-zero comoving neutrino density $\Omega_\nu h^2 = 0.00062$.

6.1 Mean multiplicative bias

We first seek to quantify the bias that would be present in a cosmic shear analysis in a survey like DES, if we were to use a simple postage stamp simulation of the sort presented in J16 and Miller et al. (2013). To this end we use the neighbour-free WAXWING data set to construct an alternative shear calibration. In Z17 we compared three methods for shear calibration using the HOOPOE simulations and found our results to be robust to the differences. We now use the fiducial (grid-based) scheme to derive an alternative set of bias corrections from WAXWING . These are then applied to the same galaxies in the matched HOOPOE simulation, and residual biases are measured in four DES-like tomographic bins. The process is very similar to the diagnostic tests in section 5 of Z17, and so we defer to that work for details of the redshift bin assignment of simulated galaxies.

Using the neighbour-free simulation we undercorrect the measurement bias by several per cent in each bin. The remeasured residual bias after calibration provides an estimate for the level of systematic that would be present were we to calibrate DES Y1 using the simpler WAXWING simulations. In the four tomographic bins used in DES Y1 we find $(\Delta m^{(1)}, \Delta m^{(2)}, \Delta m^{(3)}, \Delta m^{(4)}) = (-0.037, -0.044, -0.064, -0.073)$, and apply these biases to our mock data. The resulting shift in the shear two-point correlations is shown by the black dotted lines in the lower panel of Fig. 15. Since the calibration scheme does not explicitly include neighbour distances, but rather orders galaxies into cells of S/N and R_{gp}/R_p , this test does not include any scale-dependent neighbour effects. The calibration effectively marginalizes out d_{gn} , and the residual biases are an average over the survey. For the moment we will assume this mean shift in m is sufficient, and return to the question of scale dependence in the following section.

Our predicted cosmology constraints with weak lensing alone in DES Y1 are shown in Fig. 16. In purple we show the results of the fiducial analysis, in which the shear calibration fully captures all neighbour effects and leaves no residual multiplicative bias. The blue (solid) contours then show the impact of residual neighbour biases per bin at the level described. As we can see, even when marginalizing over m' with an (erroneously) zero-centred Gaussian prior of width $\sigma_m = 0.035$, our cosmology constraints are shifted enough to place the input cosmology outside the 1σ confidence bounds. We reiterate here that this calculation highlights the bias that *would* arise were we to naively apply a calibration of the sort used in DES SV based on neighbour-free simulations to the Y1 data. Since we are confident that the HOOPOE code captures the effects of image plane neighbours correctly (at least to first order) this is a hypothetical scenario only and not a prediction of actual bias in DES Y1. Likewise, since the recent shear analyses from KiDS and DLS also incorporate blending in their calibration simulations, we do not expect this level of neighbour-induced bias to be present in their cosmology results.

6.2 Scale dependence

It is not trivial that including a mean neighbour-induced component to m over the entire survey will be sufficient to mitigate all forms

⁴ <http://camb.info>

⁵ <https://bitbucket.org/joezuntz/cosmosis>

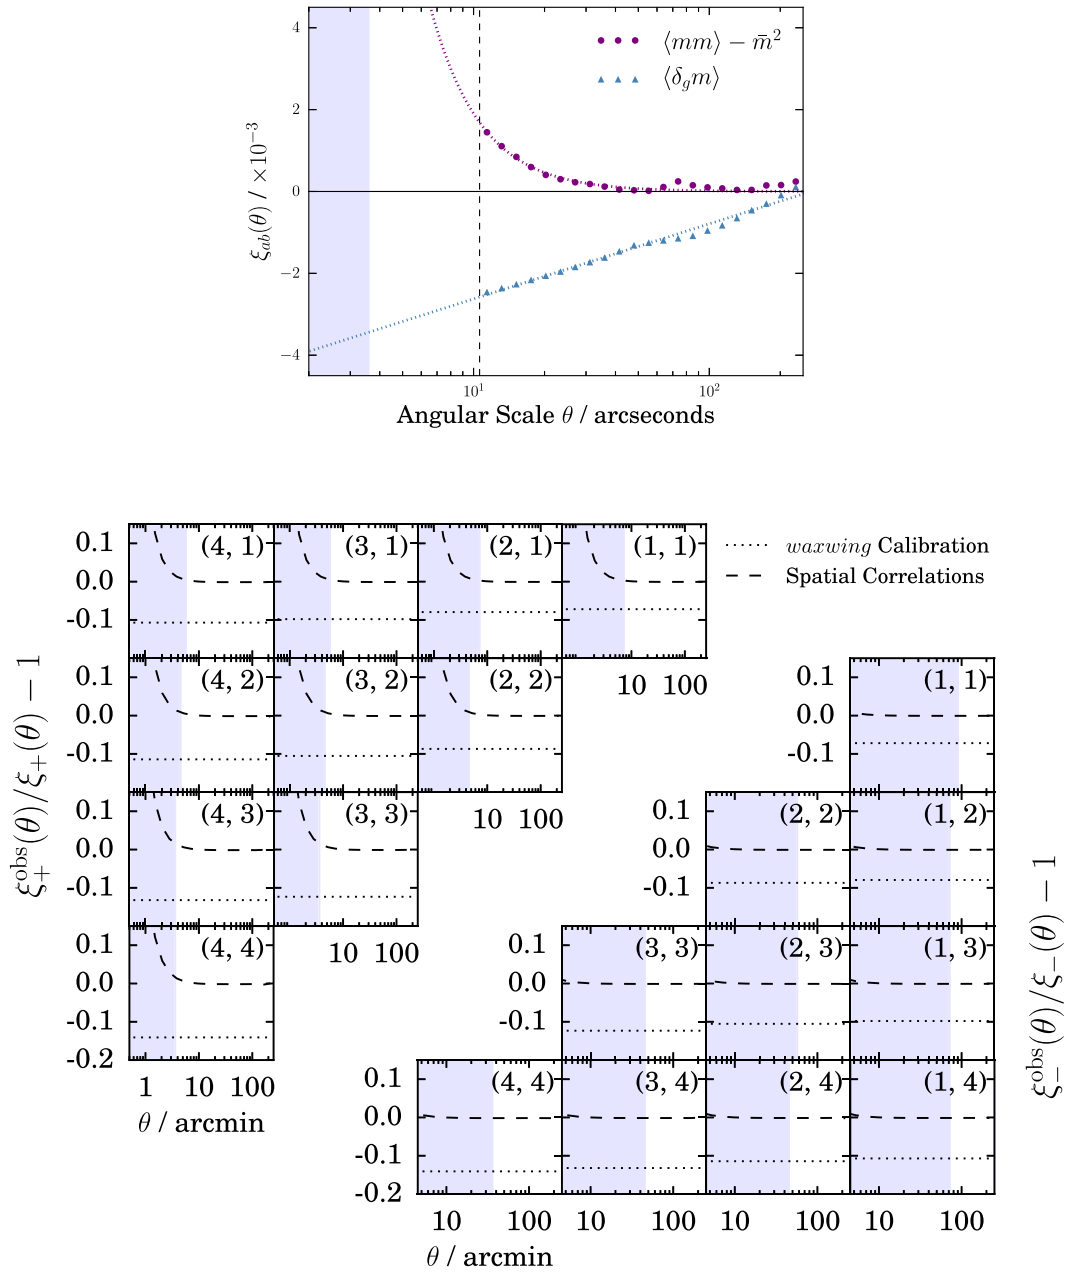


Figure 15. *Top:* The observed two-point correlation of multiplicative bias, as measured from the main HOPOE simulation set presented in this paper. Sub-patches are used to compute m in spatial patches of dimension 0.15×0.15 deg and the correlation function calculated as described in the text. The dashed vertical line shows the diagonal scale of the sub-patches, below which we do not attempt to directly measure spatial correlations. The shaded blue bands show the minimum and maximum scales used in the DES Y1 cosmic shear analysis of Troxel et al. (2017). *Bottom:* Residuals between the mock two-point shear–shear data used in this paper, before and after different forms of bias have been applied. The upper-left and lower-right triangles show the ξ_+ and ξ_- correlations, respectively, calculated using the redshift distributions of Hoyle et al. (2017). The dotted black lines, which are flat across all scales but vary between panels, show the result of calibrating our Y1 shear measurements with a simple postage stamp simulation without image plane neighbours. The dashed lines illustrate the impact of ignoring scale-dependent selection effects, which are not captured by our simulation-based calibration. The shaded blue regions of each panel show the excluded scales for each particular tomographic bin pairing.

of neighbour bias. The local mean m on a patch of sky is sensitive to spatial fluctuations in source density, which could induce scale-dependent bias on arcminute scales. Clearly, when correlating galaxy pairs on small scales one can expect a larger fraction in which the objects come from a similar image plane environment, and more often than not that environment will be densely populated. Thus the true multiplicative bias should become more negative on small scales.

Two subtly different effects emerge from this thought experiment. First, the multiplicative bias of galaxies will be spatially correlated, i.e. a correlation involving two galaxy populations ($m^i m^j$) is not just the product of the means $\bar{m}^i \bar{m}^j$. Secondly, in the small θ bins one is selecting galaxies with close partners with which to correlate, and thus oversampling the dense parts of the image. To gauge the level of these effects, we divide each simulated co-add tile into a grid of 25 square sub-patches with dimension 0.15×0.15 deg. We fit for m




Figure 16. Expected cosmology constraints from DES Y1 cosmic shear only. The purple (solid) contours show the results of calibrating using a simulation which fully encapsulates all biases in the data, leaving no residual m in the final catalogue. In blue (dash-dotted) we show the result of calibrating with an insufficiently realistic simulation, which leaves a residual bias between -0.03 and -0.08 in each of the redshift bins. For reference we mark the input cosmology with a black cross.

using the galaxies in each sub-patch and assign the resulting value to these objects. While this only allows a noisy measurement of m , it should capture the spatial variations in number density to the level of a few per cent. We next measure the two-point correlation function of multiplicative bias values assigned in this way using TREECORR,⁶ excluding galaxy pairs at angular separation smaller than the scale of the sub-patches. We refer to this bias–bias autocorrelation as ξ_{mm} , which we show as a function of angular scale in the upper panel of Fig. 15. analogously one could use the sub-patches to construct correlations between m and galaxy number density ξ_{gm} or density with density ξ_{gg} . The statistical noise on these correlations is significantly lower than that on the individual sub-patches by virtue of the large simulation footprint. Note that in Fig. 15 we subtract \bar{m}^2 , measured from all galaxies in the simulation, from the measured ξ_{mm} . If there were no θ dependence the correlation $\langle m^i m^j \rangle$ should simply average to the square of the global mean in all scale bins. As we can see from the circular points in this figure, scales larger than the diagonal size the sub-patches (shown by the vertical dashed line) exhibit non-negligible excess ξ_{mm} . One obvious question is whether this could be the result of finite binning error, which scatters galaxies in the same sub-patch into different θ bins. To verify this is not the case we repeat the measurement as before, but halve the parameter controlling binning error tolerance (‘bin slop’) and obtain the same results.

To extend this measurement down to scales below the sub-patch size we must make some assumptions about the functional form of the mm correlation. We fit a power law, $\Delta\xi_{mm}(\theta) \equiv \xi_{mm}(\theta) - \bar{m}^2$:

$$\Delta\xi_{mm}(\theta) = \beta\theta^{-\alpha}, \quad (5)$$

which is shown by the dotted purple line in this figure. This provides a qualitatively good fit to the measured points, but as we can see implies a rather dramatic inflation on small scales.

In the limited range over which we have a non-zero measured correlation, however, a linear function of θ (truncated at




Figure 17. The same as Fig. 16, but now showing the impact of residual scale-dependent selection bias. The two sets of confidence contours represent different assumptions about the small-scale extrapolation of the ξ_{mm} correlation, as outlined in the Section 6.2. In green (dashed) we show a mildly optimistic case, using the linear fit shown in Fig. 15. The pink dotted contours show a (strongly pessimistic) power-law extrapolation. The dark blue solid line makes identical assumptions to the pink, but incorporates small-scale information, to a minimum separation of $\theta_{\min}^+ = 0.5$ arcmin in $\xi_+(\theta)$ and $\theta_{\min}^- = 4.2$ arcmin in $\xi_-(\theta)$. As in Fig. 16 the input cosmology is shown by a black cross.

$\theta = 27$ arcmin) also provides a reasonable by-eye fit. The small-scale extrapolation in this case is significantly milder. The $\langle m\delta_g \rangle$ and $\langle \delta_g \delta_g \rangle$ measurements are linear with θ to good approximation, and so we use linear fits to extrapolate them below the patch size.

Assuming the bias per tomographic bin can be written as the sum of a redshift-dependent contribution (i.e. a scale invariant mean in each bin), and a scale-dependent term, one can write the correlation per bin as $m^i m^j = \bar{m}^i \bar{m}^j + \Delta\xi_{mm}(\theta)$. A more complete derivation of this expression can be found in Appendix A. The first part can be extracted from the DES Y1 calibration, and we can fit for $\Delta\xi_{mm}(\theta)$ as described above. A set of modified ξ_{\pm}^{ij} are thus computed. These appear in the lower panels of Fig. 15 as dashed lines. As we can see, the scale cuts of Troxel et al. (2017) (excluded scales are shaded in blue) are sufficiently stringent to remove almost all of the visible scale dependence. Though reassuring for the immediate prospects of DES Y1, this will not trivially be true for all future (or indeed ongoing) lensing surveys. It is thus important that the effects we identify here are properly understood at a level beyond the resources of the current paper. These biased data are then passed through our likelihood pipeline to gauge the cosmological impact, which is shown in Fig. 17. In the linear case (dashed green) there is no discernable bias in the σ_8 Ω_m pair; even the much harsher power-law extrapolation (pink dotted) induces only an incremental shift along the degeneracy direction. In both cases the input cosmology still sits comfortably within the 1σ confidence contour.

Finally we test the impact of relaxing the stringency of our scale cuts. The minimum scales used for ξ_+^{ij} and ξ_-^{ij} are shifted downwards to 0.5 and 4.2 arcmin, respectively, irrespective of bin pair, which are the cut-off values used in fiducial cosmic shear analysis of Hildebrandt et al. (2017). This increases the size of our datavectors considerably. Incorporating smaller angular scales will clearly improve the constraining power of the data to an extent. Primarily the effect is to shorten the lensing degeneracy ellipse, cutting out much of the peripheral curvature, but it also reduces the width in the S_8 direction. These scales, however, contain biased information, which

⁶ <http://rmjarvis.github.io/TreeCorr>

induces tension between the small and large angular scales. With the strongest (power law) scale dependence considered, the input cosmology is displaced marginally along the degeneracy curve, though it remains comfortably within the 1σ confidence bound.

7 CONCLUSIONS

The Dark Energy Survey represents the current state of the art in cosmological weak lensing. Multiband imaging down to 24th magnitude across 1500 deg² of the southern sky has yielded hitherto unparalleled late-time constraints on the basic parameters of the Universe (see Dark Energy Survey Collaboration 2017; Troxel et al. 2017).

In this paper we have used one of two DES Y1 shear catalogues, and large-area simulations based upon them, to quantify the impact of image plane neighbours on both ensemble shear measurements, and on the inferred cosmological parameters.

In order to properly mitigate the influence of galaxy neighbours, and thus avoid drawing flawed conclusions about cosmology from the data, it is important to first understand the mechanisms by which they enter the shape measurement. Using a simple toy model of the galaxy-neighbour system we have shown that shear bias can arise even when the distribution of neighbours is isotropic (i.e. there is no preferred direction). This is the result of a small difference in the impact of the same neighbour, when it is placed on or away from the axis of the shear. We have furthermore shown that the resulting multiplicative shear bias m can be either positive or negative, depending on the model parameters. With slight modifications to the toy model, whereby we Monte Carlo sample input parameters from the joint distribution of the equivalent properties measured in DES Y1, we have shown that a mild negative m is dominant when marginalizing over a realistic ensemble of neighbours. This was seen to be strongly dependent on the distance of the neighbour, and to be mitigated but not eliminated by basic cuts on the centroid position of the best-fitting model.

Using the DES Y1 `HOPOE` simulations, which were also used to derive shear calibration corrections for the Y1 `IM3SHAPE` catalogue of Z17, we have presented a detailed study of the ensemble effects of galaxy neighbours. In this analysis we have identified four mechanisms for neighbour bias, which we call flux contamination, selection effects, bin shifting, and neighbour dilution. All can be understood in intuitive terms, resulting from close-by or moderately close neighbours. Our results from the full simulation are consistent with the toy model calculation. Though we have shown strong dependence on distance to the nearest neighbour (and thus on number density) we found only weak sensitivity to neighbour brightness, when averaged across broad bins of magnitude. In addition to this, cuts on the DES Y1 catalogue sufficient to null the impact of the detectable neighbours would result in a degradation of over 20 per cent in source number density. We cannot recommend such measures for a code like `IM3SHAPE`, in part because the data contain correlations between shear and number density. Unless the link is preserved in the calibration simulations, such selection could conceivably induce additional bias towards low shear.⁷

Our investigation also assessed the impact of the faintest galaxies, which are not reliably detected but none the less contribute flux to the survey images. Via two different routes, first using a spin-off

neighbour-free resimulation, and also using a subset of images simulated again with the sub-detection galaxies missing, our findings suggested a net contribution to the multiplicative bias budget of $m \sim -0.01$.

Unlike most earlier works on shear measurement, we have propagated these findings to the most meaningful metric for cosmic shear: bias on the inferred cosmological parameters. The study we have presented here uses the DES Y1 cosmology pipeline, as well as real non-Gaussian shear covariance matrices and photometric redshift distributions to implement MCMC forecasts. In the first case considered, the data included a (different) multiplicative bias in each redshift bin, designed to approximate the residual m that would arise were we to calibrate DES Y1 with a simple neighbour-free simulation. Even marginalizing over m with a prior of $\mathcal{N}(0, 0.035)$ this scenario was demonstrated to result in a shift in the favoured cosmology towards low clustering amplitude of more than 1σ .

Finally, we have explored a second possible source of measurement bias arising from the link between number density and neighbour bias. This enters two-point measurements as an additional correlation between the multiplicative bias in galaxy pairs at small angular separation. In the final section we have measured such a correlation from the `HOPOE` mock images. With the most pessimistic small-scale extrapolation, this was found to result in a shift in the best-fitting cosmology of under 1σ in the negative S_8 direction, which is not remedied by marginalizing over m . A less dramatic, though still considerable, increase in the correlation strength on small scales was demonstrated to result in no discernable cosmological bias.

Both of these effects are of primary concern for the next generation of cosmological surveys. By the end of their lifetime KiDS, DES, and HSC are set to offer lensing-based cosmological constraints comparable to the CMB. The first, dominant, effect can be remedied relatively easily by calibrating our shear measurements with sufficiently complex image simulations. Indeed, the most recent shear constraints of Hildebrandt et al. (2017), Köhlinger et al. (2017) and Troxel et al. (2017) have done just that. Unfortunately, the correct treatment of scale-dependent bias is not as clear, though it should be captured at some level by the per-galaxy responses upon which `METACALIBRATION` relies. Though further statements about the likely small scale dependence of the mm correlation are beyond the scope of the present study, understanding this intricate topic will be crucial for future surveys if we are to fully exploit the constraining power of the data. The massive simulation efforts of LSST and Euclid, combined with advancement in neighbour mitigation using techniques such as multi-object fitting will be invaluable in this task. With the enhanced understanding these will provide and the exquisite data of the next generation surveys, the coming decade will be an exciting time for cosmology.

ACKNOWLEDGEMENTS

We thank Nicolas Tessore, Catherine Heymans, and Rachel Mandelbaum for various insights that contributed to this work. We are also indebted to the many DES ‘eyeballers’ for lending their holiday time to help us understand and validate our simulations. The `HOPOE` simulations were generated using the National Energy Research Scientific Computing Center (NERSC) facility, which is maintained by the U.S. Department of Energy. The likelihood calculations were performed using NERSC and the Fornax computing cluster, which was funded by the European Research Council. SLB acknowledges support from the European Research Council in the form of a Consolidator Grant with number 681431. Support for DG was provided

⁷ Although the sister catalogue to Y1 `IM3SHAPE` uses a form of internal calibration, which should allow one to correct for the additional selection bias.

by NASA through Einstein Postdoctoral Fellowship grant number PF5-160138 awarded by the *Chandra* X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for NASA under contract NAS8-03060.

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 University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, Texas A&M University, and the OzDES Membership Consortium.

The DES data management system is supported by the National Science Foundation under Grant Numbers AST-1138766 and AST-1536171. The DES participants from Spanish institutions are partially supported by MINECO under grants AYA2015-71825, ESP2015-88861, FPA2015-68048, SEV-2012-0234, SEV-2016-0597, and MDM-2015-0509, some of which include ERDF funds from the European Union. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya. Research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) including ERC grant agreements 240672, 291329, and 306478. We acknowledge support from the Australian Research Council Centre of Excellence for All-sky Astrophysics (CAASTRO), through project number CE110001020.

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

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 As-

tronomy (AURA) under a cooperative agreement with the National Science Foundation.

REFERENCES

- Albrecht A. et al., 2006, preprint ([arXiv:astro-ph/0609591](https://arxiv.org/abs/astro-ph/0609591))
- Bacon D. J., Refregier A. R., Ellis R. S., 2000, MNRAS, 318, 625
- Bernstein G. M., Armstrong R., 2014, MNRAS, 438, 1880
- Bernstein G. M., Jarvis M., 2002, AJ, 123, 583
- Bertin E., Fouqué P., 2010, Astrophysics Source Code Library, record ascl:1010.066
- Bridle S. L., Kneib J.-P., Bardeau S., Gull S. F., 2002, in Natarajan P., ed., Proc. Yale Cosmology Workshop, The Shapes of Galaxies and their Dark Halos. p. 38
- Bridle S. et al., 2009, Ann. Appl. Stat., 3, 6
- Brouwer M. M. et al., 2017, MNRAS, 466, 2547
- Chang C. et al., 2013, MNRAS, 434, 2121
- Chang C. et al., 2015, Phys. Rev. Lett., 115, 051301
- Chang C. et al., 2017, preprint ([arXiv:1708.01535](https://arxiv.org/abs/1708.01535))
- Clampitt J. et al., 2017, MNRAS, 465, 4204
- Dark Energy Survey Collaboration, 2017, preprint ([arXiv:1708.01530](https://arxiv.org/abs/1708.01530))
- Dark Energy Survey Collaboration, 2016, Phys. Rev. D, 94, 022001
- Diehl H. T. et al., 2014, in Peck A. B., Benn C. R., Seaman R. L., eds, Proc. SPIEConf. Ser. Vol. 9149, Observatory Operations: Strategies, Processes, and Systems V. SPIE, Bellingham, p. 91490V
- Drlica-Wagner A. et al., 2017, ApJS, preprint ([arXiv:1708.01531](https://arxiv.org/abs/1708.01531))
- Fenech Conti I., Herbonnet R., Hoekstra H., Merten J., Miller L., Viola M., 2017, MNRAS, 467, 1627
- Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2013, preprint ([arXiv:1306.2144](https://arxiv.org/abs/1306.2144))
- Flaugher B. et al., 2015, AJ, 150, 150
- Harnois-Déraps J., Munshi D., Valageas P., van Waerbeke L., Brax P., Coles P., Rizzo L., 2015, MNRAS, 454, 2722
- Hartlap J., Hilbert S., Schneider P., Hildebrandt H., 2011, A&A, 528, A51
- Herbonnet R., Buddendiek A., Kuijken K., 2017, A&A, 599, A73
- Heymans C. et al., 2006, MNRAS, 368, 1323
- Heymans C. et al., 2012, MNRAS, 427, 146
- Heymans C. et al., 2013, MNRAS, 432, 2433
- Hildebrandt H. et al., 2017, MNRAS, 465, 1454
- Hirata C., Seljak U., 2003, MNRAS, 343, 459
- Hirata C. M., Seljak U., 2004, Phys. Rev. D, 70, 063526
- Hoekstra H., Herbonnet R., Muzzin A., Babul A., Mahdavi A., Viola M., Cacciato M., 2015, MNRAS, 449, 685
- Hoekstra H., Viola M., Herbonnet R., 2017, MNRAS, 468, 3295
- Hoyle B. et al., 2017, MNRAS, preprint ([arXiv:1708.01532](https://arxiv.org/abs/1708.01532))
- Huff E., Mandelbaum R., 2017, preprint ([arXiv:1702.02600](https://arxiv.org/abs/1702.02600))
- Jarvis M., Jain B., Bernstein G., Dolney D., 2006, ApJ, 644, 71
- Jarvis M. et al., 2016, MNRAS, 460, 2245 (J16)
- Jee M. J., Tyson J. A., Hilbert S., Schneider M. D., Schmidt S., Wittman D., 2016, ApJ, 824, 77
- Joachimi B., Bridle S. L., 2010, A&A, 523, A1
- Kacprzak T., Zuntz J., Rowe B., Bridle S., Refregier A., Amara A., Voigt L., Hirsch M., 2012, MNRAS, 427, 2711
- Kacprzak T., Bridle S., Rowe B., Voigt L., Zuntz J., Hirsch M., MacCrann N., 2014, MNRAS, 441, 2528
- Kacprzak T. et al., 2016, MNRAS, 463, 3653
- Kaiser N., 1994, in Durret F., Mazure A., Tran Thanh Van J., eds, Proc. XXIXth Recontre De Moriond, XiVth Moriond Astrophysics Meetings, Clusters of Galaxies. p. 269
- Kaiser N., Squires G., Broadhurst T., 1995, ApJ, 449, 460
- Kaiser N., Wilson G., Luppino G. A., 2000, eprint ([arXiv:e-prints](https://arxiv.org/abs/e-prints))
- Kilbinger M. et al., 2013, MNRAS, 430, 2200
- Kitching T. et al., 2011, Ann. Appl. Stat., 5, 2231
- Köhlinger F. et al., 2017, MNRAS, 471, 4412
- Krause E., Eifler T., 2017, MNRAS, 470, 2100
- Krause E. et al., 2017, preprint ([arXiv:1706.09359](https://arxiv.org/abs/1706.09359))
- Kuijken K., 1999, A&A, 352, 355

- Lewis A., 2009, MNRAS, 398, 471
Mandelbaum R. et al., 2014, ApJS, 212, 5
Mandelbaum R. et al., 2015, MNRAS, 450, 2963
Maoli R., Van Waerbeke L., Mellier Y., Schneider P., Jain B., Bernardeau F., Erben T., Fort B., 2001, A&A, 368, 766
Massey R. et al., 2007a, MNRAS, 376, 13
Massey R. et al., 2007b, ApJS, 172, 239
Miller L. et al., 2013, MNRAS, 429, 2858
Prat J. et al., 2017, Phys. Rev. D, preprint (arXiv:1708.01537)
Refregier A., Bacon D., 2003, MNRAS, 338, 48
Refregier A., Kacprzak T., Amara A., Bridle S., Rowe B., 2012, MNRAS, 425, 1951
Rhodes J., Refregier A., Groth E. J., 2001, ApJ, 552, L85
Schmidt F., Rozo E., Dodelson S., Hui L., Sheldon E., 2009, ApJ, 702, 593
Scoville N. et al., 2007, ApJS, 172, 1
Sheldon E. S., Huff E. M., 2017, ApJ, 841, 24
Simon P., Schneider P., 2017, A&A, 604, A109
Simpson F. et al., 2013, MNRAS, 429, 2249
Troxel M. A. et al., 2017, Phys. Rev. D, preprint (arXiv:1708.01538)
Van Waerbeke L. et al., 2000, A&A, 358, 30
Van Waerbeke L. et al., 2013, MNRAS, 433, 3373
Voigt L. M., Bridle S. L., 2010, MNRAS, 404, 458
Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2013, Phys. Rep., 530, 87
Whittaker L., Brown M. L., Battye R. A., 2015, MNRAS, 454, 2154
Wittman D. M., Tyson J. A., Kirkman D., Dell'Antonio I., Bernstein G., 2000, Nature, 405, 143
Zuntz J., Kacprzak T., Voigt L., Hirsch M., Rowe B., Bridle S., 2013, MNRAS, 434, 1604
Zuntz J. et al., 2017, MNRAS, preprint (arXiv:1708.01533) (Z17)

APPENDIX A: DERIVATION OF A TWO-POINT MODIFIER FOR SCALE-DEPENDENT BIAS

In the following we set out a brief derivation of the analytic modifications to account for scale-dependent neighbour effects the shear-shear two-point correlations used in the earlier section. We do not claim that this is a precise calculation of the sort that could be used to derive a robust calibration. Rather it is an order of magnitude estimate to allow us to assess the approximate size of the cosmological bias these effects could induce in the data.

First, with complete generality it is possible to write the i component of the measured shear at angular position θ as

$$\gamma_i^{\text{obs}}(\theta) = [1 + m_i(\theta)] \gamma_i(\theta), \quad (\text{A1})$$

where γ_i is the underlying true shear, which is sensitive to cosmology only. Extending this to the level of a two-point correlation between two populations α and β this implies:

$$\begin{aligned} \xi_i^{\text{obs},\alpha\beta}(\theta) &\equiv \left\langle \gamma_i^{\text{obs},\alpha}(\theta') \gamma_i^{\text{obs},\beta}(\theta' + \theta) \right\rangle_{\theta} \\ &= \left\langle [1 + m_i^{\alpha}(\theta')] [1 + m_i^{\beta}(\theta' + \theta)] \tilde{\gamma}_i^{\alpha}(\theta') \tilde{\gamma}_i^{\beta}(\theta' + \theta) \right\rangle_{\theta}. \end{aligned} \quad (\text{A2})$$

Note that the observed shear used in a particular bin correlation is now weighted by the overdensity of galaxies in the image, in addition to the calibration bias, such that

$$\tilde{\gamma}_i^{\alpha}(\theta) \equiv \left[1 + \delta_g^{\alpha}(\theta) \right] \times \gamma_i^{\alpha}(\theta). \quad (\text{A3})$$

Expanding each of the terms one finds

$$\begin{aligned} \xi_i^{\text{obs},\alpha\beta}(\theta) &= \left\langle \gamma_i^{\alpha}(\theta') \gamma_i^{\beta}(\theta' + \theta) \right\rangle_{\theta} + \left\langle m_i^{\alpha}(\theta') \gamma_i^{\alpha}(\theta') \gamma_i^{\beta}(\theta' + \theta) \right\rangle_{\theta} \\ &\quad + \left\langle m_i^{\beta}(\theta' + \theta) \gamma_i^{\alpha}(\theta') \gamma_i^{\beta}(\theta' + \theta) \right\rangle_{\theta} \end{aligned}$$

$$\begin{aligned} &+ \left\langle \delta_g^{\alpha}(\theta') \gamma_i^{\alpha}(\theta') \gamma_i^{\beta}(\theta' + \theta) \right\rangle_{\theta} \\ &+ \left\langle \delta_g^{\beta}(\theta' + \theta) \gamma_i^{\alpha}(\theta') \gamma_i^{\beta}(\theta' + \theta) \right\rangle_{\theta} \\ &+ \left\langle m_i^{\alpha}(\theta') m_i^{\beta}(\theta' + \theta) \gamma_i^{\alpha}(\theta') \gamma_i^{\beta}(\theta' + \theta) \right\rangle_{\theta} \\ &+ \left\langle \delta_g^{\alpha}(\theta') \delta_g^{\beta}(\theta' + \theta) \gamma_i^{\alpha}(\theta') \gamma_i^{\beta}(\theta' + \theta) \right\rangle_{\theta}. \end{aligned} \quad (\text{A4})$$

The terms contributing to the measured two-point shear correlation, then, is sensitive to both spatial correlations between the m in different galaxies and to the correlations with the source density. Note that we have chosen to neglect a higher order (six-point) term. In reality there will also be a connection between galaxy density and shear, but we will follow the normal convention and assume the contribution is small enough to be neglected. In simple terms, an excess in the $\langle mm \rangle$ term above the product of the mean m values independently could arise because galaxy pairs separated on small scales tend to come from *similar* image plane environments. In contrast the density weighted correlations $\langle \delta_g m \rangle$ would be zero, but for a simple observation; selecting a random galaxy with a suitable correlation pair at a distance θ is *not* the same as unconditionally selecting a random galaxy. In the small-scale bins we will oversample the dense regions, where m tends to be larger (see Section 6.2). The angular brackets here indicate averaging over all galaxy pairs separated by θ . If we can assume the bias is independent of the underlying cosmology the above expression simplifies significantly:

$$\begin{aligned} \xi_i^{\text{obs},\alpha\beta}(\theta) &= (1 + \bar{m}_i^{\alpha} + \bar{m}_i^{\beta} + \left\langle m_i^{\alpha}(\theta') m_i^{\beta}(\theta' + \theta) \right\rangle_{\theta}) \\ &\quad + \left\langle \delta_g^{\alpha}(\theta') m_i^{\beta}(\theta' + \theta) \right\rangle_{\theta} + \left\langle m_i^{\alpha}(\theta') \delta_g^{\beta}(\theta' + \theta) \right\rangle_{\theta} \\ &\quad + \left\langle \delta_g^{\alpha}(\theta') \delta_g^{\beta}(\theta' + \theta) \right\rangle_{\theta} \times \xi_i^{\alpha\beta}(\theta | \mathbf{p}), \end{aligned} \quad (\text{A5})$$

with $\xi_i^{\alpha\beta}$ being the true correlation function of cosmological shears $\langle \gamma_i \gamma_i \rangle$, which is contingent on the underlying cosmological parameters \mathbf{p} . It can be shown that

$$\begin{aligned} \xi_{+}(\theta) &\equiv \left\langle \gamma_{+}(\theta') \gamma_{+}(\theta' + \theta) \right\rangle_{\theta} \pm \left\langle \gamma_{\times}(\theta') \gamma_{\times}(\theta' + \theta) \right\rangle_{\theta} \\ &= \left\langle \gamma_1(\theta') \gamma_1(\theta' + \theta) \right\rangle_{\theta} \pm \left\langle \gamma_2(\theta') \gamma_2(\theta' + \theta) \right\rangle_{\theta} \\ &= \xi_1(\theta) + \xi_2(\theta), \end{aligned} \quad (\text{A6})$$

and so one can use equation (A5) to construct the observed ξ_{\pm} correlation functions

$$\begin{aligned} \xi_{\pm}^{\text{obs},\alpha\beta}(\theta) &= \left(1 + \bar{m}^{\alpha} + \bar{m}^{\beta} + \left\langle m^{\alpha}(\theta') m^{\beta}(\theta' + \theta) \right\rangle_{\theta} \right. \\ &\quad + \left\langle \delta_g^{\alpha}(\theta') m^{\beta}(\theta' + \theta) \right\rangle_{\theta} + \left\langle m^{\alpha}(\theta') \delta_g^{\beta}(\theta' + \theta) \right\rangle_{\theta} \\ &\quad \left. + \left\langle \delta_g^{\alpha}(\theta') \delta_g^{\beta}(\theta' + \theta) \right\rangle_{\theta} \right) \xi_{\pm}^{\alpha\beta}(\theta | \mathbf{p}). \end{aligned} \quad (\text{A7})$$

The i subscript has been discarded here under the assumption that m_1 and m_2 are approximately equal for a given set of galaxies.

Next, let us imagine that we have a measured datavector. Our measurements are biased, but we will assume it is possible to devise a correction that recovers the true cosmological signal precisely. Our observed datavector is then just,

$$\xi_{\pm}^{\text{obs},\alpha\beta}(\theta) = \Upsilon^{\text{tr},\alpha\beta} \xi_{\pm}^{\alpha\beta}(\theta | \mathbf{p}), \quad (\text{A8})$$

which follows trivially from equation (A7). Since we do not trivially know $\Upsilon^{\text{tr},\alpha\beta}$ *ab initio* (that is why we need simulations!) we can only construct a best-estimate approximation. By applying a correction factor to the raw measurements we construct a best-estimate

datavector:

$$\xi_{\pm}^{\text{BE},\alpha\beta}(\theta) = \frac{1}{\Upsilon^{\text{BE},\alpha\beta}} \xi_{\pm}^{\text{obs},\alpha\beta}(\theta) = \frac{\Upsilon^{\text{tr},\alpha\beta}}{\Upsilon^{\text{BE},\alpha\beta}} \xi_{\pm}^{\alpha\beta}(\theta|\mathbf{p}). \quad (\text{A9})$$

Of course, if our best correction is perfect then the ratio goes to unity, and we recover the underlying cosmology. Since we apply corrections to the single-galaxy shears we will assume $\Upsilon^{\text{BE},\alpha\beta}$ includes the $\langle \delta_g \delta_g \rangle$ term, but neglects the correlations involving m . We then can write:

$$\Upsilon^{\text{BE},\alpha\beta} = \left(1 + \bar{m}^\alpha + \bar{m}^\beta + \bar{m}^\alpha \bar{m}^\beta + \left\langle \delta_g^\alpha(\theta') \delta_g^\beta(\theta' + \theta) \right\rangle_\theta \right). \quad (\text{A10})$$

We can measure the mean bias in each bin that would be obtained from the calibration directly. As we show in Z17, using the full DES Y1 HOOPOE catalogues, these biases are ~ -0.08 to -0.20 .

Finally, assume that although m clearly varies between redshift bins, the strength of the correlation does not. That is, the bias–bias term is the product of the mean m s (which varies between z bins) plus a scale-dependent shift (which does not). One then has

$$\langle m^\alpha(\theta') m^\beta(\theta' + \theta) \rangle_\theta = \bar{m}^\alpha \bar{m}^\beta + \Delta \xi_{mm}(\theta). \quad (\text{A11})$$

The additive part can be measured directly from the simulation using sub-patches, as described earlier. The density–density correlation can be obtained in the same way. This, then, leaves only the $m \times \delta_g$ cross-correlation. This should vanish in the case of zero correlation, but it also seems reasonable to assume that the magnitude should be proportional to the mean bias \bar{m}^α in a particular bin. This allows the scale-dependent (non-tomographic) cross-correlation measured from HOOPOE to be rescaled appropriately for each bin pair:

$$\langle \delta_g^\alpha(\theta') m^\beta(\theta' + \theta) \rangle_\theta = \left(\frac{\bar{m}^\beta}{\bar{m}} \right) \xi_{gm}(\theta), \quad (\text{A12})$$

where \bar{m} is the global multiplicative bias and $\xi_{gm}(\theta) \equiv \langle m \delta_g \rangle$, each measured using all simulated galaxies. Using the above equations, with our fiducial calibration and three measured correlations, one can derive a scale-dependent modification to shear–shear two-point correlation data using equation (A8).

¹Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL, UK

²Institute for Astronomy, University of Edinburgh, Edinburgh EH9 3HJ, UK

³Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, OH 43210, USA

⁴Department of Physics, The Ohio State University, Columbus, OH 43210, USA

⁵Kavli Institute for Particle Astrophysics & Cosmology, PO Box 2450, Stanford University, Stanford, CA 94305, USA

⁶SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

⁷Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA

⁸Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109, USA

⁹Department of Physics, ETH Zürich, Wolfgang-Pauli-Strasse 16, CH-8093 Zürich, Switzerland

¹⁰Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK

¹¹Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa

¹²Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, USA

¹³LSST, 933 North Cherry Avenue, Tucson, AZ 85721, USA

¹⁴CNRS, UMR 7095, Institut d'Astrophysique de Paris, F-75014 Paris, France

¹⁵Sorbonne Universités, UPMC Univ Paris 06, UMR 7095, Institut d'Astrophysique de Paris, F-75014 Paris, France

¹⁶Laboratório Interinstitucional de e-Astronomia – LIneA, Rua Gal. José Cristino 77, Rio de Janeiro, RJ 20921-400, Brazil

¹⁷Observatório Nacional, Rua Gal. José Cristino 77, Rio de Janeiro, RJ 20921-400, Brazil

¹⁸Department of Astronomy, University of Illinois, 1002 W. Green Street, Urbana, IL 61801, USA

¹⁹National Center for Supercomputing Applications, 1205 West Clark St., Urbana, IL 61801, USA

²⁰Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, E-08193 Bellaterra (Barcelona), Spain

²¹Institute of Space Sciences, IEEC-CSIC, Campus UAB, Carrer de Can Magrans, s/n, E-08193 Barcelona, Spain

²²Department of Physics, IIT Hyderabad, Kandi, Telangana 502285, India

²³Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA

²⁴Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, E-28049 Madrid, Spain

²⁵Department of Astronomy, University of Michigan, Ann Arbor, MI 48109, USA

²⁶Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

²⁷Astronomy Department, University of Washington, Box 351580, Seattle, WA 98195, USA

²⁸Cerro Tololo Inter-American Observatory, National Optical Astronomy Observatory, Casilla 603, La Serena, Chile

²⁹Santa Cruz Institute for Particle Physics, Santa Cruz, CA 95064, USA

³⁰Australian Astronomical Observatory, North Ryde, NSW 2113, Australia

³¹Argonne National Laboratory, 9700 South Cass Avenue, Lemont, IL 60439, USA

³²Departamento de Física Matemática, Instituto de Física, Universidade de São Paulo, CP 66318, São Paulo, SP 05314-970, Brazil

³³Department of Physics and Astronomy 4242 TAMU, College Station, TX 77843, USA

³⁴Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA

³⁵Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA

³⁶Institució Catalana de Recerca i Estudis Avançats, E-08010 Barcelona, Spain

³⁷Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Av. Complutense, 40, E-28040 Madrid, Spain

³⁸Brookhaven National Laboratory, Bldg 510, Upton, NY 11973, USA

³⁹School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, UK

⁴⁰Instituto de Física Gleb Wataghin, Universidade Estadual de Campinas, 13083-859 Campinas, SP, Brazil

⁴¹Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

⁴²Institute of Cosmology & Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK

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