The Dark Energy Survey Year 3 and eBOSS: constraining galaxy intrinsic alignments across luminosity and colour space

S. Samuroff [1,2][*] R. Mandelbaum,[2,3] J. Blazek,[1] A. Campos,[2,3] N. MacCrann,[4] G. Zacharegkas,[5]
A. Amon,[6,7] J. Prat,[5,8] S. Singh,[2,3] J. Elvin-Poole,[9] A. J. Ross,[10] A. Alarcon,[11] E. Baxter,[12] K. Bechtol,[13]
M. R. Becker,[11] G. M. Bernstein,[14] A. Carnero Rosell,[15,16,17] M. Carrasco Kind,[18,19] R. Cawthon,[20]
C. Chang,[5,8] R. Chen,[21] A. Choi,[22] M. Crocce,[23,24] C. Davis,[25] J. DeRose,[26] S. Dodelson,[2,3] C. Doux,[14,27]
A. Drlica-Wagner,[5,8,28] K. Eckert,[14] S. Everett,[29] A. Ferté,[29] M. Gatti,[14] G. Giannini,[5,8,30] D. Gruen,[31]
R. A. Gruendl,[18,19] I. Harrison,[32] K. Herner,[28] E. M. Huff,[29] M. Jarvis,[14] N. Kuropatkin,[28] P.-F. Leget,[25]
P. Lemos,[33] J. McCullough,[25] J. Myles,[25,34] A. Navarro-Alsina,[35] S. Pandey,[14] A. Porredon,[10,36]
M. Raveri,[37] M. Rodriguez-Monroy,[38] R. P. Rollins,[39] A. Roodman,[25,40] G. Rossi,[41] E. S. Rykoff,[25,40]
C. Sánchez,[24,14] L. F. Secco,[5,8] I. Sevilla-Noarbe,[38] E. Sheldon,[42] T. Shin,[43] M. A. Troxel,[21]
I. Tutusaus,[23,24,43] N. Weaverdyck,[44,26] B. Yanny,[29] B. Yin,[2] Y. Zhang,[45,46] J. Zuntz,[36] M. Aguena,[16]
O. Alves,[45] J. Annis,[28] D. Bacon,[47] E. Bertin,[48,49] S. Bocquet,[31] D. Brooks,[50] D. L. Burke,[34,40]
J. Carretero,[30] M. Costanzi,[51,52,53] L. N. da Costa,[16] M. E. S. Pereira,[54] J. De Vicente,[38] S. Desai,[55]
H. T. Diehl,[28] J. P. Dietrich,[31] P. Doel,[50] I. Ferrero,[56] B. Flaugher,[28] J. Frieman,[28,5] J. García-Bellido,[57]
S. R. Hinton,[58] D. L. Hollowood,[59] K. Honscheid,[10] D. J. James,[60] K. Kuehn,[61,62] O. Lahav,[50]
J. L. Marshall,[63] P. Melchior,[64] J. Mena-Fernández,[38] F. Menanteau,[18,19] R. Miquel,[65,30] J. Newman,[66]
A. Palmese,[67,2] A. Pieres,[16,67,68] A. A. Plazas Malagón,[64] E. Sanchez,[38] V. Scarpine,[28] M. Smith,[69]
E. Suchyta,[70] M. E. C. Swanson,[50,60] G. Tarle,[45] and C. To[10] (DES Collaboration)

Affiliations are listed at the end of the paper

ABSTRACT

We present direct constraints on galaxy intrinsic alignments (IAs) using the Dark Energy Survey Year 3 (DES Y3), the Extended Baryon Oscillation Spectroscopic Survey (eBOSS), and its precursor, the Baryon Oscillation Spectroscopic Survey (BOSS). Our measurements incorporate photometric red sequence (redMaGiC) galaxies from DES with median redshift $z \sim 0.2$–1.0, luminous red galaxies from eBOSS at $z \sim 0.8$, and also an SDSS-III BOSS CMASS sample at $z \sim 0.5$. We measure two-point IA correlations, which we fit using a model that includes lensing, magnification, and photometric redshift error. Fitting on scales $6 \, \mathrm{Mpc} \, h^{-1} < r_\mathrm{p} < 70 \, \mathrm{Mpc} \, h^{-1}$, we make a detection of IAs in each sample, at 5σ–22σ (assuming a simple one-parameter model for IAs). Using these red samples, we measure the IA–luminosity relation. Our results are statistically consistent with previous results, but offer a significant improvement in constraining power, particularly at low luminosity. With this improved precision, we see detectable dependence on colour between broadly defined red samples. It is likely that a more sophisticated approach than a binary red/blue split, which jointly considers colour and luminosity dependence in the IA signal, will be needed in future. We also compare the various signal components at the best-fitting point in parameter space for each sample, and find that magnification and lensing contribute ~ 2–18 per cent of the total signal. As precision continues to improve, it will certainly be necessary to account for these effects in future direct IA measurements. Finally, we make equivalent measurements on a sample of emission-line galaxies from eBOSS at $z \sim 0.8$. We constrain the non-linear alignment amplitude to be $A_1 = 0.07^{+0.32}_{-0.42}$ ($|A_1| < 0.78$ at 95 per cent CL).

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

1 INTRODUCTION

The study of cosmic shear as a probe of the large-scale structure of the Universe has developed rapidly over the past decade. Although its

[*] E-mail: s.samuroff@northeastern.edu

potential was recognized some time ago (see e.g. Jain & Seljak 1997; Hu 1999), only more recently have high-precision cosmological constraints been possible. In the past 10 yr, data sets have grown to the point where weak lensing measurements alone have roughly comparable power to constrain certain cosmological parameters as the cosmic microwave background (CMB) temperature fluctuations. Galaxy weak lensing and the CMB are both sensitive to the amplitude of the matter power spectrum in the low-redshift Universe, S_8. Although lensing allows one to probe the late-time matter field directly, the primary temperature anisotropies of the CMB provide a somewhat more complicated route, relying on an extrapolation from the surface of last scattering to the present day. Ever since the results of Heymans et al. (2013) lensing measurements have given a typically lower S_8 than the CMB; interestingly, this finding holds across multiple lensing surveys, whose members have implemented their own independent, well-tested, blind analyses (Kilbinger et al. 2013; Dark Energy Survey Collaboration 2016; Jee et al. 2016; Hildebrandt et al. 2017, 2020; Troxel et al. 2018; Hikage et al. 2019; Hamana et al. 2020; Asgari et al. 2021; Amon et al. 2022; Secco, Samuroff et al. 2022). The current level of (dis)agreement in the full parameter space, as assessed using various different metrics (Lemos et al. 2021), is at the level of up to ~2.5σ (although it differs significantly between surveys and probe combinations).

Future lensing surveys will have much smaller statistical uncertainties compared with the current generation, which will greatly increase the precision of weak lensing measurements. This will, in turn, improve our constraining power and help us make sense of the apparent tensions in the literature. It will also, however, require a much tighter control of modelling errors in order to avoid our analyses becoming dominated by systematic uncertainties. Although much progress has been made in recent years, and the methods for mitigating systematics are highly sophisticated, we still have some way to go, as a field. One outstanding gap in our understanding is the treatment of intrinsic alignments (IAs; Joachimi et al. 2015; Kiessling et al. 2015; Kirk et al. 2015; Troxel & Ishak 2015).

IAs are shape correlations induced not by cosmological lensing, but by local interactions, which can confuse the interpretation of the cosmic shear measurement. Most obviously, galaxies that are physically close by to each other experience the same background tidal field, which couples their intrinsic shapes, inducing what are known as II correlations. Additionally, GI (shear-intrinsic) correlations are generated by the fact that the same foreground matter that interacts with foreground galaxies also lenses background sources. A significant amount of literature over the past few years has focused on developing analytic models for IAs, which allow them to be forward modelled and marginalized in cosmological analyses (Catelan, Kamionkowski & Blandford 2001; Mackey, White & Kamionkowski 2002; Hirata & Seljak 2004; Bridle & King 2007; Blazek, Vlah & Seljak 2015; Blazek et al. 2019; Vlah, Chisari & Schmidt 2020; Fortuna et al. 2021a).

Perhaps the most well-established approach is an analytic formalism that assumes the intrinsic shapes of galaxies are linear in the background tidal field, and frozen at the point of galaxy formation (Catelan et al. 2001; Hirata & Seljak 2004). What became known as the linear alignment (LA) model predicts both the GI and II power spectra and is, by convention, normalized such that the free amplitude A_1 is very roughly one for a typical lensing source sample (i.e. a mixed colour sample, dominated by blue galaxies at $z \lesssim 1$). A few years later, Hirata et al. (2007) and Bridle & King (2007) introduced a modification, whereby the full non-linear matter power spectrum is used in place of the linear version in the LA model equations. This has been shown to improve the performance of the

model on scales ~a few h^{-1} Mpc (Blazek, McQuinn & Seljak 2011). More recently, Blazek et al. (2015, 2019) take further steps along this route. The perturbative model developed in those papers (known as the Tidal Alignment and Tidal Torque model, TATT) extends the LA model to include higher order terms. Although in principle there are specific physical mechanisms for how correlations that are, for example, quadratic in the tidal field arise, in practice the model is agnostic to the underlying physics. An alternative approach, which is more closely connected with the physics on sub-halo scales, is to use a version of the halo model. The basic concept was introduced a decade ago (Schneider & Bridle 2010), and more recently Fortuna et al. (2021a) took significant steps towards developing a practical implementation.

Although a useful tool for learning about IAs, pure theory cannot provide a complete picture. Real data are very much necessary for properly understanding their behaviour in the real universe. Broadly, measurements can be classified as *direct* (i.e. using a statistic that is dominated by IAs, with little or no contribution from lensing), or *simultaneous* (i.e. where IAs contribute only a small part of the total signal, and are inferred alongside cosmological and other parameters). There have also been studies that have sought to do something in between, using particular combinations of lensing data to try to isolate an IA signal (e.g. Zhang 2010; Blazek et al. 2012). By this definition, almost all cosmic shear studies to date are simultaneous IA measurements. Although comparison is complicated by non-trivial differences in the samples and measurement methods, as well as the high-dimensional model space, such constraints have typically found IA amplitudes for mixed lensing samples in the range $A_1 \sim 0.1$–1 (Troxel et al. 2018; Hildebrandt et al. 2020; Asgari et al. 2021; Secco et al. 2022). A smaller number of works have attempted to understand how IAs enter simultaneous measurements in more detail. For example, Heymans et al. (2013) split the CFHTLenS source catalogue into early and late types and perform independent cosmic shear analyses; they report $A_1 \sim 0$ in the bluer population, and $A_1 \sim 5$ in early types (albeit with large error bars). Several years later, Samuroff et al. (2019) implemented a similar colour-split methodology to explore IAs in DES Y1, this time analysing red and blue galaxies along with their cross-correlations simultaneously. Assuming the non-linear alignment (NLA) model, that work found qualitatively similar results, with blue galaxies consistent with zero alignments and $A_1 \sim 3$ in the red population. Using the TATT model, it found the quadratic alignment amplitude A_2 to be <0 at the level of ~2σ in both colour samples.

Direct measurements are typically restricted by the need for precise estimates for the redshifts of individual galaxies, and for this reason have tended to focus on bright red samples. A number of such studies have been carried out over the years (Hirata et al. 2007; Joachimi et al. 2011; Singh, Mandelbaum & More 2015; Johnston et al. 2019; Fortuna et al. 2021b), and the alignment strength as a function of luminosity is relatively well measured in brighter populations. Since these samples tend to have compact redshift distributions, any given study only weakly (if at all) constrains the redshift evolution of IAs. In the case of bluer galaxies, a handful of direct measurements have been attempted (Mandelbaum et al. 2011; Tonegawa et al. 2018; Johnston et al. 2019), but the samples here are typically small; though they make null detections and place upper limits on the IA amplitude, the error bars are wide enough to allow for significantly non-zero values. Although analogous IA measurements can be and have been made on hydrodynamic simulations, these are limited by finite box size, difficulty in constructing realistic galaxy samples, and the accuracy of the simulations themselves (Codis et al. 2015; Hilbert et al. 2017; Samuroff, Mandelbaum & Blazek 2021).

This paper falls into the category of direct measurements, and represents the first such exercise using DES. We use a combination of DES redMaGiC (or the 'red sequence Matched filter Galaxy Catalog'; photometric, but with precise per-galaxy redshift estimates) and the overlapping BOSS and eBOSS surveys (spectroscopic) to measure IA correlations in physical space, which we then fit using a range of IA models. This work follows implicit IA constraints from DES cosmic shear (Amon et al. 2022; Secco et al. 2022) and galaxy–galaxy lensing (Prat et al. 2022; Sánchez, Prat et al. 2022). We should note that, while they use the same DES catalogues, the samples in these earlier works are significantly different from ours, and so we do not expect the IA signal to be the same.

The paper is structured as follows. In Section 2, we describe the various data sets used in this work. Section 3 then outlines measurements on these data, including redshifts, calibrated galaxy shapes, and two-point correlations. In Section 4, we set out the model used fit those measurements, and discuss our analysis choices such as priors and scale cuts; a range of validation tests of that theory pipeline, using real and simulated data, are outlined in Section 5. Our main results are discussed in Section 6. We conclude in Section 7.

2 DATA

In this section, we briefly describe the data sets used in this work, and how the various galaxy samples are defined.

2.1 The Dark Energy Survey Year 3

The Dark Energy Survey is a 6-yr programme encompassing ∼5000 square deg of the Southern sky using the Víctor Blanco telescope in Chile. The approximate footprint is shown (purple) in Fig. 1. Images were taken in five photometric bands (*g*, *r*, *i*, *z*, and *Y*) with a nominal depth of magnitude *r* ∼ 24.1 (at full Y6 depth); although all five are used for redshift estimation, galaxy shape measurements are limited to *riz* due to difficulties in accurately estimating the point spread function (PSF) in the *g* band (Jarvis et al. 2021), and shallower imaging in the *Y* band (45 s exposures as opposed to 90 s in *griz*). The work described in this paper is based on data collected during the first 3 yr of operation (the Y3 data). These data cover the full area at slightly less than the full depth, with an average of about five exposures per galaxy. A description of the image processing and reduction pipeline, including background subtraction and object detection, can be found in Morganson et al. (2018) and Sevilla-Noarbe, Bechtol et al. (2021). The photometric data set, before any further cuts, is known as the GOLD catalogue (Sevilla-Noarbe et al. 2021). In Y3, this has a limiting magnitude *i* ∼ 23. Per-galaxy photometry measurements are obtained using multi-object fitting (Drlica-Wagner et al. 2018), and shapes are measured using the METACALIBRATION algorithm (Gatti, Sheldon et al. 2021; see Section 3.1).

2.2 BOSS and eBOSS

We also use galaxies from both the Baryon Oscillation Spectroscopic Survey (BOSS) and the Extended Baryon Oscillation Spectroscopic Survey (eBOSS) in this paper, and so we discuss both briefly here. The former is a spectroscopic sample collected as part of SDSS-III (Eisenstein et al. 2011). Imaging in five photometric bands (*ugriz*) and spectroscopy for BOSS were performed using the 2.5 m Sloan Telescope at Apache Point Observatory in New Mexico (Gunn et al. 2006; Smee et al. 2013). The observing program took place

between autumn 2009 and spring 2014, and covered more than 1.5M galaxies across 10 000 square deg of high-latitude sky. One can find a description of the BOSS spectrographs and other details in Dawson et al. (2013).

The eBOSS data are slightly more recent, being taken over the period between summer 2014 and spring 2019, as part of SDSS-IV (Blanton et al. 2017; see also Dawson et al. 2016 for a discussion of the differences between BOSS and eBOSS). Again, spectroscopy relied on the BOSS spectrographs on the Sloan Telescope. The method for target selection differs slightly depending on the nature of the galaxy sample. Emission-line galaxies (ELGs) were targeted from the DECam Legacy Survey (DECaLS), which is deeper than SDSS, detections, and photometry. DECaLS was carried out using the DECam on the Blanco telescope, and covers an area of 6700 square deg in the region −20 deg < *δ* < +30 deg to a 5*σ* limiting magnitude of 24.7, 23.9, and 23.0 in the *g*, *r*, and *z* bands, respectively (compared with 22.8, 22.3, and 20.4 for SDSS; Delubac et al. 2017). Luminous red galaxy (LRG) targets, on the other hand, were selected using SDSS *riz* imaging and infrared sky maps from the Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010).

The BOSS and eBOSS footprints are divided in two approximately equal area regions; the ELG and LRG samples used in this work come from the Southern Galactic Cap (SGC), which contains roughly 600 square deg of overlap with the DES footprint.

2.3 Primary Galaxy samples

In this paper we consider five distinct samples (for a summary, see Table 1). These are:

(i) **eBOSS Emission Line Galaxies**: Our eBOSS ELG sample contains ∼100 000 galaxies in the SGC region. The target selection process in described in more detail in Raichoor et al. (2017), and further discussion can be found in Tamone et al. (2020) and Raichoor et al. (2021). The cuts that define the sample are relatively complex, and include a *g*-band magnitude limit at *g* = 22.8 mag, as well as selection in colour space designed to limit the redshift range. In total, we have 92 954 galaxies, with a mean redshift of about 0.8. Although this sample covers a similar range in redshift to the LRGs described below, they are significantly bluer than any of our other samples (both in apparent and rest-frame magnitudes), and are thus not expected to exhibit strong IAs.

(ii) **eBOSS Luminous Red Galaxies**: A sample of LRGs from eBOSS DR16, drawn from the SGC region. Selection is performed using the criteria outlined in Prakash et al. (2016) (see also Ross et al. 2020; Bautista et al. 2021; Rossi et al. 2021 for details on the LRG sample); the cuts are primarily on extinction-corrected magnitudes and colours (*z* < 19.95 mag). The redshift coverage is similar to that of the ELG sample, with a mean of *z* ∼ 0.8, and the total number of galaxies is 22 244.

(iii) **BOSS CMASS**: The CMASS selection algorithm is described in detail by Reid et al. (2016) (see their section 3.3 and references therein). We additionally impose redshift cuts at *z* < 0.6 to ensure there is no overlap with the eBOSS LRGs, and at *z* > 0.4 to remove outliers below the intended CMASS redshift range as they may have atypical colours and luminosities compared to those within the intended redshift range. This leaves us with 49 820 galaxies.

(iv) **redMaGiC high-*z* (RMH)**: A sample of red sequence galaxies from the DES Y3 redMaGiC catalogue. These objects are selected using the algorithm outlined in Rozo et al. (2016). In brief, all detected galaxies are fitted using a red sequence template, yielding a best-fitting redshift, z_{redmagic}, and a derived luminosity *L*, as well as a
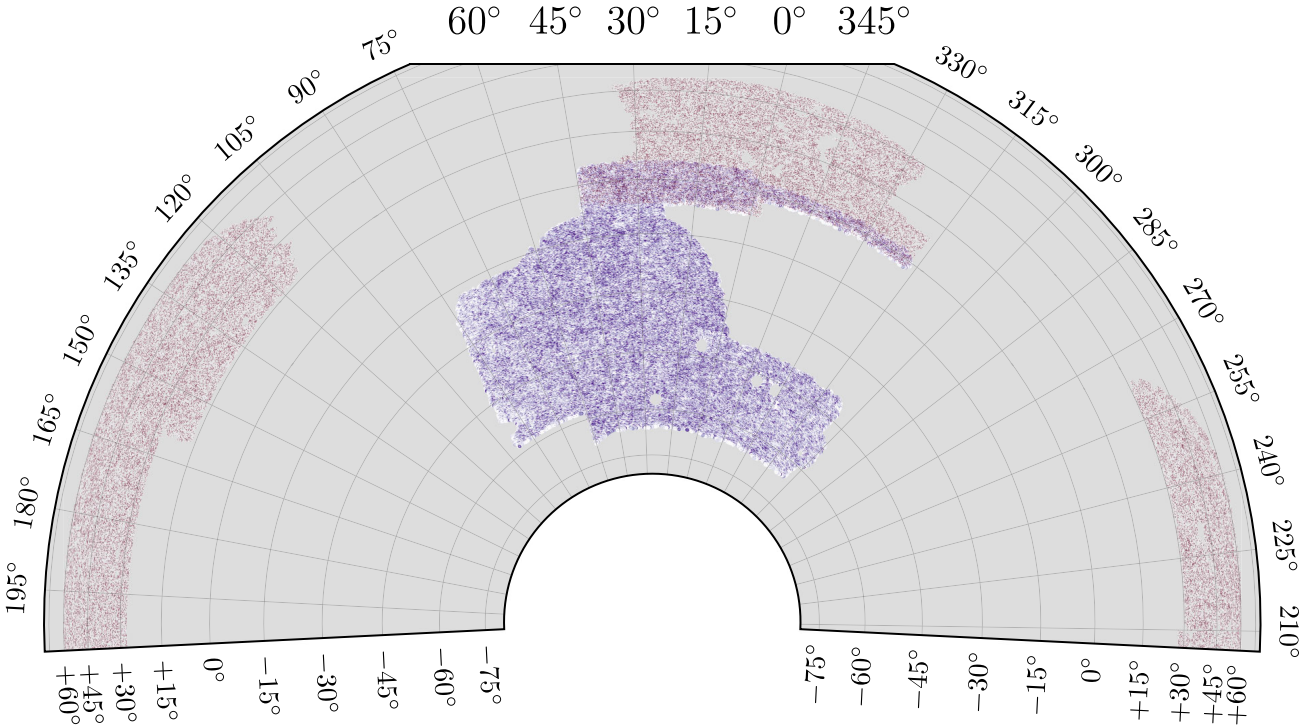
Figure 1. The approximate footprint of the Dark Energy Survey (purple) and SDSS eBOSS (pink). The overlap in the Stripe 82 region (across the beak and head of the hummingbird shape) is roughly 600 square deg.

Table 1. A summary of the properties of the shape samples used in this work. From left, we show the total number of galaxies after cuts, the area of the footprint in square degrees, and the mean comoving number density (averaged over redshift). The value f_D indicates the number of galaxies in the density tracer sample relative to the size of the shape sample. Note that this reflects both the impact of METACALIBRATION cuts and, in the SDSS samples, the geometric selection. The ellipticity dispersion σ_e is defined according to Heymans et al. 2012 (see also Gatti et al. 2021 equation 13). Also shown are the ensemble mean redshift $\langle z \rangle$, the selection response for each sample (the shear response due to the galaxy shape catalogue cuts; see Sheldon & Huff 2017), and the mean rest-frame colour and r-band magnitude. The means incorporate the weights described in Section 3.2.

	No. of Galaxies	f_D	Area [sq. deg.]	$\bar{n}_c$ [$10^4\ h^3$ Mpc^{-3}]	σ_e	Mean redshift $\langle z \rangle$	$R_S \times 10^4$	$\langle M_r - M_z \rangle$	$\langle M_r \rangle$
redMaGiC high-z	754 574	1.13	4203	3.22	0.28	0.78	4.8	0.43	−21.6
redMaGiC low-z	1.61 M	1.14	4203	7.64	0.26	0.46	23.1	0.43	−20.9
eBOSS ELGs	92 954	1.10	604	1.95	0.24	0.84	− 4.2	−	−
eBOSS LRGs	22 244	3.03	604	0.55	0.26	0.75	− 4.7	0.37	−21.9
CMASS	49 820	3.35	604	4.11	0.23	0.52	− 4.0	0.36	−21.8

corresponding χ^2. If the galaxy falls above a minimum L and below a maximum χ^2, it is included in the catalogue. The χ^2 threshold is a function of redshift, such that the comoving density is constant (Rozo et al. 2016, section 3.3). This process gives a set of bright red galaxies with both well-constrained per-object photometric redshifts ($\sigma_z/(1 + z) < 0.02$) and well-understood redshift error. Our high-z sample consists of the upper two redshift bins of the lens sample used in Dark Energy Survey Collaboration (2022), cut at $z > 0.6$. The luminosity threshold is $L_{\min} = L_*$, where L_* is a characteristic luminosity, as defined in Rozo et al. (2016), section 3.1. The sample before shape cuts comprises ∼0.8 M galaxies over 4203 square deg. The redshift distribution is relatively compact and peaks at a similar value to our eBOSS samples at $z \sim 0.8$.

(v) **redMaGiC low-z (RML)**: Our low-z redMaGiC sample is defined in a similar way to redMaGiC high-z, with key differences. Primarily, the luminosity threshold is lower at $L_{\min} = 0.5L_*$ (Dark Energy Survey Collaboration 2022). A cut on z_{redmagic} is imposed at $z < 0.6$, equivalent to the three lower lens bins from Dark Energy

Survey Collaboration (2022). Without shape cuts, the catalogue contains 1.84 M objects, with a median redshift of $z \sim 0.5$.

In each case we define density and shape tracer selections. The former uses all galaxies passing the baseline cuts described above, and also are not required to be within the DES-SDSS overlap. In each sample, we obtain galaxy shape estimates by matching galaxies to the DES Y3 METACALIBRATION catalogue (Gatti et al. 2021). We construct a KDTree of METACALIBRATION galaxy angular positions, which is used to locate the nearest DES neighbour for each eBOSS or redMaGiC object. A match tolerance of 1 arcsec is imposed to exclude spurious matches, and objects outside the overlap region between the two surveys. To obtain a subset of galaxies with reliable shapes, we then impose the cuts recommended by Gatti et al. (2021) (their section 4.2), which includes selections based on size and signal-to-noise ratio, as well as a cut designed to remove binary star contamination. We show the estimated redshift distributions $n(z)$ for each sample in the top panel of Fig. 2 (see Section 3.4 for more detail

Figure 2. Top: The estimated redshift distributions of the various shape samples used in this work. Each $n(z)$ is normalized to integrate to 1 over the redshift range shown. The $n(z)$s for the spectroscopic samples are shown as shaded curves, and are estimated as the histogram of single-galaxy z estimates. The unshaded $n(z)$s are estimated by stacking random samples from the redMaGiC redshift PDFs. **Bottom:** The same, but showing comoving number density as a function of redshift. Note that n_c is weakly cosmology-dependent, and so we assume the fiducial cosmology specified in Section 4. Note that both the $n(z)$ and $n_c(z)$ are qualitatively the same for the density tracer samples. The shape cuts remove galaxies, but do not change the shape or mean redshift of these distributions significantly.

about how these are estimated), and the comoving number density $n_c(z)$ in the bottom.

2.4 Comparison sample: LOWZ

In addition to the five catalogues discussed above, we also use BOSS LOWZ (Dawson et al. 2013) as a reference sample. The point of including these data is to test our measurement and inference pipelines by comparing against the baseline analysis of Singh et al. (2015). LOWZ is a convenient choice for this, since there are relatively detailed published results using a very similar methodology to our own.

LOWZ is a sample of LRGs from BOSS DR11. The sample covers a footprint of 9243 square deg and is approximately volume-limited over the redshift range $0.16 < z < 0.36$; a sharp cut-off is imposed at these bounds. Unlike the other samples, we do not match to Y3 METACALIBRATION to obtain shape estimates, but simply use the pre-existing catalogues (Reyes et al. 2012). For all other catalogue-level quantities (redshifts, $k + e$-corrected magnitudes etc), we likewise use the pre-computed columns (see Singh et al. 2015 for details). After cuts, the LOWZ shape and density tracer samples contain 159 620 and 173 854 galaxies, respectively.

3 MEASUREMENTS

3.1 Shapes

The galaxy shapes for all samples except LOWZ are obtained by matching to the DES Y3 METACALIBRATION catalogue. Discussion of the shape measurement algorithm, and catalogue level tests, can be found in Gatti et al. (2021). The basic measurement is a maximum likelihood fit of an elliptical Gaussian to each galaxy. This process uses a Markov chain Monte Carlo, and is performed over multiple

exposures and in bands *riz* simultaneously. In order to calibrate biases due to image noise, model bias, and other effects, the fit is repeated four times using artificially PSF-deconvolved and resheared images, a technique known as METACALIBRATION. For details of how the METACALIBRATION corrections are applied in this particular context see Section 3.6; for the general case and validation on simulations see Huff & Mandelbaum (2017), Sheldon & Huff (2017), and Zuntz et al. (2018).

3.2 Galaxy weights

For galaxy clustering and galaxy–shape measurements, we use the recommended weights for each sample. Descriptions of these can be found in Raichoor et al. (2017) and Ross et al. (2020) (for eBOSS), Reid et al. (2016) (for CMASS), and Rodríguez-Monroy et al. (2022) (redMaGiC). These are designed to correct for correlations between the observed galaxy number density and various survey properties, which can be induced by systematics. For the SDSS samples, there are additional weights designed to account for fibre collisions and redshift failures, which are combined as per the references above.

It is worth noting briefly that previous works (see e.g. Ross et al. 2020) identified a possible systematic due to variations in the redshift distributions of the eBOSS samples within the SGC and NGC regions, which is not explicitly corrected by the weights. Although relatively mild for eBOSS LRGs, it was found to be significant enough to need correcting for in a Redshift Space Distortion (RSD) analysis using the ELG sample (Tamone et al. 2020; Bautista et al. 2021; de Mattia et al. 2021). We do not, however, believe this to be a significant concern for our analysis, given the fact that our IA constraint (from the projected shape–galaxy and shape–shape correlations, $w_{g+} + w_{++}$) is constrained to the DES-eBOSS overlap region, which is a relatively small part of the overall eBOSS SGC footprint. Although we do use the full area for the galaxy–galaxy two-point measurement w_{gg}, given that the result from ELGs is essentially a null detection (see Section 6.1), we do not expect a small systematic affecting the galaxy bias to be a significant factor.

The shape catalogues for the different samples are all ultimately subsets of Y3 METACALIBRATION, and so we adopt the inverse variance weights discussed in Gatti et al. (2021).

3.3 Magnification coefficients

In addition to imprinting a coherent pattern in their shapes, lensing by large-scale structure also modulates the observed brightness and size of galaxies, an effect known as magnification. In order to model the impact on our galaxy number counts, we require an estimate for the slope of the faint end of the galaxy luminosity function for each of our density tracer samples (see e.g. Mandelbaum et al. 2005; Elvin-Poole et al. 2023 and Joachimi & Bridle 2010's appendix A). Our fiducial estimates are derived via what we refer to as the 'flux-only' method (Elvin-Poole et al. 2023). In the cases of eBOSS, CMASS, and LOWZ the process is straightforward. For a particular catalogue containing N_0 galaxies, with a given pre-existing selection function, we apply a small achromatic shift δm to the observed magnitudes. We reapply the magnitude cuts using this perturbed catalogue, and count how many galaxies are lost to the bright-end cut δN_-. The sign of δm is then flipped, and the process repeated to estimate the number shifted up over the faint-end threshold δN_+. The total change in number counts is then simply:

$$\delta N(\delta \kappa) = \delta N_+ - \delta N_- \tag{1}$$

with $\delta\kappa = 0.5(10^{-\delta m/2.5} - 1)$. For small perturbations we can measure the slope of $\delta N(\delta\kappa)/N_0$ with $\delta\kappa$ numerically. This gives us a quantity Elvin-Poole et al. (2023) referred to as C_{sample}, which describes the linear response of the observed galaxy number density to a small change in κ. We define a quantity referred to as the magnification coefficient as $\alpha = C_{sample}/2$ (see Section 4 for how this enters the theory predictions).

For redMaGiC, the sample selection is more complex. For this reason, we start with the Y3 GOLD catalogue (Sevilla-Noarbe et al. 2021), perturb the magnitudes, and rerun the redMaGiC algorithm for each δm. We then estimate α in the same way as before. We find $\alpha^{RML} = 1.101$ for redMaGiC low-z and $\alpha^{RMH} = 1.719$ for redMaGiC high-z. For our LRG, ELG and CMASS samples we find $\alpha^{LRG} = 2.020$, $\alpha^{ELG} = 1.177$, and $\alpha^{CMASS} = 0.529$, respectively.

In addition to the flux-only estimates, we have alternative values, derived using an algorithm called BALROG (Suchyta et al. 2016; Everett et al. 2022): $\alpha^{RMH} = 2.11 \pm 0.32$ and $\alpha^{RML} = 0.20 \pm 0.29$ for redMaGiC high-z and low-z. BALROG works by inserting additional synthetic galaxies into real photometric images. By running the detection and measurement processes on the altered BALROG images, one can sample the selection function of the survey and explore effects such as magnification and blending. Although these, in principle, capture size selection effects that the flux-only numbers cannot (see Elvin-Poole et al. 2023 for discussion), they are also relatively noisy. We also have BALROG estimates for the redMaGiC samples only, and not CMASS/eBOSS. We thus use the flux-only estimates as our fiducial choice; we do, however, confirm that in the two redMaGiC samples our basic conclusions are unaffected by this choice (see Section 6.3).

3.4 Redshift distributions

3.4.1 Spectroscopic redshifts for BOSS and eBOSS

For details on the BOSS and eBOSS spectroscopic redshift pipelines see Comparat et al. (2016), Hutchinson et al. (2016), and Bolton et al. (2012). In brief, galaxy spectra are collected using the BOSS spectrographs on the Sloan Telescope (Smee et al. 2013); the instrument has two identical spectrographs, each of which has a red and a blue camera, collectively covering the wavelength range 360–1040 nm, and 1000 optical fibres, 3 arcsec in diameter, and with a collision scale of 62 arcsec (corresponding to a physical scale of $\sim 0.6~h^{-1}$ Mpc at $z = 0.8$). Each object is observed in multiple exposures, which are 15 min in duration and can be distributed across several nights. All good data for a particular galaxy are co-added together during the spectroscopic data reduction process. Fits are made to each observed spectrum using a number of templates and combinations of templates evaluated for all allowed redshifts. A point estimate redshift is then obtained by maximizing the likelihood. The estimated redshift distributions used in our theory modelling of the CMASS and eBOSS samples (the shaded curves in the top panel of Fig. 2) are, then, histograms of these point estimates. Note that in making these histograms, we apply the METACALIBRATION weights described in Section 3.2.

3.4.2 Photometric redshifts

Unlike with the SDSS samples, we do not have spectroscopic redshifts for our DES redMaGiC samples. Rather, for each galaxy, we have a redshift PDF, which is obtained using DES photometry. The redMaGiC algorithm (Rozo et al. 2016) relies on the fact that red sequence galaxies have a relatively tight magnitude–colour–redshift re-

lation, which can be calibrated using overlapping spectroscopic data (Cawthon et al. 2022). Individual galaxies are fit using the process described in Section 2.3. Where it is necessary to have point redshift estimates (e.g. for the binning in Section 3.6), we use the value that maximizes the likelihood, $z_{redmagic}$. We follow Dark Energy Survey Collaboration (2022) and estimate the ensemble $n(z)$s by stacking samples from the full non-Gaussian redshift PDFs (see also Porredon et al. 2021 for discussion). These are shown in Fig. 2 (upper panel).

In addition to the $n(z)$ for each sample and point estimates themselves, our modelling also requires an estimate for the per-galaxy redshift *uncertainty* as a function of redshift. In the cases of eBOSS and CMASS, the spectral resolution allows very precise redshift estimates, and so we can assume this to be negligible. In the case of redMaGiC we obtain error estimates using a representative subsample of the Y3 redMaGiC catalogues with spectra (see Pandey et al. 2022). Specifically, we divide the sample into bins of z_{spec}, and within each bin we evaluate the histogram $p([z_{samp} - z_{spec}]/[1 + z_{spec}])$, where z_{samp} are the PDF draws used in estimating the $n(z)$ above. Since we have four PDF samples per galaxy[1], we compute four histograms, and average them, giving us a noisy estimate for the redshift error in the bin centred on z_{spec}. We find that the histograms are well approximated by a Gaussian distribution, and so we fit each histogram to obtain a width σ_z. This process leaves us with $\sigma_z(z)$, an estimate for the redshift scatter as a function of redshift, which we interpolate and incorporate into the modelling described in Section 4.4.2. Although there is some slight variation with redshift, a constant $\sigma_z/(1 + z) \sim 0.01$ is a reasonable approximation, with $\sigma_z/(1 + z) < 0.02$ over the range $z = 0.2$–1.1 (see Porredon et al. 2021, and in particular their fig. 1).

3.5 Luminosities, colours, and absolute magnitudes

To obtain rest-frame absolute magnitudes for our galaxy catalogues, we first convert the best-fitting r-band fluxes from METACALIBRATION to apparent magnitudes, $r = 30 - 2.5\log f_r$. The corresponding absolute magnitude is then given by

$$M_r^i = r^i - 5\left(\log D_l(z^i) - 1\right) - K(z^i), \tag{2}$$

where the index i denotes a galaxy, z^i is the best point estimate redshift for that galaxy, and D_l is the corresponding luminosity distance. Note that D_l *is in units of* $pc\,h^{-1}$. We calculate a $k + e$-correction K for each galaxy based on the redshift using the stellar synthesis models of Bruzual & Charlot (2003). In brief, we employ two models: one assuming a passively evolving spectral energy distribution, and the other passive but with a single instantaneous burst of star formation at $z = 9.84$. These models give us predicted colours and a $k + e$-correction as a function of z. For each galaxy i, we then compare the observed $r - i$ colour with the model predictions; if the observed colour is redder than the predicted one from the passive model, we use that model. If it is bluer than the one from the passive plus star formation burst model, then we use that one. Otherwise, we calculate a weighted average of the two $k + e$-corrections. In all cases, we correct the magnitudes to $z = 0$. Note that these star

---

[1]This a product of how the redMaGiC lens redshift distributions were estimated for DES Y3. Instead of saving the full non-Gaussian PDF as a function of redshift, four Monte Carlo samples were saved per galaxy. These were then combined to give the estimate for the distribution for the overall sample $n(z)$. Note this is different from the methodology for the fiducial Y3 lens sample, MAGLIM. See e.g. section C1 of Pandey et al. (2022) and section B2 of Dark Energy Survey Collaboration (2022) for details.

Figure 3. The r-band luminosity distributions of the red samples used in this work. We divide each sample, with the exception of eBOSS LRGs, into roughly equal number bins in luminosity, as indicated by the shading. See Section 3.5 for details.
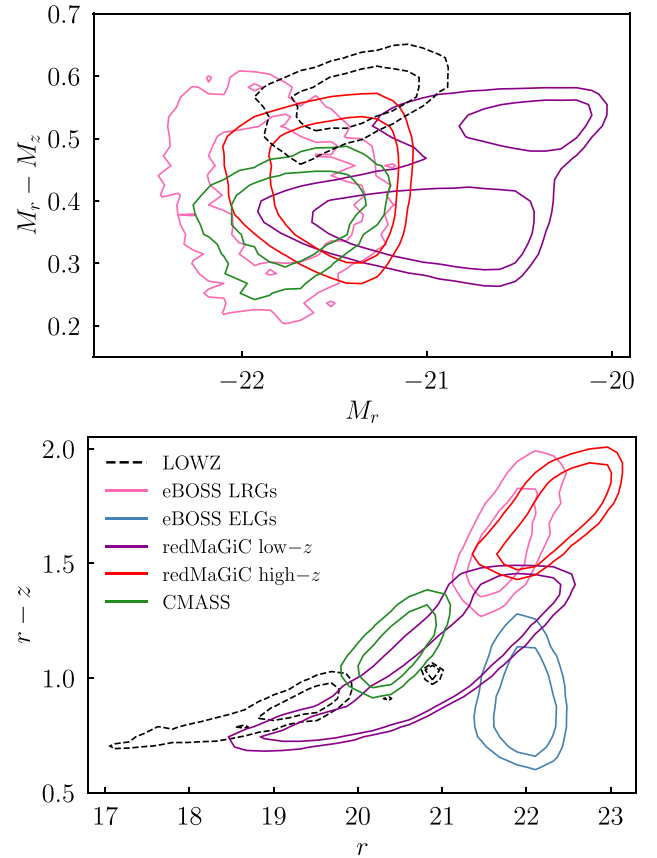

Figure 4. **Top:** The rest-frame colour–magnitude diagram for the samples used in this work. The quantities here are $k + e$-corrected magnitudes in the DES filters. In the main galaxy samples (solid lines), these are estimated using METACALIBRATION fluxes. For LOWZ, which is kept as a validation sample, we use the pre-computed absolute magnitudes described in Singh et al. (2015). The contours are defined relative to the peak, at $0.5 \times$, and $0.25 \times$ the maximum density. **Bottom:** The equivalent colour–magnitude space, but using apparent magnitudes. Note that we do not have $k + e$-corrections for eBOSS ELGs, so they appear in the lower panel only.

formation models are designed to describe elliptical galaxies, and we do not apply them to our ELG sample.

The above procedure is based on the assumption that the overall stellar population in a given galaxy sample is a mixture of two sub-populations, such that the observed colours are a linear combination of the colours of those components; these observed colours are therefore subject to a linear combination of the associated $k + e$-corrections. Note that in practice the templates do not differ enormously over the redshift range of our samples. Indeed, we recompute the $k + e$-corrections using the two models separately, and find no significant change in the distributions shown in Fig. 3.

The luminosity relative to a pivot L_0 is then given by $\log{(L_r/L_0)} = -(M_r - M_0)/2.5$, where M_0 is a fixed reference magnitude; we adopt

a value $M_0 = -22$ for the sake of comparability with previous results. For the purposes of constraining trends in alignment properties, we subdivide our red galaxy samples into luminosity bins. These are shown in Fig. 3, and are defined such that they contain roughly equal numbers of galaxies (with the exception of the bright end of redMaGiC high-z, where the signal-to-noise ratio was sufficient to allow us to further split the highest L bin in two.). Between them, our four samples cover a range of roughly $\log{(L_r/L_0)} = [-0.9, 0.4]$, with redMaGiC low-$z$ in particular providing excellent coverage of the fainter end. We also show the rest-frame colour–magnitude diagram for these red samples (as well as LOWZ) in the top panel of Fig. 4. As can be seen here, although we group these samples together as 'red', there is some significant variation in colour at fixed luminosity. We will return to this in the context of our main results in Section 9. The lower panel shows the same colour–magnitude space, but using apparent magnitudes. Here, the distributions are relatively elongated, primarily due to the colour–redshift degeneracy; that is, a galaxy of given rest-frame magnitude and colour observed at high redshift will appear both fainter and redder than the same object observed at low redshift.

3.6 Two-point correlations

Following a number of previous IA studies, our primary measurements are constructed using a modified Landy–Szalay estimator (Landy & Szalay 1993). For two-point galaxy clustering, this has the form

$$\xi_{gg}(r_\mathrm{p}, \Pi) = \frac{(D - R_D)(D - R_D)}{R_D R_D}. \tag{3}$$

The measurement is made on a grid of line-of-sight and perpendicular (comoving) separation, r_p and Π. For a particular sample of galaxies, we have a density tracer catalogue and a second catalogue of random points tracing the same footprint and redshift distribution. DD, $R_D D$, and $R_D R_D$ are the weighted counts of galaxy–galaxy, galaxy–random, and random–random pairs in a given bin of r_p and Π. To reduce shot noise, the randoms R_D are oversampled relative to the actual data by a factor of >10.

Similarly, one can estimate the shape–density cross correlation:

$$\xi_{g+}(r_\mathrm{p}, \Pi) = \frac{S_+(D - R_D)}{R_D R_S}, \tag{4}$$

where again D represents the density sample, and R_S and R_D are randoms matching the shape and density samples, respectively. The shape–shape correlation is constructed in a similar way:

$$\xi_{++}(r_\mathrm{p}, \Pi) = \frac{S_+ S_+}{R_S R_S}. \tag{5}$$

We also define

$$S_+ S_+ \equiv \sum_{\alpha, \beta; \beta \neq \alpha} e_+(\beta | \alpha) e_+(\alpha | \beta), \tag{6}$$

$$S_+ D \equiv \sum_{\alpha, \beta; \beta \neq \alpha} e_+(\beta | \alpha). \tag{7}$$

Here, the sum runs over galaxies (or random points) at a given separation drawn from the two catalogues; $e_+(\beta | \alpha)$ is the tangential ellipticity component of galaxy i, defined by the separation vector with galaxy β. The quantity $S_+ R_D$ is the same, but using the positions of random points in place of galaxies. One can write down a set of analogous equations for $\xi_{g\times}$, $\xi_{+\times}$, and $\xi_{\times\times}$, which are identical to the above, but with galaxy ellipticities rotated by 45 deg. Any astrophysical contribution to these, however, is expected to be negligible (due to parity arguments) and for this reason they are commonly used for null testing.

The $\xi(r_\mathrm{p}, \Pi)$ measurements are then projected along the line of sight as

$$w_{ab}(r_\mathrm{p}) = \int_{-\Pi_{\max}}^{\Pi_{\max}} \xi_{ab}(r_\mathrm{p}, \Pi) \mathrm{d}\Pi. \tag{8}$$

We use TREECORR[2] (Jarvis, Bernstein & Jain 2004) for all two-point measurements with bin_slop = 0.0. We use 20 logarithmically spaced bins in r_p, over the range 0.1–200 h^{-1} Mpc. For the line-of-sight binning we set $\Pi_{\max} = 100$ h^{-1} Mpc, with 20 linearly spaced bins between $\pm\Pi_{\max}$. The resulting data vectors are shown in Fig. 5. For parts of our analysis, we also use data vectors in bins of luminosity. These are shown in Fig. 6. The shaded regions here indicate physical scales excluded from our fiducial analysis. Further discussion of the fits to these data can be found in Section 6. The choice of $\Pi_{\max}$ is driven by signal-to-noise considerations in the photometric samples, and follows Singh et al. (2015). We also note that since we are including a galaxy–galaxy lensing (Section 4) term

---

[2]http://rmjarvis.github.io/TreeCorr (version 4.1.1)

---

in our model, our choice here is not limited by the need to suppress such contributions.

The DES shape catalogues use a technique called METACALIBRATION for accurately inferring an underlying shear signal from galaxy shape estimates. We apply response corrections in exactly the same way as in the DES Y3 cosmology analyses (Gatti et al. 2021). That is, we have a mean scale-independent factor $\langle R \rangle$, which is applied at the level of the two-point measurements as

$$\xi_{g+} \to \frac{1}{\langle R \rangle} \xi_{g+}, \quad \xi_{++} \to \frac{1}{\langle R \rangle^2} \xi_{++}, \tag{9}$$

where the angular brackets indicate an average over galaxies, and $\langle R \rangle = \langle R_\gamma + R_\mathrm{sel} \rangle$, or the sum of a shear response and a selection term for the shape sample in question.

Note that R_sel corrects only for shape catalogue selection cuts. Since we do not have either redMaGiC re-runs or the eBOSS selection on sheared images, any bias induced by the basic sample selection is not included in this correction. That said, the estimated selection response for an early/late split of the DES Y1 catalogues was found to be of the order of 10^{-4}, which is easily subdominant to our uncertainties (see Samuroff et al. 2019 Section 4.1). We thus judge it safe to ignore this missing correction for the purposes of our analysis.

4 MODELLING AND ANALYSIS CHOICES

The following section sets out our theoretical modelling choices. Our aim here is to connect an observed joint data vector, $w_{gg} + w_{g+} + w_{++}$, with underlying physical quantities, which can be calculated from theory. Each of these data vector components is a combination of two observable fields $\hat{\delta}_g$ and $\hat{\gamma}$, or the observed galaxy overdensity in counts and shapes. If we assume the former is the sum of contributions from gravitational clustering and magnification, $\hat{\delta}_g = \delta_g + \delta_\mu$, and the latter is the combination of intrinsic shape alignments and lensing, $\hat{\gamma} = \gamma^I + \gamma^G$, we have a total of four correlations contributing to each observable. The sections below will set out how we evaluate these model ingredients

In reality, the observed shear is weighted by the overdensity of shape galaxies, $\hat{\gamma} \to (1 + \delta_{g,S})\hat{\gamma}$ (see e.g. Hirata & Seljak 2004 equation (6)). This contributes an additional IA term, which is explicitly included in TATT (although not NLA; Blazek et al. 2019). For conciseness, we absorb this factor into the definition of γ^I when discussing the TATT model. Note that since the overdensity weighting applies to the total observed shear, not just the intrinsic component, it also gives rise to terms that scale as $\gamma^G \times \delta_{g,S}$, an effect known as source clustering. An analogous effect called source magnification enters in a similar way. These extra terms, however, are expected to be small at the level of projected observables, and so we neglect them here (see e.g. Krause et al. 2021 section 5B and Schmidt et al. 2009 for discussion).

We start in Sections 4.1–4.3 by describing how we calculate the 3D power spectra that enter each of our models. Section 4.4 then sets out how these are combined and projected to give predictions for the observable correlations. We discuss how the covariance matrix of the data is estimated in Section 4.5. Finally, Section 4.6 discusses how we choose a set of scale cuts, which restrict our analysis to the regime where our model is thought to be sufficient.

When it is necessary to assume a background cosmology, we use a flat Lambda cold dark matter (ΛCDM) model $\mathbf{p}_\mathrm{cos} = (\Omega_\mathrm{m},$

Figure 5. IA correlations measured from DES Y3 and eBOSS. The rows show (top/bottom) position–shape and shape–shape correlations, as defined in the text. The shaded bands indicate scales excluded from our NLA (light grey) and TATT (darker grey) fits. The strong scale dependence in the measured w_{g+} on small scales is thought to arise from a combination of one-halo IA correlations and non-linear galaxy bias, both of which become significant at $r_p < \sim 1\,h^{-1}$ Mpc. The solid lines are the best-fitting NLA predictions for each data set at $r_p > 6\,h^{-1}$ Mpc. We also show the sum of the terms arising from lensing and magnification separately (the dashed lines); for clarity, this sub-dominant contribution is scaled up by a factor of 10. Note that the vertical axes vary between panels.

Figure 6. The measured galaxy–shape correlations from redMaGiC and eBOSS LRGs. The columns (left to right) show bins of $k + e$-corrected r-band luminosity (as defined in Table 4). The rows show (top to bottom, colours the same as in Fig. 5) redMaGiC low-z (purple), redMaGiC high-z (red), CMASS (green), and LRGs (pink). Note that the luminosity bins are **not** the same in the three cases; we use the LX notation for convenience, but the bin edges and widths are defined for a particular sample (shown in the upper left of each panel; see also Section 3.5 for discussion of how the luminosity bins are defined). We also show the best-fitting NLA model prediction for each measurement (solid line). As above, points within the shaded grey regions are excluded from the fits. Note that we do not fit the TATT model to our luminosity-binned measurements and so, unlike in Fig. 5, there is only one set of grey bands shown here.

Ω_b, σ_8, h, n_s, $\Omega_\nu h^2$) = (0.3, 0.048, 0.82, 0.69, 0.97, 0.00083[3]). Although our results are not strongly cosmology dependent, we do quantify the impact of this choice in Section 6. The linear matter power spectrum is computed using the Boltzmann code CAMB[4] (Lewis, Challinor & Lasenby 2000), with non-linear corrections using HALOFIT (Takahashi et al. 2012) (this choice is not expected to have a significant bearing on our results, given our scale cuts; see the tests in Section 4.6). Parameter inference is performed within COSMOSIS[5] (Zuntz et al. 2015) using the MULTINEST nested sampling algorithm[6] (Feroz et al. 2019).

4.1 IA power spectra

To model the power spectra of the intrinsic alignment GI and II signals (respectively, the correlation between γ^I and γ^G, and γ^I with itself), we use the TATT model (Blazek et al. 2019). The basic idea is that the intrinsic shape field γ^I can be expressed as an expansion in powers of the background tidal field s and matter overdensity δ:

$$\gamma_{ij}^I = C_1 s_{ij} + C_{1\delta}\delta s_{ij} + C_2 \sum_k s_{ik}s_{kj} + \cdots \qquad (10)$$

Note that while δ is a scalar at any given position **x**, γ^I and s are 3×3 tensors. The above expansion can be propagated to the two-point level to give expressions for P_{GI} and P_{II} (see Blazek et al. 2019). Our implementation of the TATT model is identical to that of Secco et al. (2022) and Krause et al. (2021). We refer the reader to those papers for specifics, and in particular Section D2 and equations (21) and (22) of Secco et al. (2022) for the full expressions.

The TATT model has three free parameters, which we refer to as A_1, A_2, and b_{TA}. One can also parametrize the redshift dependence of all the contributions, if desired, as in previous cosmological analyses. Since our individual samples do not have a particularly wide redshift range, however, this is not an especially useful thing to do in our case. On the other hand, one can look at the evolution *between* samples. Considering galaxies with similar colour and luminosity properties in Appendix E, we find no evidence for z evolution over a significantly wider range than the coverage of any one of our samples alone. The two amplitudes modulate the strength of IA contributions that are linear and quadratic in s:

$$C_1 = -A_1\frac{\rho_{crit}\Omega_m\bar{C}_1}{D(z)} \qquad C_2 = A_2\frac{5\rho_{crit}\Omega_m\bar{C}_1}{D^2(z)}, \qquad (11)$$

where $D(z)$ is the linear growth factor, ρ_{crit} is the critical density, and $\bar{C}_1$ is a constant, which by convention is fixed to a value of $5 \times 10^{-14}M_\odot$ Mpc^2 h^{-2}.[7] The other parameter, b_{TA}, is known as the

---

[3]This corresponds to a total sum of the neutrino masses, $\sum m_\nu = 0.077$ eV.
[4]http://camb.info/
[5]https://bitbucket.org/joezuntz/cosmosis/wiki/Home; v1.6, 'des-y3' branch of cosmosis-standard-library, 'develop' branch of cosmosis
[6]v3.6; efficiency = 0.3, live_points = 500
[7]Note that there are versions of this equation in the literature that feature the rescaled growth factor $\bar{D}(z) = (1 + z)D(z)$ in place of the unweighted $D(z)$. They are, however, consistent with the formulation shown here. One can work back from e.g. the pre-factor in equation (18) of Hirata & Seljak (2004) by making the substitution $\bar{\rho} = \rho_{crit} \times \Omega_m(1 + z)^3$ and $a = 1/(1 + z)$. The factors of $(1 + z)$ cancel, and we are left with the left-hand expression in equation (11). Note that there is a subtle distinction in the normalization convention. Although Hirata & Seljak (2004) normalize $\bar{D}$ to equal 1 during the epoch of matter domination, more recent work (including this one and

Table 2. The free parameters and priors for the models discussed in this paper. The upper three rows are IA model parameters, while the lower two describe galaxy bias. We include two sets of priors here: one for our TATT model analyses (which extends to $r_p > 2\ h^{-1}$ Mpc), and one for NLA ($r_p > 6\ h^{-1}$ Mpc).

Parameter	Description	TATT Prior	NLA Prior
A_1	Lin. IA amplitude	U[−8, 8]	U[−8, 8]
A_2	Quadratic IA amplitude	U[−8, 8]	δ[0]
b_{TA}	Density wt. coefficient	U[−6, 6]	δ[0]
b_1	Lin. galaxy bias	U[0, 3]	U[0, 3]
b_2	Second-order galaxy bias	U[−3, 3]	U[−3, 3]

density weighting coefficient, and controls the size of the $C_{1\delta}$ term above as $C_{1\delta} = b_{TA}C_1$.

We also consider a nested subspace of the full TATT model. The simplest subspace, known as the NLA model (Bridle & King 2007; Hirata et al. 2007), has only one free parameter, A_1. The γ^I field is assumed to be purely linear in the tidal field (effectively setting $C_2 = 0$, $C_{1\delta} = 0$ in equation (10)), and so the IA power spectra have the same shape as the non-linear matter power spectra, but with a scaling factor applied.

In all samples considered, we vary the IA parameters with wide flat priors, as given in Table 2.[8]

4.2 Galaxy power spectrum

Galaxy bias, or the mapping between the matter and galaxy overdensity fields, is an important source of uncertainty in any analysis that relies on galaxy–shape correlations. Similar to γ^I in Section 4.1, one can expand the galaxy overdensity in terms of δ (McDonald 2006; Baldauf et al. 2010; Saito et al. 2014):

$$\delta_g = b_1\delta + \frac{1}{2}b_2\left(\delta^2 - \langle\delta^2\rangle\right) + \frac{1}{2}b_{s^2}\left(s^2 - \langle s^2\rangle\right) + b_{3nl}\psi + \cdots \qquad (12)$$

Here, ψ is the sum of several different third-order terms with the same scaling (see Saito et al. 2014). On large enough scales, it is often sufficient to assume a simple linear relation $\delta_g = b_1\delta$; in this case the galaxy power spectrum is simply

$$P_{\delta_g}(k, z) = b_1^2 P_\delta(k, z), \qquad (13)$$

where the galaxy bias b_1 depends on the galaxy sample, but is independent of wave number. P_δ is the non-linear matter power spectrum. Unfortunately, we see evidence of the need for a more sophisticated approach in some of our samples. This is discussed further in Section 4.6, where we see that redMaGiC high-z and eBOSS LRGs favour a more complicated bias model, even on relatively large scales.

Using equation (12) one can write down a more complete expression for P_{δ_g} (e.g. in Krause et al. 2021 equation 38). Our fiducial model for the galaxy power spectrum includes all terms in the expansion above, for which we use the implementation in FASTPT

Bridle & King 2007) have tended to impose $D(z = 0) = 1$. This has a roughly 30 per cent impact on the magnitude of C_1.
[8]Note that these differ slightly from those used in Dark Energy Survey Collaboration (2022). Although the DES Y3 priors were chosen to be uninformative for that particular sample, we are considering significantly different (often much redder) populations of galaxies. We thus opt to allow for more extreme IA values.

(McEwen et al. 2016). Assuming co-evolution, however, we can reduce the number of free parameters to two, with $b_{s^2} = -4/7(1 - b_1)$ and $b_{3nl} = b_1 - 1$ (see Saito et al. 2014, and also Pandey et al. 2020 and Krause et al. 2021 for further discussion). For all samples, we marginalize over these galaxy bias parameters with wide flat priors $b_1 = [0.05, 3]$, $b_2 = [-3, 3]$.

For the power spectra entering w_{g+} we assume linear bias (despite using equation (12) for P_{δ_g}):

$$P_{\delta_g I}(k, z) = b_1 P_{GI}(k, z), \tag{14}$$

and similarly,

$$P_{\delta_g \delta}(k, z) = b_1 P_\delta(k, z), \tag{15}$$

where b_1 is the same as in equation (12) above. In principle, non-linear galaxy bias, and also various cross terms between TATT parameters and higher order bias are expected to contribute to w_{g+}. In all cases considered here, however, w_{g+} has significantly lower signal-to-noise ratio than the equivalent galaxy–galaxy correlations, and thus the latter dominate the fits for galaxy bias. To within the level of uncertainty the TATT model is able to sufficiently describe the potential impact of correlations between non-linear galaxy bias and IA through the free b_{TA} parameter (see the similarity of these non-linear terms in Blazek et al. 2015). For fits using the NLA model, we exclude scales where non-linear bias correlations are significant (see Section 4.6 for discussion of how the scale cuts are chosen). Although b_{TA} cannot absorb the non-linear bias contributions to the galaxy–galaxy lensing signal $\delta_g G$ so easily, this term is subdominant on all scales and for all samples (~ 5–10 per cent of the total w_{g+} signal; see Section 6.3). We test the impact on the galaxy–galaxy lensing signal by substituting the non-linear bias expression from equation (12) into $P_{\delta_g \delta}$, in place of the approximation in equation (15). Using the best fit bias parameters for each of our samples, we find a roughly 10 per cent change in $w_{\delta_g G}$ on scales $<6h^{-1}$ Mpc; compared to the full signal including IA, however, the impact is at the sub-percentage level. Implementing a fully consistent non-linear model is a work in progress, but we do not expect this to have a significant impact given the statistical uncertainties in current data sets.

4.3 Magnification and lensing power spectra

As well as contributions due to galaxy clustering and intrinsic shape correlations, magnification can have an effect on direct IA measurements. Its impact is to alter the observed galaxy number density in a patch of sky as $\hat{\delta}_g = \delta_g + \delta_\mu$. Similarly, the observed shear in a set of galaxies has both an IA contribution, and one from cosmological lensing: $\hat{\gamma} = \gamma^I + \gamma^G$. At the two-point level, one has two additional terms in the gg correlation (galaxy-magnification and magnification-magnification; $\delta_g \delta_\mu$ and $\delta_\mu \delta_\mu$), and two in the galaxy–shape correlation (magnification-intrinsic and magnification-lensing; $\delta_\mu \gamma^I$ and $\delta_\mu \gamma^G$). Similarly, w_{++} has contributions from the standard II and GI power spectra, but also a pure cosmic shear term $\gamma^G \gamma^G$. On large scales, the additional magnification power spectra are all related to galaxy and IA power spectra via magnification coefficients α (see Table 3, and Joachimi & Bridle 2010; Joachimi et al. 2021; von Wietersheim-Kramsta et al. 2021; Elvin-Poole et al. 2023 for discussion).

A number of different methods for constraining α have been discussed in the literature. We describe how we estimate α for each density sample in Section 3.3. In short, our fiducial estimates are obtained by artificially perturbing the observed galaxy magnitudes (i.e. a flux-only estimate). For the two redMaGiC samples, we

Table 3. A summary of the various contributing terms to our observables w_{gg}, w_{g+}, and w_{++}. The kernel column lists all the possible combinations for $q_a^i q_b^j$ in equation (23) (where each q is either the lensing kernel g or the galaxy PDF p). For each one we show the kernel (either lensing efficiency or redshift distribution), and the relevant power spectrum included in the Limber integral. The pre-factors α are magnification coefficients, which are defined in Section 3.3. Note that in later sections we refer to these terms simply by their subscripts (e.g. $\mu\mu$ for magnification–magnification correlations).

Correlation	Kernel	Power Spectrum	Correlation Function
$\delta_g \delta_g$	$p^i p^j$	P_{δ_g}	w_{gg}
$\delta_\mu \delta_\mu$	$g^i g^j$	$4(\alpha^i - 1)(\alpha^j - 1)P_\delta$	w_{gg}
$\delta_\mu \delta_g$	$g^i p^j$	$2(\alpha^i - 1)P_{\delta_g \delta}$	w_{gg}
$\delta_g \gamma^I$	$p^i p^j$	$P_{\delta_g I}$	w_{g+}
$\delta_\mu \gamma^I$	$g^i p^j$	$2(\alpha^i - 1)P_{GI}$	w_{g+}
$\delta_g \gamma^G$	$p^i g^j$	$P_{\delta_g \delta}$	w_{g+}
$\delta_\mu \gamma^G$	$g^i g^j$	$2(\alpha^i - 1)P_\delta$	w_{g+}
$\gamma^I \gamma^I$	$p^i p^j$	P_{II}	w_{++}
$\gamma^G \gamma^I$	$g^i p^j$	P_{GI}	w_{++}
$\gamma^G \gamma^G$	$g^i g^j$	P_δ	w_{++}

have estimates from BALROG, which we use for validation (see Section 6.2.3).

4.4 Modelling projected correlation functions

4.4.1 Modelling spectroscopic data

Given power spectra from any model, one can convert into projected correlation functions of the sort discussed in Section 3.6 via Hankel transforms. In the case of perfect knowledge of individual galaxy redshifts (i.e. spectroscopic redshifts) one has:

$$w_{g+}^{ij}(r_p) = -\int dz \mathcal{W}^{ij}(z) \int \frac{dk_\perp k_\perp}{2\pi} J_2(k_\perp r_p) P_{\delta_g I}(k_\perp, z), \tag{16}$$

with the Roman indices indicating the two galaxy samples, and J_ν being a Bessel function of the first kind of order ν. The projection kernel is given by (see Mandelbaum et al. 2011's Appendix A)

$$\mathcal{W}^{ij}(z) = \frac{n^i(z)n^j(z)}{\chi^2(z)d\chi/dz} \times \left[\int dz \frac{n^i(z)n^j(z)}{\chi^2(z)d\chi/dz} \right]^{-1}. \tag{17}$$

In the above, $n^i(z)$ is the estimated redshift distribution for sample i, and $\chi(z)$ is the comoving line-of-sight distance corresponding to a redshift z. The other two-point correlations follow by analogy as:

$$w_{gg}^{ij}(r_p) = b_g^i b_g^j \int dz \mathcal{W}^{ij}(z) \int \frac{dk_\perp k_\perp}{2\pi} J_0(k_\perp r_p) P_{\delta_g}(k_\perp, z), \tag{18}$$

and

$$w_{++}^{ij}(r_p) = \int dz \mathcal{W}^{ij}(z)$$
$$\times \int \frac{dk_\perp k_\perp}{2\pi} \left[J_0(k_\perp r_p) + J_4(k_\perp r_p) \right] P_{II}(k_\perp, z). \tag{19}$$

In each case, the theory prediction amounts to a projection of a power spectrum along the redshift axis, and then a Bessel integral.

4.4.2 Modelling IA correlations in the presence of photo-z error

When dealing with spectroscopic galaxy samples, one can in general safely assume that the associated redshift error is much smaller than the distance scales of interest. This assumption does not hold for

photometric samples such as redMaGiC, which means the modelling is slightly more complicated. The impact of redshift error is to scatter galaxies along the line of sight; this in effect shuffles galaxies between Π bins and so redistributes power out along the line of sight. In principle the effect should wash out when integrating over a sufficiently large range in Π. In reality, however, one must choose finite Π limits, and widening the integration range to large separations is not necessarily desirable, since it can degrade the signal-to-noise ratio of the measurement. This leads to an overall suppression of the measured correlations due to photo-z scatter. Another impact of photo-z error is that it can boost additional (non-IA) signals. That is, galaxy pairs allocated to small Π bins may actually be physically quite distant. Such pairs carry little local Π signal, but they do tend to increase the lensing and magnification contributions. The consequence of this is that one must account for the Π cut-off in the model. To do so we follow the method set out in Joachimi et al. (2011), of which we provide an outline below.

To begin, we compute angular spectra from the IA and galaxy–galaxy power spectra. Incorporating all of the magnification and lensing contributions to number counts and shear, one has:

$$C^{ij}_{\hat{\delta}_g \hat{\delta}_g} = C^{ij}_{\delta_g \delta_g} + C^{ij}_{\delta_\mu \delta_\mu} + C^{ij}_{\delta_\mu \delta_g} + C^{ij}_{\delta_g \delta_\mu} \tag{20}$$

$$C^{ij}_{\hat{\delta}_g \hat{\gamma}} = C^{ij}_{\delta_g \gamma^I} + C^{ij}_{\delta_\mu \gamma^I} + C^{ij}_{\delta_g \gamma^G} + C^{ij}_{\delta_\mu \gamma^G} \tag{21}$$

$$C^{ij}_{\hat{\gamma}\hat{\gamma}} = C^{ij}_{\gamma^I \gamma^I} + C^{ij}_{\gamma^I \gamma^G} + C^{ij}_{\gamma^G \gamma^I} + C^{ij}_{\gamma^G \gamma^G}, \tag{22}$$

where the subscripts δ_μ, γ^I, δ_g and γ^G indicate magnification, intrinsic shape, gravitationally induced galaxy overdensity and gravitational shear[9] Implicitly, the Π term here is the E-mode autocorrelation, $C^{\mathrm{EE}}_{\gamma^I \gamma^I}$. In principle, one could also include $C^{\mathrm{BB}}_{\gamma^I \gamma^I}$, which can be calculated assuming a particular IA model. We do not include this in our model because (a) typically any IA induced B modes are small (Hirata & Seljak 2004; Blazek et al. 2019) and (b) they contribute only to w_{++}, where the signal-to-noise ratio of our measurements is low. The Limber integrals used to compute each of the angular power spectra then have the form:

$$C^{ij}_{ab}(\ell|z_1, z_2) = \int_0^{\chi_{\mathrm{hor}}} \mathrm{d}\chi' \frac{q^i_a\left(\chi'|\chi(z_1)\right) q^j_b\left(\chi'|\chi(z_2)\right)}{\chi'^2}$$
$$\times P_{ab}\left(k = \frac{\ell + 0.5}{\chi'}, z(\chi')\right). \tag{23}$$

The kernel q is either the lensing efficiency g, or the error distribution p, according to Table 3. The power spectrum P_{ab} corresponding to a given $C(\ell)$ are also shown in Table 3. Here, $p_i(z'|z)$ is the conditional probability distribution for the true redshift of a galaxy from sample i, which has a best-estimate redshift at z. The estimates for p at any given z are obtained using the method described in Section 3.4. Note that this is different from the more common form of the Limber integral in the context of cosmological lensing, which uses the ensemble redshift distribution $n(z)$, not the per-galaxy PDF. One can then transform from harmonic to angular space as follows:

$$\xi^{ij}_{ab}(\theta|z_1, z_2) = \frac{1}{2\pi} \int_0^\infty \mathrm{d}\ell \, \ell J_\nu(\ell\theta) C_{ab}(\ell|z_1, z_2), \tag{24}$$

where the order of the Bessel function ν depends on the type of correlation ($\nu = 0$ for $ab = \hat{\delta}_g \hat{\delta}_g$, $\nu = 2$ for $ab = \hat{\delta}_g \hat{\gamma}$ or $\nu = (0, 4)$ for $ab = \hat{\gamma}\hat{\gamma}$), as in Section 4.4.1. As argued in Joachimi et al. (2011), from here one can obtain the photometric correlation function $\xi^{ij}_{ab}(r_{\mathrm{p}}, \Pi, z_{\mathrm{m}})$ using a simple coordinate transformation (see equation A11 in that paper, which also defines $z_{\mathrm{m}} = (z_1 + z_2)/2$). Finally, the projected correlation function as a function of perpendicular physical separation is expressed as,

$$w^{ij}_{ab}(r_{\mathrm{p}}) = \int_{-\Pi_{\mathrm{max}}}^{\Pi_{\mathrm{max}}} \mathrm{d}\Pi \int \mathrm{d}z_{\mathrm{m}} \mathcal{W}^{ij}(z_{\mathrm{m}}) \xi^{ij}_{ab}(r_{\mathrm{p}}, \Pi, z_{\mathrm{m}}) \tag{25}$$

With these ingredients, the recipe for generating a theory prediction for the cross correlation between photometric samples i and j is as follows.

  (i) Choose an initial value of Π and z_{m}. Use Joachimi et al. (2011)'s equation A11 to obtain z_1 and z_2, and evaluate the per-galaxy error distributions for the two samples at these redshift values.

  (ii) Carry out the Limber integral in equation (23) with these error distributions to obtain $C(\ell|z_1, z_2)$.

  (iii) Carry out the Hankel transform in equation (24) with the appropriate Bessel kernel to obtain $\xi(\theta|z_1, z_2)$.

  (iv) Perform the coordinate transform, such that $\xi(\theta|z_1, z_2) \rightarrow \xi(r_{\mathrm{p}}|\Pi, z_{\mathrm{m}})$.

  (v) Repeat the above steps with varying Π and z_{m}, to give a three-dimensional grid $\xi(r_{\mathrm{p}}, \Pi, z_{\mathrm{m}})$.

  (vi) Integrate over the redshift kernel $\mathcal{W}(z)$ and then over line-of-sight separation with the appropriate Π_{max} to obtain $w_{ab}(r_{\mathrm{p}})$.

We confirm that our implementation of this method returns the same results as equations (16)–(19) in the limit of narrow photo-z distributions and wide Π bounds. We also verify that, with a matching cosmology and set of input parameters, our modelling code can reproduce fig. 5 from Joachimi et al. (2011). Our fiducial modelling set-up is to use the steps above to predict w_{g+} and w_{++}.

For w_{gg}, however, it is not sufficient to assume RSDs have negligible impact (see Appendix A and Fig. A1). For this reason, we instead choose to use a sum over multipoles to obtain the anisotropic galaxy–galaxy correlation ξ_{gg}, which we then integrate over Π (equation (A1)-(A6)). We do, however, still need to account for lensing, magnification and photo-z suppression. Unlike with w_{g+}, where the combined impact of these effects are seen to have some non-trivial scale dependence, this is much less true for w_{gg}; using the recipe set out above, we generate theory data vectors for each sample with and without photo-z scatter, lensing, and magnification, finding that correction factor, $a(r_{\mathrm{p}}) = w^{\mathrm{full}}_{gg}/w_{gg}$, is flat with r_{p} to good approximation over scales $2 < r_{\mathrm{p}} < 70h^{-1}$ Mpc. Given this, we derive a single multiplicative factor for each sample, which we apply to the theory predictions as $w_{gg} \rightarrow a w_{gg}$. We obtain $a_{\mathrm{RMH}} = 0.83$ and $a_{\mathrm{RML}} = 0.87$ for our two redMaGiC samples, respectively.

4.5 Covariance matrix

We estimate the covariance of our data using an analytic prescription. This approach has a number of advantages over data-based estimators such as jackknife, which have been widely used in the past (Hirata et al. 2007; Joachimi et al. 2011; Mandelbaum et al. 2011; Singh et al. 2015; Fortuna et al. 2021b). For example, it can be used on large scales where jackknife breaks down, and it is unaffected by noise in the data. Note, though, we are assuming here that the covariance of our data is dominated by the Gaussian component, and

---

[9]Note that in the following sections, when discussing the various contributions to our observables, we will refer to some of these terms by their subscripts (i.e. μ for magnification, I for intrinsic shapes, G for lensing-induced shear). This helps to simplify the notation by avoiding too many double-subscripts and makes the discussion clearer later on.

any connected 4pt and super sample covariance contributions are negligible for our purposes (see e.g. Takada & Jain 2009; Takada & Hu 2013).

The covariance is assumed to be dominated by two components: one from cosmic variance, and one from shape and shot noise $\mathrm{Cov} = \mathrm{Cov}^{\mathrm{CV}} + \mathrm{Cov}^{\mathrm{SN}}$. For any two elements of our data vector in scale bins centred on $r_{\mathrm{p},m}$ and $r_{\mathrm{p},n}$, the cosmic variance part is given by:

$$
\begin{aligned}
\mathrm{Cov}^{\mathrm{CV}} & \left[w_{ij}(r_{\mathrm{p},m}) w_{kl}(r_{\mathrm{p},n}) \right] \\
& = \frac{1}{\mathcal{A}(z_c)} \int \frac{k \mathrm{d}k}{2\pi} \Theta_{ij}(k r_{\mathrm{p},m}) \Theta_{kl}(k r_{\mathrm{p},n}) \\
& \times \left[P_{ik}(k) P_{jl}(k) + P_{il}(k) P_{kj}(k) \right],
\end{aligned}
\tag{26}
$$

where the lower indices define the tracer type (i.e. g or $+$). The term $\Theta_{ij}(x)$ is a Bessel function of the first kind (or a sum of two); specifically J_2, J_0 and $J_0 + J_4$ for $ij = g+$, gg and $++$ respectively. The power spectra are $P_{\delta_g \mathrm{I}}$ for $ij = g+$, P_{δ_g} for gg and P_{II} for $++$. Note that we do not include secondary contributions from magnification and lensing, but this is not expected to significantly change our results. The pre-factor $\mathcal{A}$ is the projected comoving area of the footprint (including masking), at a characteristic redshift z_c.

The noise contribution is simply given by (Schneider et al. 2002; Hu & Jain 2004; Joachimi & Bridle 2010):

$$
\mathrm{Cov}^{\mathrm{SN}} \left[w_{g+}(r_{\mathrm{p},m}) w_{g+}(r_{\mathrm{p},n}) \right] = \delta_{mn} \frac{\sigma_\epsilon^2}{N_{\mathrm{p}}^{g+}},
\tag{27}
$$

$$
\mathrm{Cov}^{\mathrm{SN}} \left[w_{gg}(r_{\mathrm{p},m}) w_{gg}(r_{\mathrm{p},n}) \right] = \delta_{mn} \frac{1}{N_{\mathrm{p}}^{gg}},
\tag{28}
$$

$$
\mathrm{Cov}^{\mathrm{SN}} \left[w_{++}(r_{\mathrm{p},m}) w_{++}(r_{\mathrm{p},n}) \right] = \delta_{mn} \frac{\sigma_\epsilon^4}{N_{\mathrm{p}}^{++}},
\tag{29}
$$

for our three observables respectively. Since our measurements with TREECORR give us the number of galaxy pairs in each bin $N_{\mathrm{p}}(r_{\mathrm{p},m})$ without extra computational cost, we use these exact numbers here. The shape dispersion σ_ϵ is measured for each sample, using the Heymans et al. (2012) definition, and incorporating the correct response weighting (see Gatti et al. 2021 equation 13).

We perform initial fits using a preliminary covariance matrix, which we then replace with an updated version with the best-fitting values of A_1, b_1, and b_2 entering equation (26) above. Since w_{g+} is shape noise dominated for all samples and on all but the largest scales (and w_{++} entirely so, on all scales), this update makes little difference to the final IA parameter constraints.

We compare our analytic predictions with jackknife estimates in Appendix C, and find good agreement on scales $2 \, h^{-1} \, \mathrm{Mpc} < r_{\mathrm{p}} < 70 \, h^{-1} \, \mathrm{Mpc}$.

4.6 Scale cuts

We impose scale cuts on all three of our measured correlations when fitting, to mitigate model uncertainty. In brief, our minimum scales are $r_{\mathrm{p,\,min}} = (2, 6, 6) \, h^{-1} \, \mathrm{Mpc}$ for NLA and $r_{\mathrm{p,\,min}} = (2, 2, 2) \, h^{-1} \, \mathrm{Mpc}$ for TATT (where the ordering is w_{gg}, w_{g+}, w_{++}). For the latter two this is driven by the fact that we know our IA models start to break down on certain scales, and rely on assumptions that are valid only in specific regimes (NLA on scales above $\sim 5 - 10 \, h^{-1} \, \mathrm{Mpc}$, TATT down to $\sim 1-2 \, h^{-1} \, \mathrm{Mpc}$; Bridle & King 2007; Blazek et al. 2015). The motivation behind the w_{gg} scale cuts is discussed in more detail in Section 4.6.1.

We also impose an upper cut at $70 \, h^{-1} \, \mathrm{Mpc}$, a choice motivated by the null tests in Appendix B. This maximum scale is applied to all three correlations for all samples. Large-scale systematics, most prominently PSF modelling error, are known to modulate galaxy

number counts at large r_{p}, but are difficult to model analytically. We thus choose to remove the affected scales.

4.6.1 Galaxy clustering

Since scale cuts are designed to mitigate modelling uncertainty, the choice of $r_{\mathrm{p,\,min}}$ for w_{gg} is unavoidably connected to the choice of galaxy bias model. We first seek to test whether there are a set of cuts that will allow us to use a simple, scale-independent linear bias model. For each of our samples, one can estimate an effective bias

$$
b_g'(r_{\mathrm{p}}) = \sqrt{\frac{\hat{w}_{gg}(r_{\mathrm{p}})}{w_{\delta\delta}(r_{\mathrm{p}})}},
\tag{30}
$$

where $\hat{w}_{gg}$ is the measured projected galaxy–galaxy correlation. The matter–matter correlation in the denominator is the theoretical prediction, and so assumes a particular cosmology; we test the impact of switching between reasonably different cosmologies (specifically, the best-fitting values from DES Y1 and Planck 2018), and find our results are only very weakly sensitive to this choice. For each sample, we fit b_g' twice, once using a scale-independent constant $b_g'(r_{\mathrm{p}}) = b$, and again using a linear-exponential function $b_g'(r_{\mathrm{p}}) = a e^{-r_{\mathrm{p}}} + b$. Although this is not a physically motivated bias model, it has qualitatively the correct behaviour, increasing rapidly on small scales and converging to a constant on large scales. The exact form was motivated by fig. D1 of Samuroff et al. (2021), where the bias in IllustrisTNG is seen to scale roughly as $e^{-r_{\mathrm{p}}}$ plus a constant. In each case, we compute the Bayesian Information Criterion (Schwarz 1978),

$$
\mathrm{BIC} = k \log N_{\mathrm{pts}} + \chi^2,
\tag{31}
$$

where k is the number of model parameters (either 1 or 2, in the constant/log-linear cases, respectively) and N_{pts} is the number of data points included in the fit. The χ^2 for model M is computed using the full data covariance matrix, as $\chi_M^2 = [\hat{w}_{gg} - b_{g,M}' w_{\delta\delta}] \mathbf{C}^{-1} [\hat{w}_{gg} - b_{g,M}' w_{\delta\delta}]$. The difference $\Delta \mathrm{BIC}$ then gives us an indicator of which model is preferred by the data – that is, whether linear bias is sufficient, in a statistical sense, to describe the measured w_{gg}. We repeat this process using a range of lower scale cuts $r_{\mathrm{p,\,min}}$, resulting in the curves shown in Fig. 7. Although eBOSS ELGs and redMaGiC low-z appear to be relatively consistent with a linear bias model, even down to small scales, this is not true in all of our samples. The picture is slightly different in the case of redMaGiC high-z and eBOSS LRGs, with the latter in particular preferring the more complicated bias scaling for almost any choice of minimum scale.

The above test indicates that, at least for some of our samples, even at relatively large scales (above $6 \, h^{-1} \, \mathrm{Mpc}$), the linear bias approximation does not provide a good description of the data. Motivated by these findings, our fiducial scale cuts are as follows. We fit w_{gg} for all samples down to $r_{\mathrm{p,\,min}} = 2 \, h^{-1} \, \mathrm{Mpc}$, with a model that includes non-linear galaxy bias (as described in Section 4.4). At $2 \, h^{-1} \, \mathrm{Mpc}$ we are still well outside the one-halo regime, even for the largest objects in our samples, and so the perturbative expansion in equation (12) may still be sufficiently accurate. To help further validate this choice, we perform additional fits to w_{gg} alone, using very large scales ($> 10 \, h^{-1} \, \mathrm{Mpc}$) and linear bias. For each sample, we calculate the shift relative to the b_1 value obtained using the fiducial set-up, and verify that it is not sufficient to produce an appreciable bias in w_{g+}.

Figure 7. The difference in Bayesian Information Criterion between two simple models of effective galaxy bias, as a function of the lower scale cut $r_{p,\,\mathrm{min}}$ applied to w_{gg}. Negative values indicate that the linear bias assumption is justified by the data; positive values, on the contrary, indicate that the dark-matter-only correlation $w_{\delta\delta}$ with a scale-independent multiplicative factor is not a good representation of w_{gg}. The different colours show fits to the galaxy–galaxy correlation from various DES and eBOSS samples. As in Fig. 5, the shaded region shows scales excluded in our fiducial analysis set-up.

4.6.2 IA correlations

For w_{g+} and w_{++}, $r_{p,\,\mathrm{min}} = 6\ h^{-1}\,\mathrm{Mpc}$ for our NLA fits (see Section 6), primarily driven by IA modelling uncertainty. The TATT model allows us to push to slightly smaller scales, and so here we adopt $r_{p,\,\mathrm{min}} = 2\ h^{-1}\,\mathrm{Mpc}$.

We test the robustness of our chosen w_{g+} cuts to a number of unmodelled effects in Fig. 8. Specifically, we generate theory data vectors containing (a) a matter power spectrum contaminated with OWLS-like baryonic feedback (considered as an extreme scenario; see the next paragraph); (b) a one-halo IA signal, and (c) projection effects in the 3D correlation function. In each case, we choose a reference IA model $A_1 = 3.5$ (roughly the NLA best fit for redMaGiC in Section 6). We do this using both redMaGiC high-z and redMaGiC low-z redshift distributions, since these are representative of the range in z covered by our samples. The fractional differences in Fig. 8 are calculated relative to a reference data vector, which does not contain the contamination.

For (a), we follow Krause et al. (2021) in taking the OWLS-AGN scenario (Schaye et al. 2010; van Daalen et al. 2011) as an upper limit on the extremity of baryon feedback in the late-time matter power spectrum (see also Secco et al. 2022 fig. 5 and the accompanying discussion). The baseline matter power spectrum from CAMB is modified in such a way as to preserve the original cosmology but introduce high-k distortions which mimic the impact of baryons in the OWLS hydrodynamic simulations (Dark Energy Survey Collaboration 2016, equation 8). As we can see in Fig. 8, baryonic feedback is entirely negligible at $r_p > 6\ h^{-1}\,\mathrm{Mpc}$. Its impact increases rapidly in the intermediate ($2\ h^{-1}\,\mathrm{Mpc} < r_p < 6\ h^{-1}\,\mathrm{Mpc}$) window, but is still only $\sim 2-5$ per cent at $r_p = 2\ h^{-1}\,\mathrm{Mpc}$, which is well below the level of statistical error on these scales. If we take the TATT best fit from each redMaGiC sample (fit on scales $r_p > 2\ h^{-1}\,\mathrm{Mpc}$; see Section 6.2.2), we obtain $\Delta\chi^2_{\mathrm{RMH}} = 0.05$ and $\Delta\chi^2_{\mathrm{RML}} = 0.45$.

We carry out a similar exercise with small scale alignments. To get a rough gauge of the impact of one-halo contributions, we use the

fitting formulae provided by Schneider & Bridle (2010). We choose to update the overall amplitude of the model to the value found by Singh et al. (2015) ($a_h = 0.08$); since this matches LOWZ LRGs, which are somewhat brighter and redder than any of our samples, we expect this to be an upper estimate for the impact of intra-halo physics. Shown by the dotted line in Fig. 8, we again see the effect to be vanishing on scales $r_p > 6\ h^{-1}\,\mathrm{Mpc}$ and slightly larger but still subdominant to errors at $r_p > 2\ h^{-1}\,\mathrm{Mpc}$.

Finally, we also test the impact of a kind of projection effect that induces anisotropy in $\xi_{g+}(r_p, \Pi)$ (dashed lines in Fig. 8). First identified by Singh & Mandelbaum (2016), the idea is that galaxy alignments along the Π direction cannot be measured using shapes measured in 2D projected space; the result is a suppression of the observed alignment signal that scales as $f = r_p/(r_p^2 + \Pi^2)^{\frac{1}{2}}$. As one might intuitively expect from the geometry, f diverges from 1 as Π increases, at fixed r_p. Looking at Fig. 8 we can see that the impact is primarily at large r_p. To understand this, consider the fact that particular Π scales do not contribute equally to the projected correlation at all r_p (for an illustration of this, see Singh & Mandelbaum 2016 fig. 10c). That is, at $r_p = 1\ h^{-1}\,\mathrm{Mpc}$, even considerable suppression at large Π matters very little; that regime contributes almost nothing to the line-of-sight integral, since ξ_{g+} scales approximately as $1/r^2$, which at small r_p and large Π is essentially $1/\Pi^2$. At $r_p = 30\ h^{-1}\,\mathrm{Mpc}$, on the other hand, large Π scales contribute much more. Although the geometric suppression at a given Π is less important for this larger r_p value, the background scaling of ξ_{g+} dominates, and so the overall impact on w_{g+} is larger. This can be modelled in an analogous way to RSDs in galaxy clustering. Although not included in our fiducial model, we can assess the impact in the NLA case using the recipe set out in Singh & Mandelbaum (2016), section 2.3. Fortunately we see the impact is largely contained at separations above $r_p > 70\ h^{-1}\,\mathrm{Mpc}$, which are already removed by the upper r_p cut. Within the range of scales used for our fits, the impact is comfortably smaller than our error bars.

5 PIPELINE TESTING

In this section, we describe the various tests of the analysis pipeline, and the measurements themselves. These include tests of the theory predictions by comparing different code implementations. We seek to validate the pipeline by reanalysing an existing data set and comparing with published results. Finally, we discuss null tests on the data, designed to be sensitive to residual systematics.

5.1 Reanalysing LOWZ

For the purposes of validating our measurements and demonstrating comparability with previous results, we partially reanalyse the BOSS LOWZ catalogues of Singh et al. (2015) (see also Section 2). LOWZ makes a good test data set for several reasons – not least that it has documented, relatively high-signal-to-noise ratio w_{g+} measurements in the literature, and the redshift catalogue is publicly available. We repeat all measurement steps downstream from shape catalogues using our pipeline, and then fit the resulting correlation functions with our modelling set-up. At the level of data vectors, our pipeline can reproduce the galaxy clustering and galaxy–shape correlations, w_{gg} and w_{g+}, of Singh et al. (2015) to sub- per cent precision on scales $[2, 70]\ h^{-1}\,\mathrm{Mpc}$.

We also analyse the LOWZ data on large scales, and compare our results to those of Singh et al. (2015); when matching their analysis choices exactly (NLA model, linear bias, $r_p > 6\ h^{-1}\,\mathrm{Mpc}$), we recover
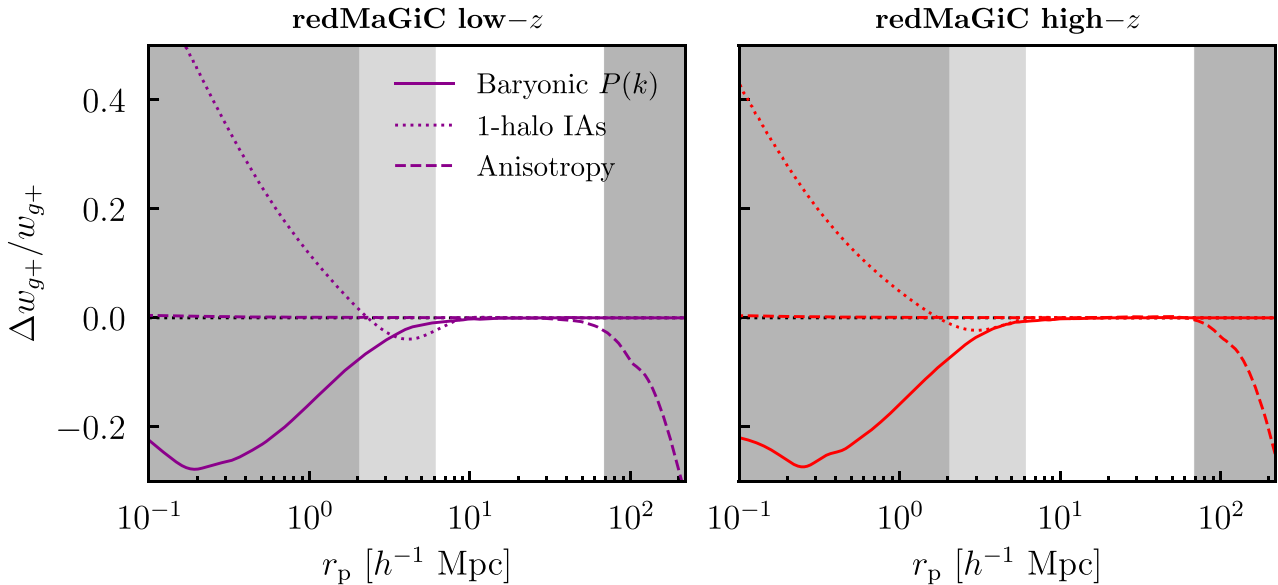
Figure 8. The impact of higher order effects on galaxy–shape correlations. We include both redMaGiC high-z and redMaGiC low-z to show that our conclusions hold across the redshift range of our samples. As before, the shaded bands indicate scales removed in our fiducial NLA (light grey) and TATT (dark grey) analyses. We include here the impact of baryonic feedback, as represented by OWLS-AGN (solid line); the impact of one-halo alignment physics, as represented by the model of Schneider & Bridle (2010) (dotted), and signal boosting due to anisotropy in the 3D correlation (Singh & Mandelbaum 2016; dashed). In all cases, the unmodelled effects appear at the level of a few per cent or less on the scales we use, which is below the level of our statistical precision ($\sigma_{w_{g+}}/w_{g+} \sim 10$–$15$ per cent on scales $r_{\rm p} > 2\ h^{-1}$ Mpc).

their reported best fit in the $A_1 - b_1$ plane to $\ll 1\sigma$. Our fiducial analysis configuration differs from the published LOWZ paper in a number of ways. Most significantly, these include:

(i) Our assumed cosmology is that set out in Section 4, instead of WMAP9 (Hinshaw et al. 2013). This results in a slight increase in the amplitude of the matter power spectrum, which in turn results in a slightly lower alignment amplitude. Note that our fiducial cosmology includes massive neutrinos, which modify the non-linear $P(k)$ slightly. The difference in h also alters the measurement of the two-point functions earlier in the pipeline (via the redshift to distance conversion), although this difference is minimal.

(ii) Our fiducial data vector includes galaxy–galaxy, galaxy–shape, and shape–shape correlations, whereas Singh et al. (2015) include only the former two.

(iii) We use the Takahashi et al. (2012) version of halofit to compute the non-linear matter power spectrum, whereas Singh et al. (2015) use a slightly older release (Smith et al. 2003).

(iv) We use an analytic calculation to estimate the data vector co-variance matrix, instead of jackknife. While the two agree relatively well, slight differences in the relative weighting of different scales in both w_{gg} and w_{++} are apparent.

(v) We include contributions from lensing and magnification in our model. Although this has little impact on a low-redshift spectroscopic data set such as LOWZ, it has a larger bearing on our eBOSS and redMaGiC samples.

We show a more detailed comparison at the parameter level in Appendix C. In short, when matching the analysis choices of Singh et al. (2015), we can reproduce their published IA results almost exactly. Switching to our fiducial NLA set-up produces a very similar result, with a small reduction in the size of the error bars.

5.2 Null tests

A number of systematics (e.g. PSF modelling errors) can lead to a non-zero mean shear. Unlike multiplicative biases, we can look for such effects directly using the data. We find no evidence of such a signal in any of the samples considered here, with $|\langle e_i \rangle| \sim 10^{-4}$ in all cases. A number of other tests for systematic induced signals are presented in Gatti et al. (2021); they find no evidence for correlations between the response-corrected shear and PSF shape and size, or for a statistically significant B-mode signal.

We also measure one additional null signal. Constructing $w_{g \times}$ involves the same basic quantities as w_{g+}, but measuring the shape component that is rotated 45 deg with respect to the radial/tangential direction. Like lensing, astrophysical processes such as IAs, to first order, should induce only tangential/radial correlations.[10] Non-zero detection of a cross signal, then, is a red flag for residual measurement systematics. For all samples considered, we find these additional measurements to be consistent with zero within the scales $r_{\rm p} < 70\ h^{-1}$ Mpc. Details of the measurements can be found in Appendix B.

6 RESULTS

This section presents the results of our analyses on the various samples. Although we will focus on IAs, it is worth bearing in mind that each analysis also includes two free galaxy bias parameters. The

---

[10] Although some IA models predict a non-zero B-mode contribution (see e.g. Catelan et al. 2001; Hirata & Seljak 2004; Blazek et al. 2019), which translates into correlations in the cross ellipticity component, such effects appear only in the II alignment spectra. Given that our constraints are dominated by g + correlations, these terms are thought to be easily subdominant to noise in current surveys.

Figure 10. 68 per cent and 95 per cent confidence intervals for the TATT model constraints from our two DES Y3 redMaGiC samples, on scales $r_p > 2\ h^{-1}$ Mpc. As before, we show redMaGiC high-z in red and redMaGiC low-z in purple. For comparison, the marginalized NLA constraints on A_1 are shown as dashed curves in the upper left panel. Note that the scales here differ from the fiducial NLA analyses, and so the best fits are different from those shown in Table 4. The dotted black lines mark the zero-points for the non-NLA parameters.

constraints on the bias parameters are strongly dominated by w_{gg}, and so they contribute relatively little to the marginal uncertainties on IA parameters. The bias does, however, also enter w_{g+}, and so it is important to model it accurately. In every case, the linear bias falls within the bounds of expectation from previous studies ($b_1 \sim$ 1.5–2.0, depending on the sample), and b_2 is small ($b_2 \sim 0$–0.3). We note that all samples appear to be fit reasonably well by our model (as quantified by the best χ^2 obtained from fits to the joint $w_{gg} + w_{g+} + w_{++}$ data vector). For more detail on the bias constraints, see Appendix D. It is also worth bearing in mind that all parameters (bias and IA) are constrained within the prior bounds. As we note below, although some samples provide only weak constraints, the priors in Table 2 are sufficiently wide to allow the contours to close in all cases. In Section 6.1, we discuss our results on ELGs, which amount to a null detection. We then move on to the various red samples in Section 6.2, presenting first large-scale results using NLA in Section 6.2.1, and then extending to slightly smaller scales with TATT in Section 6.2.2. Section 6.3 then considers more carefully the level of contribution from lensing and magnification.

6.1 Emission-line galaxies

Our first, and perhaps easiest to interpret, results are based on eBOSS ELGs. The data vector is shown in blue in Fig. 5. We fit the NLA model on large scales (the unshaded region in Fig. 5), and obtain a null detection,

$$A_1^{\mathrm{ELG}} = -0.42^{+0.50}_{-0.50} \qquad (r_p > 6h^{-1}\mathrm{Mpc}) \tag{32}$$

with $\chi^2/\mathrm{dof} = 1.17$ (with a corresponding p-value $p = 0.32$). This is expected, given the sample: a non-zero IA signal has never been

detected in ELGs (or in any colour-selected sample of blue galaxies more generally; Mandelbaum et al. 2011; Johnston et al. 2019; Samuroff et al. 2019). The additional (non δ_gI) terms are also seen to be small, for a number of reasons: first, the magnification coefficient is small $(\alpha - 1) \sim 0.1$, for ELGs, which scales down the μI and μG contributions. Second, the limits of the line-of-sight integral tend to suppress the lensing contributions to the signal; integrated out to $\Pi_{\max} = 1000\ h^{-1}$ Mpc, μG tends to dominate on larger scales. In practice, however, with integral limits at $\pm 100\ h^{-1}$ Mpc, the largest term by some way is δ_gI, with $w_{\delta_g I}/(w_{\delta_g G} + w_{\mu G} + w_{\mu I}) \sim 14$ at 6 h^{-1} Mpc (as evaluated at the best-fitting parameters). Similarly, for the shape–shape correlation, the ratio of II to other terms is ~ 18. The end result is a combined best-fitting theory prediction that is below the level of shape noise.

Since the signal-to-noise is relatively low, and there is no visible structure in w_{g+}, we also repeat our NLA fits with slightly less stringent cuts, $r_p > 2\ h^{-1}$ Mpc. This tightens the bounds on the alignment amplitude to

$$A_1^{\mathrm{ELG}} = 0.07^{+0.32}_{-0.42} \qquad (r_p > 2h^{-1}\mathrm{Mpc}) \tag{33}$$

Indeed, even considering scales down to 0.1 h^{-1} Mpc in Fig. 5, we still see no hints of signal in w_{g+} or w_{++}. Computing the null χ^2 on all scales $r_p < 70\ h^{-1}$ Mpc, we find $\chi^2 = 17.5$ for 16 data points ($p = 0.42$). This is interesting, since it suggests that there is not a strongly scale-dependent one halo (1h) signal of the sort seen in the redMaGiC and CMASS samples (or at least, not one that is detectable above the level of shape noise).

In terms of sample, the closest results in the literature are those of Mandelbaum et al. (2011) and Tonegawa et al. (2018). These both use blue ELGs, from WiggleZ and Subaru, respectively, and also make null detections of $A_1^{\mathrm{M11}} = 0.15^{+1.03}_{-1.07}$ and $A_1^{\mathrm{T18}} = 0.49^{+3.56}_{-3.56}$, respectively. Our results tighten the null constraint, imposing an upper limit of $|A_1| < 0.78$ at 95 per cent CL. In terms of redshift, our eBOSS ELG measurements sits between the earlier two ($z \sim 0.8$, compared with $z \sim 0.5$ for WiggleZ and $z \sim 1.4$ for Subaru). It is worth exercising some caution here, however, since in both cases it is not clear that the sample matches ours closely. In particular, the Mandelbaum et al. (2011) sample is a relatively bright selection of starburst galaxies with specific colour cuts (see their section 3.1). That said, the best-fitting bias values are relatively similar to our own ($b_1 \sim 1.4$ for eBOSS ELGs, $b_1 \sim 1.5$ for WiggleZ). The Tonegawa et al. (2018) sample on the other hand has both a complicated spectroscopic selection function, and additional shape catalogue cuts that remove ~ 50 per cent of objects. Although all three results (including our own) present IA results consistent with zero in blue-ish samples across a range of redshifts, it is not clear the results are directly equivalent.

It is also worth stressing here that although very different from (and much bluer than) an LRG or redMaGiC type sample, eBOSS ELGs are not exactly representative of a typical weak lensing catalogue either. Indeed, the eBOSS ELG selection is designed to facilitate a high S/N BAO measurement, and is based on the presence of particular emission lines, as determined via a complex set of magnitude and colour cuts (Raichoor et al. 2017). The completeness of the sample given the cut is difficult to quantify (Guo et al. 2019). In contrast, lensing samples tend to have much simpler (if any) colour selection, and cuts designed to minimize lensing measurement biases and optimize a weak lensing measurement. The two are designed for different scientific purposes, and so we should not expect them to match. For this reason, caution is required when trying to extrapolate these results.
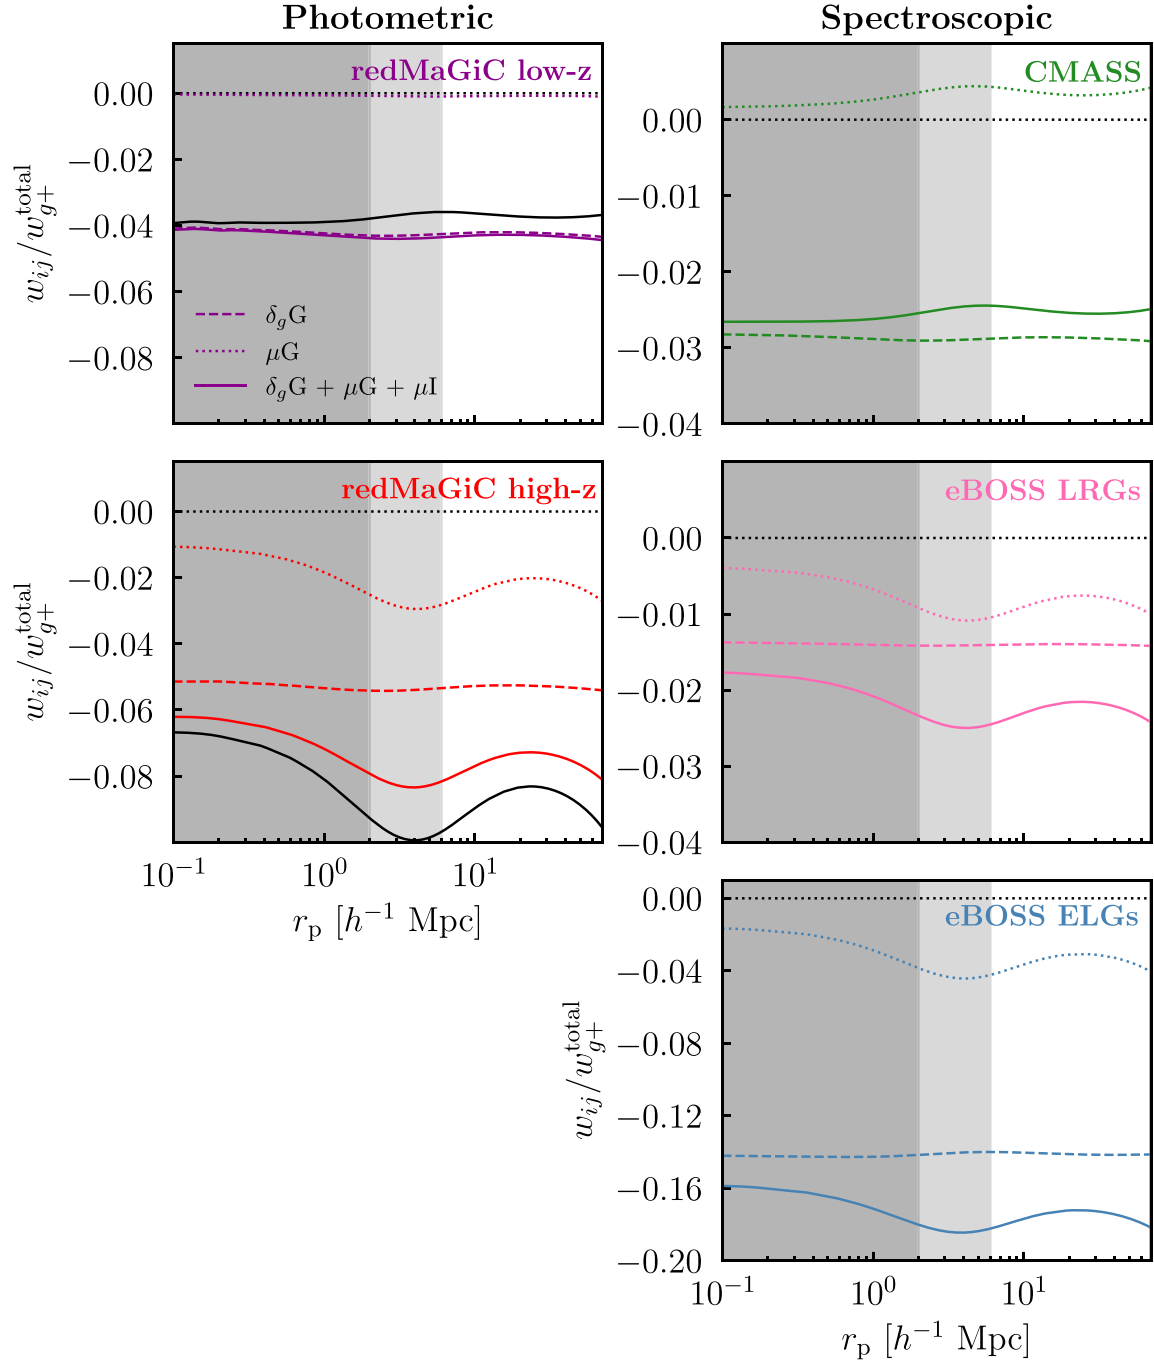
Figure 11. The impact of lensing and magnification on our best-fitting theory data vectors. Each contribution to w_{g+} is shown as a fraction of the total signal ($\delta_g I + \delta_g G + \mu G + \mu I$). Clockwise from top left, we have redMaGiC low-z (purple), CMASS (green), eBOSS LRGs (pink), eBOSS ELGs (blue), and redMaGiC high-z (red). In the left-hand column we show our two photometric samples, which are more strongly affected by lensing and magnification. On the right are our three spectroscopic samples. Within each column, the samples are arranged vertically by redshift, starting with low-redshift samples at the top. As in previous figures, the shaded bands represent scales excluded from fits to these data vectors. Different line styles indicate different signal components, with the sum of the non $\delta_g I$ terms shown as a solid line. Note that the μI term is negligible on all scales and for all samples, and so we do not show it separately. In the two redMaGiC cases, we include a second solid line (black). This demonstrates the impact of using an alternative estimate for the magnification coefficients α, obtained using BALROG source injection. Unlike in the other samples, the CMASS μG term is positive. This is because the magnification coefficient for this particular sample is small ($\alpha < 1$, see Section 3.3); in physical terms this means the geometric effect from magnification, which expands the apparent area of a patch of sky, outweighs the boost in number density due to the brightening of the apparent galaxy fluxes. Note that the scale on the vertical axis differs between panels.

Table 4. NLA model constraints from the various red galaxy samples discussed in this work, as a function of r-band luminosity. The upper four rows refer to the full samples, while the lower 12 refer to subsets in luminosity bins, as defined in Section 3.5 and Fig. 6. The first three columns show the mean luminosity, redshift, and rest-frame colour of each sample. The final two are the posterior mean IA amplitude, and the corresponding p-value. In all cases the model is seen to provide a reasonable fit to the data. These fits were performed using large scales only ($r_p > 6 h^{-1}$ Mpc) in w_{g+} and w_{++}.

Sample	$\langle L_r \rangle / L_0$	$\langle z \rangle$	$\langle M_r - M_z \rangle$	A_1	$p(> \chi^2)$
RMH (all L)	0.68	0.78	0.43	$3.54^{+0.18}_{-0.18}$	0.20
RML (all L)	0.40	0.46	0.43	$3.34^{+0.13}_{-0.13}$	0.12
CMASS (all L)	0.84	0.52	0.36	$2.72^{+0.47}_{-0.47}$	0.10
LRGs (all L)	1.20	0.75	0.39	$5.78^{+1.19}_{-1.19}$	0.89
RMH $L0$	0.49	0.77	0.44	$3.47^{+0.34}_{-0.34}$	0.11
RMH $L1$	0.63	0.77	0.42	$3.01^{+0.31}_{-0.30}$	0.52
RMH $L2$	0.83	0.78	0.44	$4.13^{+0.33}_{-0.34}$	0.15
RMH $L3$	1.16	0.78	0.45	$5.22^{+1.11}_{-1.11}$	0.40
RML $L0$	0.19	0.32	0.50	$1.95^{+0.23}_{-0.23}$	0.28
RML $L1$	0.26	0.43	0.42	$2.86^{+0.29}_{-0.29}$	0.22
RML $L2$	0.35	0.47	0.40	$3.01^{+0.52}_{-0.51}$	0.62
RML $L3$	0.49	0.50	0.40	$4.39^{+0.20}_{-0.20}$	0.79
RML $L4$	0.73	0.53	0.40	$5.00^{+0.20}_{-0.21}$	0.14
CMASS $L0$	0.52	0.48	0.42	$2.23^{+0.92}_{-0.92}$	0.47
CMASS $L1$	0.79	0.52	0.36	$3.00^{+0.62}_{-0.61}$	0.47
CMASS $L2$	1.24	0.55	0.29	$3.78^{+0.49}_{-0.49}$	0.70

6.2 Red Galaxies

6.2.1 Constraints on large-scale IAs

We next consider our other galaxy samples, redMaGiC, eBOSS LRGs, and CMASS, which we fit on large scales (again, $>6 h^{-1}$ Mpc) using NLA. In each case, we find a clear detection, with our three-parameter model of the joint data vector (A_1, b_1, b_2) providing a good χ^2/dof. The constraints and the goodness-of-fit statistics can be found in the upper four rows of Table 4. Defining the signal-to-noise according to equation (2) of Becker et al. (2016),[11] we find $S/N = 22$ in redMaGiC low-z and $S/N = 18$ in redMaGiC high-z. Given the smaller area, the detections in our SDSS samples are slightly weaker, at $S/N = 6$ for CMASS and $S/N = 5$ for eBOSS LRGs. The best-fitting model predictions can be seen in Fig. 5. As we saw with the ELGs in Section 6.1, CMASS and eBOSS LRGs are dominated by the primary IA signal (δ_gI for w_{g+}, II for w_{++}). For redMaGiC, however, the picture is slightly different; photo-z scatter tends to increase the maximum distance galaxies can be physically separated (by shifting well-separated objects below Π_{max}), and so boosts the lensing and magnification terms. With redMaGiC low-z this is partly cancelled out by the fact that α is very close to 1, and that the mean redshift is relatively low. These things are less true for redMaGiC high-z, however, and so we see a stronger δ_gG contribution. The δ_gI signal is also slightly stronger, however, and

---

[11]The expression is $S/N = (\hat{\mathbf{w}}\mathbf{C}^{-1}\mathbf{w}^{\mathrm{model}})/(\mathbf{w}^{\mathrm{model}}\mathbf{C}^{-1}\mathbf{w}^{\mathrm{model}})^{\frac{1}{2}}$, where $\hat{\mathbf{w}}$ is the observed (noisy) data vector, $\mathbf{w}^{\mathrm{model}}$ is the best-fitting theory prediction, and $\mathbf{C}^{-1}$ is the inverted covariance matrix. This is slightly different from the common definition using $\hat{\mathbf{w}}$ only, which is known to be biased high if noise is present.

Figure 9. IA strength as a function of $k + e$-corrected r-band luminosity in red galaxies. By convention, the luminosities are defined relative to a pivot L_0, which corresponds to an absolute magnitude $M_r = -22$. Open points show previous results from the literature (as indicated in the legend). For illustrative purposes we also show two power-law fits from the literature, one fit to GAMA + SDSS at low-mid L (dotted blue), and the other to MegaZ at high L (solid purple). The filled points show our measurements on redMaGiC low-z (stars; five points), redMaGiC high-z (red dots, four points), eBOSS LRGs (pink triangle; one point), and CMASS (green upside down triangles; three points). Note that all results shown here assume the one-parameter NLA model.

the additional terms still account for only order of a few per cent of the total signal.

Another useful exercise is to divide the samples into luminosity bins, and map out the dependence of the alignment signal. For each of the bins shown in Figs 3 and 6, we fit an NLA amplitude (we also fit for galaxy bias, but since we only split the shape sample by luminosity, that does not change significantly between L bins). The results are shown in Fig. 9, with numerical parameter constraints in Table 4. Note that we also include a selection of previous measurements from the literature, denoted by open points. There are a number of trends worth considering here. First, taken naively, our results are consistent with the qualitative picture of a broken power-law dependence on luminosity: the trend in A_1 below $\log L/L_0 \sim -0.2$ is much shallower than above it. The $(L/L_0)^{\beta}$ power-law parametrization was first introduced as an empirical scaling by Joachimi et al. (2011). Although there is little physical motivation, it has been adopted relatively widely both in direct IA measurements (Singh et al. 2015; Johnston et al. 2019; Fortuna et al. 2021b) and in forecasts (Krause, Eifler & Blazek 2016; Fortuna et al. 2021a), as it was simple to implement and appeared to fit the available data relatively well. In recent years, the picture has become more complicated, as evidence has begun to emerge of a weaker relationship at low L (see e.g. Johnston et al. 2019). Again, however, this is empirical, and there is no first-principles reason to expect a double power law in particular (or any other form). Our results appear to reinforce that evidence. Our redMaGiC low-z sample in particular provides a significant improvement in the constraints on the

fainter end of the A_1–L relation (by a factor of 3 or more in the error bars). At the brighter end, our CMASS, LRG, and redMaGiC high-z samples also appear qualitatively consistent with previous results, following a considerably steeper slope. Taken at face value, given the error bars, we could interpret this as ruling out a single power law with relatively high significance. We stress, however, that it is worth being cautious here. Despite all being 'red', there are differences between the composition of the samples, as we will come to below – it is possible these population differences may be partly responsible for the apparent trends in L–A_1 space. For this reason, we do not present the best-fitting constraints on β, but rather a more qualitative discussion of how to interpret our results.

It is noticeable that CMASS (and to a lesser extent redMaGiC high-z) tends to lie roughly $\sim$1–3σ below the best-fitting single power law from the literature (the purple line in Fig. 9). Since the error bars here are mostly shape noise dominated and the luminosity bins are disjoint, the points should also be uncorrelated to good approximation, meaning it is unlikely this is random scatter. We can perhaps understand these trends in terms of colour differences. In the upper panel of Fig. 4 we can see that CMASS is considerably bluer than LOWZ. Indeed, while most extreme for CMASS, all of our samples tend to peak lower than LOWZ in $M_r - M_z$ space. Not only this, there is also some difference in colour between luminosity bins for a given sample. For example, the galaxies in the upper CMASS bin ($L2$), on average, have slightly bluer rest-frame magnitudes than the lower two bins. Although for the sake of convenience, it has been useful to split galaxies into binary 'red' and 'blue' categories, our results suggest that this may be an oversimplification for modelling purposes, given the precision of current data sets. They suggest that a more sophisticated modelling may be needed, which accounts for colour and luminosity (and potentially other properties such as satellite fraction) simultaneously. This will be the focus of future work.

We can compare these results with those of Singh et al. (2015), who consider colour bins within the LOWZ LRG sample. Although that work reported no clear trend across five bins in $g - i$ rest-frame colour, it should be noted that LOWZ covers a fairly narrow range in colour space (see the black contour in Fig. 4). Even the bluest bin in that paper still represents a relatively bright red sample compared with the galaxies considered here. It seems plausible that our wider coverage allows us to see a trend that is not detectable in a relatively homogeneous sample like LOWZ. We thus consider the two results qualitatively consistent.

One other feature worth mentioning, although we do not seek to quantify it, is the behaviour of w_{g+} on very small scales. Both red-MaGiC samples appear to exhibit a strongly luminosity-dependent 1h contribution to w_{g+} (Fig. 5, purple and red). Noticeably, the RML L3 and RMH L0 bins, though having very similar mean luminosities, have qualitatively different one halo signals. In the case of eBOSS LRGs and CMASS we see no such trends, but this is quite possibly simply the result of low S/N, even on smaller scales. In all of these cases, it is worth bearing in mind that the density tracer samples differ. Although the differences between the small scale behaviour of w_{g+} in the various bins/samples *could* be a result of the 1h IA signals, they could also be partly down to differences in the small-scale galaxy–halo connection. A full halo model that includes IA could potentially distinguish these effects – we leave this topic for future investigation.

We also briefly test for redshift evolution in our red galaxy samples. Since we are interested in isolating inherent evolution in the IA signal (as opposed to changes in sample composition), we compare samples at roughly the same luminosity. Specifically, we define two narrow bins in Fig. 9 (one in the low-luminosity

regime, at $\log L/L_0 \sim -0.3$, and the other at higher L, $\log L/L_0 \sim -0.05$). Plotting out $A_1(z)$ in these two slices, we find no evidence for redshift evolution over the range $z = [0.25, 0.8]$. Although we see the same trend with CMASS being slightly lower than other samples at the same L, there is no evidence that this is the result of an underlying redshift trend. Since there is no statistically significant correlation, we do not include the figure in the main body of the paper; for completeness, however, it is shown along with redshift power-law constraints in Appendix E, Fig. E1.

6.2.2 Model comparison: NLA and TATT

In Section 6.2.1, we explored the behaviour of IAs on very large scales, where the NLA model is thought to be sufficient. We next turn to a slightly different question: on what scales precisely does the simple model break down? Based on theory, there is thought to be an intermediate regime, outside the one halo regime, but where higher order correlations (such as those included in the TATT model) become significant. It is still, however, and open question as to *how* significant and on exactly *which* scales.

To explore this, we repeat our analysis with the minimum scale reduced slightly to $r_p > 2\ h^{-1}$ Mpc in w_{g+} and w_{++}. Note that the cuts on w_{gg} are fixed, and the modelling there does not change (i.e. there are always two free galaxy bias parameters). For each sample, we fit both NLA (one free IA parameter) and TATT (three IA parameters). Our results are summarized in Fig. 10.

Unsurprisingly given Fig. 5, CMASS, eBOSS ELGs, and eBOSS LRGs provide very broad constraints on the TATT parameters, and so are not shown. The signal-to-noise in these samples is still relatively low. Although one can fit a single amplitude relatively well, there is little constraining power left for the shape of the correlation function. The picture is slightly different, however, in our redMaGiC samples. The cosmological volume here is significantly larger, and the S/N higher. Starting with the slightly larger sample, redMaGiC low-z, we find

$$A_2^{\mathrm{RML}} = 1.24^{+0.74}_{-0.80} \qquad b_{\mathrm{TA}}^{\mathrm{RML}} = -0.43^{+0.18}_{-0.13}. \tag{34}$$

In words, our measurements favour (albeit relatively weakly) a combination $A_2 > 0$ and $b_{\mathrm{TA}} < 0$. Comparing with the NLA fits on the same scales, we find $p(> \Delta\chi^2) = 0.01$ ($\chi^2_{\mathrm{NLA}} = 31.2$ and $\chi^2_{\mathrm{TATT}} = 21.9$).

In the case of redMaGiC high-z, the data appear to prefer a non-zero b_{TA} at the level of roughly 2.5σ, with

$$A_2^{\mathrm{RMH}} = -0.36^{+1.63}_{-1.64} \qquad b_{\mathrm{TA}}^{\mathrm{RMH}} = 0.92^{+0.41}_{-0.36}. \tag{35}$$

Again, comparing the TATT and NLA fits, we obtain $p(> \Delta\chi^2) = 0.00004$, suggesting a statistically significant preference for TATT on these scales; the respective goodness-of-fit statistics are $\chi^2_{\mathrm{NLA}} = 33.9$ and $\chi^2_{\mathrm{TATT}} = 13.6$, and $\Delta\mathrm{dof} = 2$.

Interestingly, the deviations from NLA manifest in quite different ways at the data vector level (see Fig. 5). In the case of redMaGiC high-z, the positive b_{TA} increases the power on intermediate scales ($2\ h^{-1}$ Mpc $< r_p < 10\ h^{-1}$ Mpc), resulting in a significantly flatter w_{g+}. For redMaGiC low-z, on the other hand, we see the opposite effect; a reduction in the amplitude of w_{g+} in the ($2\ h^{-1}$ Mpc $< r_p < 6\ h^{-1}$ Mpc) range is accompanied by a slight increase around 20–30 h^{-1} Mpc. This gives a slightly steeper theory prediction, which matches the shape of the measured correlation relatively well. Taken together, this amounts to a $\sim$3.8σ difference in b_{TA} between the two redMaGiC samples.

Although interesting, it is worth being cautious here. Since TATT is a relatively flexible model (albeit a physically motivated one), it

is possible that non-zero A_2 and $b_{\rm TA}$ values could arise due to other untreated systematics. Baryonic physics, for example, tends to appear on small-intermediate scales, and modulates the power spectrum in a scale-dependent way. It is also true, however, that baryons tend to *suppress* power on smaller scales; this is the case in (almost) all hydrodynamic simulations at all redshifts. Since redMaGiC high-z prefers a TATT model that does the opposite (relative to NLA), it seems unlikely that baryons are driving the non-zero $b_{\rm TA}$ value here. Even considering the redMaGiC low-z case, it seems very unlikely that we are simply seeing residual baryonic feedback. In Section 4.6, we saw that the OWLS-AGN scenario (itself an extreme case) produced at most $\Delta\chi^2 = 0.45$ at the TATT best fits quoted above. The difference between the NLA and TATT goodness of fit are more than an order of magnitude larger than this; for baryonic feedback alone to explain the non-zero TATT parameters would require a significantly more extreme scenario than OWLS-AGN.

Another possible effect here is non-linear galaxy bias. Our model for it is incomplete in the sense that while we include non-linear bias in our w_{gg} model, only b_1 enters the w_{g+} prediction. Incorporating non-linear bias, and all the TATT-bias cross terms, into the w_{g+} model is the focus of ongoing work. We can, however, make a rough estimate for the impact based on our w_{gg} fits. For redMaGiC low-z, we find the data consistent with linear bias ($b_2 = -0.09 \pm 0.07$). The equivalent value for redMaGiC high-z is slightly larger, but still small $b_2 = 0.39 \pm 0.08$. Since the additional terms contributing to w_{g+} will be proportional to b_2 multiplied by the various IA coefficients, it seems likely that they should be relatively small compared with the IA-only contributions.

Finally, we also consider the possibility that our results here could be the result of a non-local lensing contribution from small scales. Such contributions add to the galaxy–galaxy lensing (gG) term, and tend to boost its power on small to intermediate scales (see e.g. Baldauf et al. 2010; MacCrann et al. 2020). Fortunately, even considerably different halo mass profiles produce approximately the same contribution on scales well outside the virial radius, behaving effectively as an enclosed point mass and scaling as $1/r_{\rm p}^2$. To test this, we generate NLA-only theory data vectors from the NLA fits on large scales; we add a point mass term (equations 7 and 8 of MacCrann et al. 2020), and adjust δM until the NLA + PM theory prediction for w_{g+} matches the data on scales $2\ h^{-1}\,{\rm Mpc} > r_{\rm p} > 6\ h^{-1}\,{\rm Mpc}$. Although a $1/r_{\rm p}^2$ scaling *can* match the data in the redMaGiC high-z case, we find the mass required to do this is $\sim 3 \times 10^{15}\ {\rm M}_\odot\,h^{-1}$, which is much larger than the typical halo mass expected for DES redMaGiC (see, for example, Pandey et al. 2022; Zacharegkas et al. 2022). This would also require the point mass contribution to dominate gG out to scales of $\sim 20\ h^{-1}$ Mpc, which again is not thought to be realistic. Moreover, the point mass explanation should lead to an excess w_{g+} on small scales for both redMaGiC samples, contrary to the observed behaviour. For all of these reasons, we conclude that a point mass term cannot explain the deviations from NLA on small scales for redMaGiC high-z.

Overall, these tests seem to suggest a real IA signal (or, at least, a significant systematic that we have not considered). This is interesting from a modelling perspective. It implies some dependence in the TATT parameters with galaxy properties ($b_{\rm TA}$ most obviously, going from negative in redMaGiC low-z to positive in redMaGiC high-z, but also potentially A_2). As we discussed in the previous section, the redMaGiC high- and low-z samples differ in redshift, but also in galaxy properties like colour and luminosity. Although, given this, differences somewhat expected, this is not something on which there are previous results to guide us. Disentangling what

exactly is driving the differences is an interesting question, but a potentially difficult one to answer. We leave this for future work.

6.2.3 Robustness to cosmology and $X_{\rm lens}$

In this section, we seek to test the robustness of our analysis to various sources of systematic error. One such potential contaminant is the effect known as $X_{\rm lens}$. One can find extensive early discussion in the DES Y3 results papers (Dark Energy Survey Collaboration 2022; Pandey et al. 2022), but essentially $X_{\rm lens}$ is a multiplicative factor of unknown origin between the amplitudes of the galaxy–galaxy lensing and galaxy clustering measurements. This offset was seen to be scale- and redshift-independent, and to impact only Y3 redMaGiC, and not the fiducial magnitude-limited lens sample, MagLim. Subsequent tests have pointed towards a systematic in the photometry, which affects the redMaGiC selection (see Pandey et al. 2022 Sec. VG). The magnitude of $X_{\rm lens}$ is constrained relatively well by the 3 × 2pt data in Dark Energy Survey Collaboration 2022, to $X_{\rm lens} = 0.877^{+0.026}_{-0.019}$, which is roughly the size of the fractional error bar, $(w_{g+} - \sigma_{w_{g+}})/w_{g+}$, for our redMaGiC samples in the range $2\ h^{-1}\,{\rm Mpc} < r_{\rm p} < 70\ h^{-1}\,{\rm Mpc}$. Since it is scale-independent, we expect the impact to be completely degenerate with A_1. Given these things, we do not expect $X_{\rm lens}$ to have a qualitative impact on our results. Although it may modulate the best fit A_1 in our redMaGiC (not CMASS or eBOSS) at the level of ~ 10 per cent, comparison between samples is already uncertain to at least this level due to differences in colour space. Given that the TATT parameters primarily alter the shape of w_{g+}, we do not expect $X_{\rm lens}$ to alter the findings of Section 6.2.2.

We also briefly consider the impact of our choice of cosmology; to fit for IAs, we need to assume a particular set of cosmological parameters (e.g. for calculating the matter power spectrum). As discussed in Section 4, we assume a flat ΛCDM universe with massive neutrinos and a clustering amplitude σ_8 similar to that reported by Planck. For each of our samples, in addition to the best-fitting NLA data vector, we generate a second with a perturbed cosmology; for this we choose the DES Y3 1 × 2pt best fit. By comparing w_{gg} at the two cosmologies, we can compute an effective shift in large-scale bias Δb_1. This in hand, plus the observed impact on w_{g+}, we can estimate the shift in the best fit A_1. The end result is a change of at most one or two per cent. That is, the difference between plausible cosmologies is not sufficient to significantly affect our results.

6.3 Assessing the contribution of magnification and lensing

In addition to the main IA signal, our measurements have contributions from lensing and magnification (see the discussion in Section 4). These are always included in our modelling, but it is interesting to briefly discuss their effect. The fractional impact of the various terms is shown in Fig. 11. Note that the data vectors here are evaluated at the NLA best fit for each sample, and so b_1 and A_1 differ somewhat between the panels.

Consider first the two photometric samples, redMaGiC low-z and redMaGiC high-z (upper two panels, purple and red in Fig. 11). Here, we see a total non-δ_gI contribution of roughly 8 per cent and 4 per cent in the high- and low-z samples, respectively. For context, the 1σ uncertainty on A_1 in these samples is ~ 4–5 per cent (see Table 4); including these effects in the modelling of w_{g+} is clearly necessary to avoid bias at the current precision. The closest comparison in the literature is fig. 5 in Joachimi et al. (2011); again, we confirm that

our pipeline can reproduce those results. We see a similar ordering of the terms to their figure, with δ_gG dominant, followed by μG, and with μI the smallest, at sub- per cent level (although note there are some important differences between the comparison in Fig. 11 and that of Joachimi et al. 2011, and so one should not expect the details to match perfectly).

In the two left-hand panels, we also show the results of the same exercise, but using alternative estimates for the magnification coefficients from BALROG (black lines; see Section 3.3 for the actual values). Note that α is the only quantity that changes here; we do not refit the data vectors, and so the dominant δ_gI term in the denominator is fixed. The new α values are thought to include additional effects omitted by the flux-only method described in Section 3.3. Unfortunately, however, we do not have BALROG injections covering the whole DES footprint, and so the resulting estimates are noisy. We also cannot easily produce BALROG mocks closely matching our CMASS and eBOSS samples, which somewhat limits its use for our purposes. The overall impact, however, is seen to be relatively small. That is, while ignoring magnification altogether could have a significant impact on our results, the choice of one α estimate over another is unlikely to.

Although it is common to assume spectroscopic IA measurements are immune to lensing effects, we see in Fig. 11 that this is not entirely the case. The two main terms are the same as with redMaGiC: μG and δ_gG. It is interesting that the latter (dashed) still dominates; even in the case of very narrow per-galaxy redshift distributions, where $p(z) \rightarrow \delta_z$, the ensemble redshift distribution is sufficiently broad to allow a non-negligible galaxy–galaxy lensing contribution. The μG term (dotted), we also note, does not depend on the quality of the redshifts. Galaxies along the same line of sight are affected by magnification and lensing due to the same foreground structure, which modulates w_{g+} at the level of a few per cent at $z \sim 0.8$.

It is worth briefly considering what these results mean for future IA measurements. The Stage IV spectroscopic survey DESI is expected to obtain spectra for a sample of LRGs over a comparable redshift range to ours, but for a considerably wider area and greater number density ($n_c \sim 6 \times 10^{-4} \, h^3 \, \mathrm{Mpc}^{-3}$ over 14 000 square deg; Zhou et al. 2020). Similarly, one can expect at least an order of magnitude increase in the number of ELGs available for the type of measurement we perform here (Raichoor et al. 2020). Euclid and the Roman Space Telescope will also have spectroscopic instruments, which will further add to the pool of data available for IA measurements. Likewise, Stage IV photometric surveys such as Rubin, Euclid, and Roman will probe a similar selection of galaxies to DES, but over a much wider area and to a greater depth (see, for example, Euclid Collaboration 2021). These will allow measurements using a redMaGiC-like sample similar to ours, but with significantly improved S/N and finer binning in colour/redshift/luminosity. Used together – with photometry providing shape inference, and accurate redshift information from either spectroscopic data or high-quality photometric measurements – the next generation of surveys will provide a powerful tool for studying IAs. Given what we see here, it seems very likely that direct IA measurements using these upcoming data will need to account for lensing and magnification.

It is finally worth noting that any data-based estimates of magnification will have some level of noise; this is most obvious for our BALROG estimates (due to the limited area), but it is also true at some level for the flux-only method. As future data sets become more constraining it will likely be necessary to properly characterize this uncertainty, and marginalize over α either directly or analytically.

7 CONCLUSIONS

This paper presents direct constraints on IAs from the Dark Energy Survey Year 3 shape catalogues. The Y3 METACALIBRATION catalogue is used to provide shape estimates for 2.4M redMaGiC galaxies from across the DES footprint, as well as ~50 000 CMASS galaxies, ~22 000 eBOSS LRGs, and ~100 000 ELGs. We make a high significance detection of IAs in all of these samples, with the exception of eBOSS ELGs, where we place upper bounds on the magnitude of the possible alignment amplitude.

The key conclusions of this paper are:

(i) Fitting for A_1 in red galaxies, our data support the qualitative picture of a broken power law in r-band luminosity of the form $A_1 \propto L^\beta$, with β differing between high and low L. Our redMaGiC low$-z$ sample provides a significant improvement constraints at the faint end of the A_1–L relation, where the slope is shallower than at the bright end (by a factor of several; see Fig. 9).

(ii) Amongst red galaxy samples, however, we find noticeable colour dependence in the IA–luminosity trend. This is most obvious in CMASS and redMaGiC high-z, which both lie below the bulk of previous measurements at a similar L. These differences can be qualitatively explained by differences in the colour space distributions. This raises potential questions about the sufficiency of a simple red/blue binary split for modelling IAs, and whether joint modelling of luminosity and colour dependence may be needed.

(iii) We find no statistically significant signal in our ELG sample on any scale. Using the combination of w_{g+} and w_{++}, we impose an upper limit on the large-scale NLA amplitude in ELGs at $|A_1| < 0.3$ (68 per cent CL). This is an improvement on the null constraint from WiggleZ at $|A_1| < 1.03$ (Mandelbaum et al. 2011).

(iv) The one-parameter NLA model is seen to fit all of our red galaxy samples reasonably on scales $r_p > 6 \, h^{-1}$ Mpc. In our redMaGiC samples, which give the highest signal-to-noise measurements, we do see deviations from the NLA prediction in the range 2–6 h^{-1} Mpc. These deviations are more pronounced in the higher redMaGiC redshift sample.

(v) Allowing additional flexibility via the TATT model, we can obtain a good fit to both redMaGiC samples on intermediate scales, $r_p = 2$–6 h^{-1} Mpc. We thus place constraints on the additional parameters (see Fig. 10).

(vi) We show that lensing and magnification can have a potentially significant impact on direct IA measurements. The extra terms are dominated by a galaxy–galaxy lensing like contribution δ_gG, and the magnification-lensing cross correlation, δ_μG. Together they make up ~ 2–20 per cent of the total signal, depending on the sample. This is relevant for our higher S/N samples, and will certainly be significant for future measurements of a similar kind, even those relying only on spectroscopic samples.

A weak lensing cosmology analysis is underway using the Dark Energy Survey Year 6 data, and similar efforts are ongoing on the KiDS legacy and HSC Y3 results. Understanding astrophysical systematics such as IAs, on both large and small scales, will clearly be important for the success of these ongoing cosmology projects, as well as future surveys such as the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), Euclid, and the Nancy Grace Roman Space Telescope. Our ability to accurately model IAs and mitigate their impact is, however, still somewhat limited; even given detailed information about the redshift and rest-frame colour of a sample (which typically is not available in a photometric survey), we do not have sufficient a priori understanding of the physical processes to predict the IA signal for an arbitrary selection of galaxies. We can,

however, make measurements of IAs in a range of samples, and map out the dependence on galaxy properties. In this way, we can start to build up a phenomenological understanding of IAs, which will feed into the next generation of analyses. The longer term goal is to develop more accurate models of IAs on all scales but also, ideally, to derive informative priors on their parameters. This paper aims to contribute to this task using some of the most constraining current data sets.

Our results provide a small step towards a more complete understanding of IAs in lensing surveys. In particular, we present results from new data sets that allow a substantially improved constraint on the faint end of the L–A_1 relation, and at intermediate redshifts. This is important, as the extrapolation into this regime is still a significant uncertainty in both model building and model sufficiency testing for future surveys. There are also a number of natural extensions to the work presented here. One obvious example is the development of a simple model that can account for both colour and redshift dependence in our red samples simultaneously. This is the focus of ongoing work.

Fig. 1 was created using SKYMAPPER, which is available at https://github.com/pmelchior/skymapper. The contour plots were made using GETDIST (Lewis 2019).

DATA AVAILABILITY STATEMENT

All processed data products created for this work (two-point functions, covariance matrices, and matched shape catalogues) are available on request. The DES Y3 redMaGiC, METACAL, and

GOLD catalogues have been public since 2022 February, and can be found at https://des.ncsa.illinois.edu/releases/y3a2. Catalogues for the eBOSS LRG and ELG samples (including randoms, but without shapes) are available from https://data.sdss.org/sas/dr16/eboss/lss/catalogs/DR16/. For the equivalent CMASS catalogues see https://data.sdss.org/sas/dr12/boss/lss/.

REFERENCES

Amon A. et al., 2022, Phys. Rev. D, 105, 023514
Asgari M. et al., 2021, A&A, 645, A104
Baldauf T., Smith R. E., Seljak U., Mandelbaum R., 2010, Phys. Rev. D, 81, 063531
Bautista J. E. et al., 2021, MNRAS, 500, 736
Becker M. R. et al., 2016, Phys. Rev. D, 94, 022002
Blanton M. R. et al., 2017, AJ, 154, 28
Blazek J., McQuinn M., Seljak U., 2011, J. Cosmol. Astropart. Phys., 2011, 010
Blazek J., Mandelbaum R., Seljak U., Nakajima R., 2012, J. Cosmol. Astropart. Phys., 2012, 041
Blazek J., Vlah Z., Seljak U., 2015, J. Cosmol. Astropart. Phys., 2015, 015
Blazek J. A., MacCrann N., Troxel M. A., Fang X., 2019, Phys. Rev. D, 100, 103506
Bolton A. S. et al., 2012, AJ, 144, 144
Bridle S., King L., 2007, New J. Phys., 9, 444
Bruzual G., Charlot S., 2003, MNRAS, 344, 1000
Catelan P., Kamionkowski M., Blandford R. D., 2001, MNRAS, 320, L7
Cawthon R. et al., 2022, MNRAS, 513, 5517
Codis S. et al., 2015, MNRAS, 448, 3391
Comparat J. et al., 2016, A&A, 592, A121
Dark Energy Survey Collaboration, 2016, Phys. Rev. D, 94, 022001
Dark Energy Survey Collaboration, 2022, Phys. Rev. D, 105, 023520
Dawson K. S. et al., 2013, AJ, 145, 10
Dawson K. S. et al., 2016, AJ, 151, 44
de Mattia A. et al., 2021, MNRAS, 501, 5616
Delubac T. et al., 2017, MNRAS, 465, 1831
Drlica-Wagner A. et al., 2018, ApJS, 235, 33
Eisenstein D. J. et al., 2011, AJ, 142, 72
Elvin-Poole J. et al., 2023, MNRAS, 523, 3649
Euclid Collaboration, 2021, A&A, 655, A44
Everett S. et al., 2022, ApJS, 258, 15
Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2019, Open J. Astrophys., 2, 10
Fortuna M. C., Hoekstra H., Joachimi B., Johnston H., Chisari N. E., Georgiou C., Mahony C., 2021a, MNRAS, 501, 2983
Fortuna M. C. et al., 2021b, A&A, 654, A76
Gatti M. et al., 2021, MNRAS, 504, 4312
Gunn J. E. et al., 2006, AJ, 131, 2332
Guo H. et al., 2019, ApJ, 871, 147
Hamana T. et al., 2020, PASJ, 72, 16
Heymans C. et al., 2012, MNRAS, 427, 146
Heymans C. et al., 2013, MNRAS, 432, 2433
Hikage C. et al., 2019, PASJ, 71, 43
Hilbert S., Xu D., Schneider P., Springel V., Vogelsberger M., Hernquist L., 2017, MNRAS, 468, 790
Hildebrandt H. et al., 2017, MNRAS, 465, 1454
Hildebrandt H. et al., 2020, A&A, 633, A69
Hinshaw G. et al., 2013, ApJS, 208, 19
Hirata C. M., Seljak U., 2004, Phys. Rev. D, 70, 063526
Hirata C. M., Mandelbaum R., Ishak M., Seljak U., Nichol R., Pimbblet K. A., Ross N. P., Wake D., 2007, MNRAS, 381, 1197
Hu W., 1999, ApJ, 522, L21
Hu W., Jain B., 2004, Phys. Rev. D, 70, 043009
Huff E., Mandelbaum R., 2017, preprint (arXiv:1702.02600)
Hutchinson T. A. et al., 2016, AJ, 152, 205
Jain B., Seljak U., 1997, ApJ, 484, 560
Jarvis M., Bernstein G., Jain B., 2004, MNRAS, 352, 338

Jarvis M. et al., 2021, MNRAS, 501, 1282
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
Joachimi B., Mandelbaum R., Bridle S. L., 2011, A&A, 527, A26
Joachimi B. et al., 2015, Space Sci. Rev., 193, 1
Joachimi B. et al., 2021, A&A, 646, A129
Johnston H. et al., 2019, A&A, 624, A30
Kaiser N., 1987, MNRAS, 227, 1
Kiessling A. et al., 2015, Space Sci. Rev., 193, 67
Kilbinger M. et al., 2013, MNRAS, 430, 2200
Kirk D. et al., 2015, Space Sci. Rev., 193, 139
Krause E., Eifler T., Blazek J., 2016, MNRAS, 456, 207
Krause E. et al., 2021, preprint (arXiv:2105.13548)
Landy S. D., Szalay A. S., 1993, ApJ, 412, 64
Lemos P. et al., 2021, MNRAS, 505, 6179
Lewis A., 2019, preprint (arXiv:1910.13970)
Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
MacCrann N., Blazek J., Jain B., Krause E., 2020, MNRAS, 491, 5498
Mackey J., White M., Kamionkowski M., 2002, MNRAS, 332, 788
Mandelbaum R. et al., 2005, MNRAS, 361, 1287
Mandelbaum R. et al., 2011, MNRAS, 410, 844
McDonald P., 2006, Phys. Rev. D, 74, 103512
McEwen J. E., Fang X., Hirata C. M., Blazek J. A., 2016, J. Cosmol. Astropart. Phys., 2016, 015
Morganson E. et al., 2018, PASP, 130, 074501
Pandey S. et al., 2020, Phys. Rev. D, 102, 123522
Pandey S. et al., 2022, Phys. Rev. D, 106, 043520
Porredon A. et al., 2021, Phys. Rev. D, 103, 043503
Prakash A. et al., 2016, ApJS, 224, 34
Prat J. et al., 2022, Phys. Rev. D, 105, 083528
Raichoor A. et al., 2017, MNRAS, 471, 3955
Raichoor A. et al., 2020, Res. Notes Am. Astron. Soc., 4, 180
Raichoor A. et al., 2021, MNRAS, 500, 3254
Reid B. et al., 2016, MNRAS, 455, 1553
Reyes R., Mandelbaum R., Gunn J. E., Nakajima R., Seljak U., Hirata C. M., 2012, MNRAS, 425, 2610
Rodríguez-Monroy M. et al., 2022, MNRAS, 511, 2665
Ross A. J. et al., 2020, MNRAS, 498, 2354
Rossi G. et al., 2021, MNRAS, 505, 377
Rozo E. et al., 2016, MNRAS, 461, 1431
Saito S., Baldauf T., Vlah Z., Seljak U. c. v., Okumura T., McDonald P., 2014, Phys. Rev. D, 90, 123522
Samuroff S. et al., 2019, MNRAS, 489, 5453
Samuroff S., Mandelbaum R., Blazek J., 2021, MNRAS, 508, 637
Sánchez C. et al., 2022, Phys. Rev. D, 105, 083529
Schaye J. et al., 2010, MNRAS, 402, 1536
Schmidt F., Rozo E., Dodelson S., Hui L., Sheldon E., 2009, ApJ, 702, 593
Schneider M. D., Bridle S., 2010, MNRAS, 402, 2127
Schneider P., van Waerbeke L., Kilbinger M., Mellier Y., 2002, A&A, 396, 1
Schwarz G., 1978, Ann. Stat., 6, 461
Secco L. F. et al., 2022, Phys. Rev. D, 105, 023515
Sevilla-Noarbe I. et al., 2021, ApJS, 254, 24
Sheldon E. S., Huff E. M., 2017, ApJ, 841, 24
Singh S., Mandelbaum R., 2016, MNRAS, 457, 2301
Singh S., Mandelbaum R., More S., 2015, MNRAS, 450, 2195
Smee S. A. et al., 2013, AJ, 146, 32
Smith R. E. et al., 2003, MNRAS, 341, 1311
Suchyta E. et al., 2016, MNRAS, 457, 786
Takada M., Hu W., 2013, Phys. Rev. D, 87, 123504
Takada M., Jain B., 2009, MNRAS, 395, 2065
Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, ApJ, 761, 152
Tamone A. et al., 2020, MNRAS, 499, 5527
Tonegawa M., Okumura T., Totani T., Dalton G., Glazebrook K., Yabe K., 2018, PASJ, 70, 41
Troxel M. A., Ishak M., 2015, Phys. Rep., 558, 1
Troxel M. A. et al., 2018, Phys. Rev. D, 98, 043528

van Daalen M. P., Schaye J., Booth C. M., Dalla Vecchia C., 2011, MNRAS, 415, 3649

Vlah Z., Chisari N. E., Schmidt F., 2020, J. Cosmol. Astropart. Phys., 2020, 025

von Wietersheim-Kramsta M. et al., 2021, MNRAS, 504, 1452

Wright E. L. et al., 2010, AJ, 140, 1868

Zacharegkas G. et al., 2022, MNRAS, 509, 3119

Zhang P., 2010, ApJ, 720, 1090

Zhou R. et al., 2020, Res. Notes Am. Astron. Soc., 4, 181

Zuntz J. et al., 2015, Astron. Comput., 12, 45

Zuntz J. et al., 2018, MNRAS, 481, 1149

APPENDIX A: MODELLING REDSHIFT SPACE DISTORTIONS AND IA ANISOTROPY

In this appendix we set out the formalism used to estimate the impact of RSDs and anisotropic IAs on our results (see also Singh & Mandelbaum 2016). One can write the redshift-space galaxy–galaxy power spectrum in terms of the (isotropic) real space equivalent in the form

$$P_{gg,s}(k) = (1 + \beta_a \mu)(1 + \beta_b \mu) P_{gg}(k), \tag{A1}$$

with galaxy samples a, b and $\beta_a \equiv f(z)/b_{g,a}$, the ratio of the logarithmic growth rate to the linear galaxy bias. The factor μ is the cosine of the angle between mode $\mathbf{k}$ and the axis of the line of sight $\hat{\mathbf{z}}$, $\mu = \hat{\mathbf{k}} \cdot \hat{\mathbf{z}}$. This is an approximation that applies on linear scales, but begins to break down at large k (Kaiser 1987). In general, one can decompose equation (A1) in terms of Legendre polynomials $\mathcal{P}_\ell$,

$$P_{gg,s}(k) = \left[\sum_{\ell=0}^{2} \alpha_{2\ell}^{ab} \mathcal{P}_{2\ell} \right] P_{gg}(k). \tag{A2}$$

That is, the sum of monopole, quadrupole, and hexadecupole contributions. Note that in the case that $\ell = 0$, $\beta = 0$, the above reverts to the isotropic case and $P_{gg,s} = P_{gg}$. The coefficients have the same form as equations 48–50 of Baldauf et al. (2010) for even values of ℓ,

$$\alpha_{2\ell}^{ab} = \begin{cases} 1 + \frac{1}{3}(\beta_a + \beta_b) + \frac{1}{5}\beta_a\beta_b & \ell = 0 \\ \frac{2}{3}(\beta_a + \beta_b) + \frac{4}{7}\beta_a\beta_b & \ell = 1 \\ \frac{8}{35}\beta_a\beta_b & \ell = 2 \end{cases}, \tag{A3}$$

and are zero otherwise. The configuration space equivalent of equation (A2) has a similar form

$$\xi_{gg,s}(r_p, \Pi) = \sum_{\ell=0}^{2} \alpha_{2\ell}^{ab} \mathcal{P}_{2\ell} \xi_{gg,2\ell}(r_p, \Pi), \tag{A4}$$

with

$$\xi_{gg,2\ell}(r_p, \Pi) = \frac{(-1)^\ell}{2\pi} \int k^2 P_{gg}(k) j_{2\ell}(kr) \mathrm{d}k. \tag{A5}$$

The integration kernel j_μ is a spherical Bessel function of the first kind of order μ, and it is this that determines the shape of each term. Putting these pieces together, and integrating over line-of-sight separation, one finally obtains the expression

$$w_{gg}(r_p) = \sum_{\ell=0}^{2} \frac{(-1)^\ell}{2\pi} \alpha_{2\ell} \mathcal{P}_{2\ell} \int_{-\Pi_{max}}^{\Pi_{max}} \int k^2 P_{gg}(k) j_{2\ell}(kr) \mathrm{d}k \mathrm{d}\Pi \tag{A6}$$

Although RSDs themselves do not have a significant impact on w_{g+} (see Singh et al. 2015), there is an analogous effect due to the projection of 3D shapes into 2D space. This suppresses the observed alignment strength at $|\Pi| > 0$, and so alters the shape of ξ_{g+} in the $r_p - \Pi$ plane


Figure A1. The fractional impact of RSDs and projection effects on projected galaxy–galaxy and galaxy–shape correlation functions. Note that the difference is defined as the magnitude of the difference between theory predictions (in either w_{g+} or w_{gg}, as labelled) with and without RSDs/projection effects $\Delta w_{ab} = |w_{ab}^{RSD} - w_{ab}^{no\,RSD}|$. The theory predictions are generated at our fiducial cosmology and $A_1 = 1$. We show both high- and low-z redMaGiC samples to illustrate the impact of (quite significant) differences in the redshift distributions.

(see the discussion in Section 4.6.2). The impact can be modelled in a very similar way to with RSDs (Singh & Mandelbaum 2016, section 2.3 and equation 13).

Fig. A1 shows the absolute impact of the additional RSD signal and the projection effect described above (note that RSDs and IA anisotropy work in opposite directions, and so the sign of the two eddects in Fig. A1 are different). RSDs have an impact on w_{gg} at the level of tens of per cent on scales $r_p > 6\ h^{-1}$ Mpc. We thus expect to be sensitive to their impact, and include them in our fiducial model. In the case of w_{g+} we see an impact on very large scales, dropping away below $\sim 70\ h^{-1}$ Mpc. Given that we impose an upper scale cut at $r_p = 70\ h^{-1}$ Mpc, due to possible large-scale systematics, we do not consider it necessary to include IA anisotropy in our fiducial model for w_{g+}.

APPENDIX B: NULL TESTS

Before carrying out our analysis, we carried our various validation tests. Among those was a null test, constructed by repeating our w_{g+} measurements, but using shapes measured at 45 deg to the tangential/radial direction. In the absence of systematics, this should return no signal.


Figure B1. Cross shear correlations. The measurements shown are computed in the same way as w_{g+} (see equation (16)), but with the galaxy shapes rotated by 45 deg. This is meant as a null test, since to first order neither IAs nor gravitational lensing produce such correlations. The coloured shaded regions show the estimated shape + shot noise uncertainty for each sample. The grey shaded bands indicate scales discarded in our two analysis set-ups. In each case, the quoted null χ^2 is computed on all scales $r_p < 70\ h^{-1}$ Mpc.

The results for our five samples are shown in Fig. B1. The error bands here are calculated assuming shape (and shot) noise only, using the observed number of galaxy pairs in each r_p bin. In the case of the two redMaGiC samples, we see a slight increase in (negative) power on very large scales. The reason for this apparent signal is not known for certain. We treat it as an unknown systematic, and simply choose to remove the affected scales. After imposing an upper limit at $r_p < 70\ h^{-1}$ Mpc, we find $w_{g\times}$ to be consistent with zero on all surviving scales. The null χ^2 values are shown for each sample in Fig. B1. Even in the case with the worst goodness of fit, redMaGiC high z, we find a $\chi^2/\mathrm{dof} = 24.1/16$, giving a corresponding p-value $p = 0.09$.

APPENDIX C: COMPARISON WITH LOWZ

As discussed in Section 5, we carry out several layers of pipeline testing and validation using LOWZ. The LOWZ LRG sample is useful for this, in that it is a relatively well-understood data set, which gives a high signal-to-noise w_{g+} signal. Crucially, there are also published IA measurements to which we can compare (Singh et al. 2015).

In addition to the data vector level comparison described in Section 5, we also use the LOWZ data to help validate our analytic covariance estimates. LOWZ covers a broad more or less contiguous footprint, making jackknife estimates viable. We divide that footprint into 100 patches using a k-means algorithm, and iteratively remeasure the whole data vector ($w_{gg} + w_{g+} + w_{++}$) in each. The diagonal elements of the resulting jackknife covariance matrix are compared to our analytic estimate in Fig. C1. As expected the latter is somewhat smoother. The two

Figure C1. A comparison of the square root of the LOWZ covariance diagonals obtained using two methods. From top we show w_{g+}, w_{gg}, and w_{++}. In each case the shaded grey regions indicate scales excluded from our large-scale IA fits. We see very good agreement between the two estimates on the scales of interest.

diverge slightly on very large scales, where the approximations behind the jackknife method break down. On the scales of interest, however, we see very good agreement.

We also carry out an end-to-end reanalysis of LOWZ using our pipeline. Starting with galaxy catalogues and randoms, we remeasure the joint data vector. Using our analytic covariance matrix, and the modelling pipeline set out in Section 4, we obtain parameter constraints. The results of this exercise are summarized in Fig. C2. In black we show the published IA and bias results from Singh et al. (2015); note that the fits for b_1 and A_1 were performed serially, and so we have a point with error bars instead of a full contour. The open blue contour shows the result of analysing the LOWZ data using our pipeline, but with all the analysis choices matched to those of Singh et al. (2015). These are detailed in Section 5, but include the choice of cosmology and the version of HALOFIT. We see good agreement in both parameters.

The filled contours then show the impact of switching to our analysis choices, assuming the NLA and TATT models. The former (dark purple) gives a very similar A_1 constraint to the original Singh et al. (2015) analysis. This is reassuring, in the sense that it suggests the new results from our pipeline are readily comparable with those in the literature. The lighter purple contours show the impact of opening up the TATT parameter space, and also extending the minimum scale in w_{g+} and w_{++} down to 2 h^{-1} Mpc. The marginalized A_1 constraint is broadened and shifted downwards slightly, primarily due to the degeneracy with A_2. It is interesting to briefly note here that although the contours on the extra parameters (A_2 and b_{TA}) are not symmetric about zero, they are totally consistent with zero. That is, the LOWZ data do not appear to require additional terms beyond the NLA model to describe scales down to $2h^{-1}$ Mpc.

Figure C2. 68 per cent and 95 per cent confidence intervals from our reanalysis of the SDSS LOWZ data. The black cross represents the published constraint on A_1 and galaxy bias from Singh et al. (2015). In blue (open contours) we show the result of fitting the LOWZ data vector using analysis choices matched to those of Singh et al. (2015). We recover the best-fitting A_1 well. The filled contours then show the result of switching to our preferred analysis settings, using the NLA and TATT models.

APPENDIX D: CONSTRAINTS ON GALAXY BIAS

In all samples, and for all fits, we include in our model free parameters for galaxy bias. Although the bias constraints are almost entirely dominated by the w_{gg} part of the data vector, we allow bias to vary alongside our IA parameters. Justified by the exercise in Section 4.6.1, our model includes two free parameters: b_1 and b_2 (there are additional terms in the expression for P_{δ_g}, but the model is fully specified by the two values; see Section 4.2).

The main galaxy bias results from each of our samples are presented in Table D1. In each case, the model provides a reasonable fit to the joint data vector. The redMaGiC numbers here are qualitatively consistent with those presented in the upper panel of Dark Energy Survey Collaboration (2022)'s fig. 8 (the pale purple points). That there are some small differences in the actual numbers in not surprising, given the different nature of the analysis (e.g. we are assuming a particular fixed cosmology). The values are roughly in line with the expectation for these sorts of galaxy samples.

Table D1. Constraints on galaxy bias from our various density tracer samples.

Sample	b_1	b_2
redMaGiC low-z	$1.59^{+0.01}_{-0.01}$	$-0.09^{+0.07}_{-0.07}$
redMaGiC high-z	$1.81^{+0.01}_{-0.01}$	$0.39^{+0.08}_{-0.07}$
eBOSS LRGs	$2.20^{+0.03}_{-0.03}$	$0.34^{+0.17}_{-0.16}$
eBOSS ELGs	$1.37^{+0.06}_{-0.03}$	$-0.84^{+0.74}_{-0.64}$
CMASS	$1.97^{+0.02}_{-0.02}$	$0.01^{+0.13}_{-0.15}$

APPENDIX E: REDSHIFT DEPENDENCE OF ALIGNMENTS IN RED GALAXIES

In this appendix we illustrate the redshift dependence of our red samples. Fig. E1 is the counterpart to Fig. 9, but showing the trend with redshift rather than luminosity. The colour scheme for the different samples is the same in the two.

For the sake of comparability, we define two narrow bins in L_r, and consider the redshift dependence in each. As explained in Section 6.2.1, these are centred on $\log L/L_0 \sim -0.3$ and $\log L/L_0 \sim -0.05$, respectively. The samples included in these two bins are shown in the upper/lower panels of Fig. E1 (labelled 'Low L' and 'High L'). The idea here is to separate inherent evolution in redshift (i.e. in a fixed sample with unchanging observable properties) from the evolution of galaxy selection with z.

As we can see here, there is no clear trend over the baseline of the samples, in either luminosity bin. If one fits a slope in redshift of the form $A_1(z) \propto [(1+z)/(1+z_0)]^{\eta_1}$, where $z_0 = 0.62$, the results are consistent with $\eta_1 = 0$. Specifically, we find $\eta_1 = -0.37 \pm 0.94$ in the lower L bin, and $\eta_1 = -0.05 \pm 0.73$ in the upper L bin.


Figure E1. IA strength as a function of estimated mean redshift in red galaxies. To avoid complication due to evolution in galaxy properties, we take only points within two narrow bands in L/L_0. See Section 6.2.1 for discussion. We show the best-fitting power-law slopes, fit to each panel; in both cases the power-law index η_1 is consistent with zero to $\ll 1\sigma$.

[1]*Department of Physics, Northeastern University, Boston, MA 02115, USA*

[2]*McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA*

[3]*NSF AI Planning Institute for Physics of the Future, Carnegie Mellon University, Pittsburgh, PA 15213, USA*

[4]*Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, UK*

[5]*Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA*

[6]*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

[7]*Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK*

[8]*Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA*

[9]*Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*

[10]*Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, OH 43210, USA*

[11]*Argonne National Laboratory, 9700 South Cass Avenue, Lemont, IL 60439, USA*

[12]*Institute for Astronomy, University of Hawai'i, 2680 Woodlawn Drive, Honolulu, HI 96822, USA*

[13]*Physics Department, University of Wisconsin-Madison, 2320 Chamberlin Hall, 1150 University Avenue Madison, WI 53706-1390, USA*

[14]*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA*

[15]*Instituto de Astrofisica de Canarias, E-38205 La Laguna, Tenerife, Spain*

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

[17]*Departamento de Astrofísica, Universidad de La Laguna, E-38206 La Laguna, Tenerife, Spain*

[18]*Center for Astrophysical Surveys, National Center for Supercomputing Applications, 1205 West Clark St, Urbana, IL 61801, USA*

[19]*Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 W. Green Street, Urbana, IL 61801, USA*

[20]*Physics Department, William Jewell College, Liberty, MO 64068, USA*

[21]*Department of Physics, Duke University, Durham, NC 27708, USA*

[22]*NASA Goddard Space Flight Center, 8800 Greenbelt Rd, Greenbelt, MD 20771, USA*

[23]*Institut d'Estudis Espacials de Catalunya (IEEC), E-08034 Barcelona, Spain*

[24]*Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, E-08193 Barcelona, Spain*

[25]*Kavli Institute for Particle Astrophysics & Cosmology, P. O. Box 2450, Stanford University, Stanford, CA 94305, USA*

[26]*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA*

[27]*CNRS, LPSC-IN2P3, Université Grenoble Alpes, F-38000 Grenoble, France*

[28]*Fermi National Accelerator Laboratory, P. O. Box 500, Batavia, IL 60510, USA*

[29]*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr, Pasadena, CA 91109, USA*

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

[31]*University Observatory, Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstr 1, D-81679 Munich, Germany*

[32]*School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, UK*

[33]*Mila – Quebec AI Institute, 6666 St-Urbain Street, 200 Montreal, QC, H2S 3H1, Canada*

[34]*Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA*

[35]*Instituto de Física Gleb Wataghin, Universidade Estadual de Campinas, Campinas, SP 13083-859, Brazil*

[36]*Institute for Astronomy, University of Edinburgh, Edinburgh EH9 3HJ, UK*

[37]*Department of Physics, University of Genova and INFN, Via Dodecaneso 33, I-16146 Genova, Italy*

[38]*Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Av. Complutense, 40, 28040 Madrid, Spain*

[39]*Jodrell Bank Center for Astrophysics, School of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL, UK*

[40]*SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA*

[41]*Department of Astronomy and Space Science, Sejong University, 05006 209, Neungdong-ro, Gwangjin-gu, Seoul, South Korea*

[42]*Brookhaven National Laboratory, Bldg 510, Upton, NY 11973, USA*

[43]*Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794, USA*

[44]*Département de Physique Théorique and Center for Astroparticle Physics, Université de Genève, 24 quai Ernest Ansermet, CH-1211 Geneva, Switzerland*

[45]*Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA*

[46]*Cerro Tololo Inter-American Observatory, NSF's National Optical-Infrared Astronomy Research Laboratory, Casilla 603, La Serena, R5GW+PH, Chile*

[47]*Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK*

[48]*CNRS, Institut d'Astrophysique de Paris, UMR 7095, F-75014 Paris, France*

[49]*Institut d'Astrophysique de Paris, Sorbonne Universités, UPMC Univ Paris 06, UMR 7095, F-75014 Paris, France*

[50]*Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

[51]*Astronomy Unit, Department of Physics, University of Trieste, via Tiepolo 11, I-34131 Trieste, Italy*

[52]*INAF – Osservatorio Astronomico di Trieste, via G. B. Tiepolo 11, I-34143 Trieste, Italy*

[53]*Institute for Fundamental Physics of the Universe, Via Beirut 2, I-34014 Trieste, Italy*

[54]*Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, D-21029 Hamburg, Germany*

[55]*Department of Physics, IIT Hyderabad, Kandi, Telangana 502285, India*

[56]*Institute of Theoretical Astrophysics, University of Oslo. P.O. Box 1029 Blindern, NO-0315 Oslo, Norway*

[57]*Instituto de Fisica Teorica UAM/CSIC, Universidad Autonoma de Madrid, E-28049 Madrid, Spain*

[58]*School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia*

[59]*Santa Cruz Institute for Particle Physics, Santa Cruz, CA 95064, USA*

[60]*Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA*

[61]*Australian Astronomical Optics, Macquarie University, North Ryde, NSW 2113, Australia*

[62]*Lowell Observatory, 1400 Mars Hill Rd, Flagstaff, AZ 86001, USA*

[63]*George P. and Cynthia Woods Mitchell Institute for Fundamental Physics and Astronomy, and Department of Physics and Astronomy, Texas A&M University, College Station, TX 77843, USA*

[64]*Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA*

[65]*Institució Catalana de Recerca i Estudis Avançats, E-08010 Barcelona, Spain*

[66]*Department of Physics and Astronomy and PITT PACC, University of Pittsburgh, PA 15260, USA*

[67]*Department of Astronomy, University of California, Berkeley, 501 Campbell Hall, Berkeley, CA 94720, USA*

[68]*Observatório Nacional, Rua Gal. José Cristino 77, Rio de Janeiro, RJ 20921-400, Brazil*

[69]*School of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, UK*

[70]*Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA*

This paper has been typeset from a TEX/LATEX file prepared by the author.