2dFLenS and KiDS: determining source redshift distributions with cross-correlations

Andrew Johnson,[1,2] Chris Blake,[1,2]★ Alexandra Amon,[3] Thomas Erben,[4]
Karl Glazebrook,[1] Joachim Harnois-Deraps,[3] Catherine Heymans,[3]
Hendrik Hildebrandt,[4] Shahab Joudaki,[1,2] Dominik Klaes,[4] Konrad Kuijken,[5]
Chris Lidman,[6] Felipe A. Marin,[1,2] John McFarland,[7] Christopher B. Morrison,[4,8]
David Parkinson,[9] Gregory B. Poole,[10] Mario Radovich[11] and Christian Wolf[12]

[1]*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, VIC 3122, Australia*
[2]*ARC Centre of Excellence for All-Sky Astrophysics (CAASTRO)*
[3]*Scottish Universities Physics Alliance, Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK*
[4]*Argelander Institute for Astronomy, University of Bonn, Auf dem Hugel 71, D-53121 Bonn, Germany*
[5]*Leiden Observatory, Leiden University, Niels Bohrweg 2, NL-2333 CA Leiden, the Netherlands*
[6]*Australian Astronomical Observatory, North Ryde, NSW 2113, Australia*
[7]*Kapteyn Astronomical Institute, Postbus 800, NL-9700 AV Groningen, the Netherlands*
[8]*Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195, USA*
[9]*School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia*
[10]*School of Physics, University of Melbourne, Parkville, VIC 3010, Australia*
[11]*INAF – Osservatorio Astronomico di Padova, via dell'Osservatorio 5, 35122 Padova, Italy*
[12]*Research School of Astronomy and Astrophysics, The Australian National University, Canberra, ACT 2611, Australia*

ABSTRACT

We develop a statistical estimator to infer the redshift probability distribution of a photometric sample of galaxies from its angular cross-correlation in redshift bins with an overlapping spectroscopic sample. This estimator is a minimum-variance weighted quadratic function of the data: a quadratic estimator. This extends and modifies the methodology presented by McQuinn & White. The derived source redshift distribution is degenerate with the source galaxy bias, which must be constrained via additional assumptions. We apply this estimator to constrain source galaxy redshift distributions in the Kilo-Degree imaging survey through cross-correlation with the spectroscopic 2-degree Field Lensing Survey, presenting results first as a binned step-wise distribution in the range $z < 0.8$, and then building a continuous distribution using a Gaussian process model. We demonstrate the robustness of our methodology using mock catalogues constructed from *N*-body simulations, and comparisons with other techniques for inferring the redshift distribution.

Key words: surveys – cosmology: observation – large-scale structure of Universe.

1 INTRODUCTION

Current and forthcoming photometric surveys aim to image a significant fraction of the sky.[1] In doing so, they will obtain the angular positions of millions of galaxies. Realizing the scientific potential of these surveys requires an estimate of the redshift distribution of the galaxies, which is important for connecting measurements – such as tomographic weak lensing (Hu [1999]; Huterer [2002]), the Integrated Sachs–Wolfe effect and angular power spectra – to the underlying cosmological model. In this work, we investigate a method to measure galaxy redshift distributions using angular cross-correlations (CCs) with an overlapping spectroscopic sample.

We outline the approach as follows. Consider two galaxy data sets: a spectroscopic sample with a known redshift distribution, and a photometric sample with an unknown redshift distribution. The samples overlap on the sky and in redshift. Since they are sampled from the same underlying density field, we expect that they will share a positive CC function regardless of galaxy attributes such as colour and luminosity. The amplitude of the angular CC will increase with the degree of overlap of the two samples (e.g.

★E-mail: cblake@swin.edu.au

[1] For example, deep optical imaging surveys currently being completed for the science goal of weak gravitational lensing include the KiDS (de Jong et al. [2015]), the Dark Energy Survey (Abbott et al. [2016]) and the HyperSuprimeCam imaging survey. Future such surveys will include those performed by the Large Synoptic Sky Telescope (LSST) and *Euclid*.

Ho et al. 2008; Newman 2008; Erben et al. 2009). Therefore, the redshift distribution of the photometric sample can be mapped out by dividing the spectroscopic sample into adjacent redshift bins, and for each bin measuring the angular CC with the photometric sample.[2] In this work, we use this technique to constrain the redshift distribution of galaxies in tomographic bins within the Kilo-Degree Survey (KiDS) (de Jong et al. 2015; Kuijken et al. 2015) using the spectroscopic 2-degree Field Lensing Survey (2dFLenS, Blake et al. 2016) to trace the surrounding large-scale structure.

Knowledge of the redshift distribution of the source galaxies is a critical component of a weak lensing analysis because it is required to calculate the expected weak lensing signal for a given cosmological model (e.g. Huterer et al. 2006; Ma, Hu & Huterer 2006; Kitching, Taylor & Heavens 2008). Uncertainty and bias in the source redshift distribution directly propagates to derived cosmological constraints as one of the most important astrophysical systematics. The required level of systematic error control increases in severity for future surveys: for 'Stage IV' dark energy experiments (Weinberg et al. 2013) such as the LSST and *Euclid*, in order to avoid a degradation of dark energy constraints by more than 50 per cent, the mean and standard deviation of the photometric redshift distribution need to be measured to an accuracy $\sim 0.002(1 + z)$ (Huterer et al. 2006; Newman et al. 2015).

Many approaches have been proposed for determining source redshift distributions. We define direct calibration methods as those that calibrate a mapping from the flux in photometric bands to a galaxy's redshift. Template-based approaches and machine learning algorithms both lie in this category,[3] but there are various factors that make the above level of accuracy difficult to obtain, including catastrophic photometric errors, completeness requirements for spectroscopic training samples and sample variance (Bernstein & Huterer 2010; Cunha et al. 2012, 2014; Newman et al. 2015). We will discuss these factors in the subsequent section. An alternative 'indirect calibration' approach, which we pursue in this study, is provided by CC methods. In particular, we focus on the extension and application of the optimal quadratic estimation method proposed by McQuinn & White (2013) (hereafter, MW13), testing this method using both simulations and data.

The outline of this paper is as follows. Section 2 introduces the strengths and weaknesses of calibration via CCs, and highlights the previous work in the field. In Section 3, we introduce the data sets we employ in this study: the KiDS, the 2dFLenS and mock catalogues built from *N*-body simulations. Section 4 introduces the background theory and Section 5 describes the quadratic estimator we employ to measure the redshift distribution of galaxies. We validate our methodology using mock catalogues in Section 6, and present the results of applying our methodology to data in Section 7. We conclude in Section 8.

2 STRATEGIES FOR PHOTO-*z* CALIBRATION

2.1 Motivations

The key point of distinction between direct and indirect calibration approaches is that the former requires spectroscopic redshifts for a

---

[2] As we will discuss below, there are additional effects that can correlate the two samples.
[3] Examples of machine learning algorithms include SkyNet (Graff et al. 2014), TPZ (Carrasco Kind & Brunner 2013), ANNz2 (Sadeh, Abdalla & Lahav 2016) and MLPQNA (Cavuoti et al. 2015). Examples of template-based methods include bpz (Benítez 2000) and eazy (Brammer, van Dokkum & Coppi 2008).

subsample of the full photometric sample, and this subsample needs to be representative of the full sample in both colour and magnitude space (Sánchez et al. 2014; Sadeh et al. 2016). This requires that the targeted spectroscopic sample should be highly complete, i.e. a secure redshift needs to be measured for >90 per cent of the subsample (Newman et al. 2015). To achieve this level of completeness, spectroscopic redshifts are required for faint and high-redshift galaxies that are abundant in deep imaging surveys. In contrast, for a CC analysis one is free to target any tracer of overlapping large-scale structure (most usefully, the brightest galaxies), circumventing this difficulty.

Achieving a high level of spectroscopic-redshift completeness for direct calibration methods presents a significant observational challenge, as the chance of obtaining a successful redshift is dependent on an object's magnitude. Therefore, spectra are typically obtained for a non-random subsample of the target catalogue. A useful example is the DEEP2 survey conducted on the DEIMOS spectrograph at Keck Observatory: for the highest redshift quality class, secure redshifts were only obtained for 60 per cent of the galaxies (Newman et al. 2013). Considering future surveys, the severity of this problem is demonstrated by the requirement for spectroscopic follow-up suggested by Newman et al. (2015): obtaining >90 per cent redshift completeness at $i = 25.3$ (LSST depth) would require more than 100 nights on Keck.

We note that the requirements for spectroscopic follow-up can be reduced by assigning weights to galaxies during the training phase of photometric-redshift calibration (Lima et al. 2008; Cunha et al. 2012; Sadeh et al. 2016). These weights are assigned based on the colour–magnitude phase-space distribution of both the parent photo-*z* sample and the follow-up spectroscopic sample. The result is that the weighted spec-*z* sample more closely matches the photo-*z* sample in colour–magnitude space. The extent to which this approach allows one to reduce the required completeness is currently a subject of study.

The above challenges are also relevant for machine learning algorithms, as they require separate training, testing and validation samples. Similar requirements exist for template-based approaches, where – even though templates can be derived synthetically in principle – assessing the resulting accuracy of the photo-*z* estimate requires a representative spec-*z* sample, which can also aid in the construction of accurate spectral templates. Moreover, deriving a Bayesian prior used in the fitting process (e.g. bpz Benítez 2000) requires a spec-*z* sample; this prior can strongly influence the final results (Sánchez et al. 2014).

Catastrophic errors, photo-*z* estimates z_p with $|z_p - z_{true}| \sim \mathcal{O}(1)$, present an additional issue for direct calibration methods. Such errors occur because, with only broad-band flux information, there exist degeneracies in galaxy colours such as the confusion between the Lyman and Balmer breaks. General studies of the consequences of catastrophic errors are presented by Bernstein & Huterer (2010) and Hearin et al. (2010). These outliers cannot be mitigated by re-weighting the sample.

2.2 Challenges

Although photo-*z* calibration by CC avoids some of the issues listed above, there remain significant challenges for this approach, some of which we outline below.

(i) Degeneracy with galaxy bias: CCs measure the combination $b(z) \times P(z)$, where $b(z)$ is the source galaxy bias factor and $P(z)$ is the source redshift probability distribution. Therefore, galaxy bias is degenerate with the redshift distribution

(Newman 2008). Calibrating the redshift-dependent galaxy bias requires extra probes (e.g. galaxy–galaxy lensing) or assumptions (e.g. the bias varies smoothly with redshift). This likely represents the dominant issue when constraining the source redshift distribution using CCs.

(ii) Cosmological dependence: the model CC function depends on both our guess of $P(z)$ and the cosmological model. This introduces a worrying circularity, as our aim is to test the cosmological model with measurements derived using $P(z)$.

(iii) Extra source of correlations: cosmic magnification introduces additional correlations and hence can bias measurements of the redshift distribution. In addition to changing the shape of galaxies, lensing changes their size and brightness, promoting fainter galaxies into a magnitude-limited survey and correlating foreground and background objects (Bernstein & Huterer 2010; McQuinn & White 2013; Duncan et al. 2014).

(iv) Spec-z coverage: the redshift range over which the source distribution can be reconstructed is limited by the redshift and areal coverage of the spectroscopic CC samples.

2.3 Developments

In this section, we summarize recent work on photometric calibration with angular CCs.

(i) Estimators: a number of estimators have been proposed for inferring the redshift distribution of a photometric sample using an overlapping spec-z sample (Newman 2008; Matthews & Newman 2010; Schulz 2010; McQuinn & White 2013; Ménard et al. 2013; Schmidt et al. 2013). For example McQuinn & White (2013) develop a quadratic estimator, while Schulz (2010), Ménard et al. (2013) and Schmidt et al. (2013) use maximum-likelihood approaches to infer the source distribution.

(ii) Self-calibration: dividing a photometric sample into redshift bins allows one to cross-correlate between bins. This correlation allows one to determine the contamination fraction for the sample, and potentially constrain other systematic errors (Padmanabhan et al. 2007; Erben et al. 2009; Benjamin et al. 2013; Choi et al. 2016).

(iii) Applications: the methodology presented by Ménard et al. (2013) and Schmidt et al. (2013) has been applied to estimate the redshift distributions of galaxies in the Sloan Digital Sky Survey (SDSS) (Rahman et al. 2015; Rahman, Ménard & Scranton 2016), the Cosmic Infrared Background (Schmidt et al. 2015), the Canada–France–Hawaii Telescope Legacy Survey (CFHTLS) (Scottez et al. 2016), infrared sources from Wide-field Infrared Survey Explorer (*WISE*) and the Two-Micron All-Sky Survey (2MASS), and radio from the Faint Images of the Radio Sky at Twenty cm (FIRST) survey (Ménard et al. 2013). Additionally, Hildebrandt et al. (2016) present the first application of this methodology to a cosmic shear analysis.

3 DATASETS

3.1 The Kilo-Degree Survey

The KiDS is a multiband imaging survey designed for weak gravitational lensing analyses (de Jong et al. 2015). The survey is being performed at the 2.6-m VLT Survey Telescope where, using the 300-mega-pixel wide-field camera OmegaCAM, images are taken

in four filters *ugri*. KiDS aims to image ~1500 deg^2 of the sky down to a limiting *r*-band magnitude of ~25.[4]

The first and second data releases of KiDS are presented by de Jong et al. (2015) and Kuijken et al. (2015). Based on these catalogues, gravitational lensing science analyses of the 100 deg^2 overlap area with the Galaxy And Mass Assembly survey (Driver et al. 2011) were undertaken by Viola et al. (2015), Sifón et al. (2015), van Uitert et al. (2016) and Brouwer et al. (2016), using matched-aperture *ugri* colours in conjunction with bpz to derive redshift probability distributions and hence the lensing efficiencies.

We performed our analyses using the third data release of KiDS, separately using the *r*-band 'KiDS-450' (Hildebrandt et al. 2016) and *i*-band 'KiDS-800' (Amon et al. in preparation) imaging data sets. Hildebrandt et al. (2016) carried out a careful analysis of different methods for calibrating the source redshift distribution and settled on a direct photo-z calibration scheme, verified by estimates based on clustering in a few square degrees of overlapping deep spectroscopic fields. Morrison et al. (2016) have recently presented a determination of the KiDS redshift distribution using small-scale CCs.

3.2 The 2-degree Field Lensing Survey

We map the large-scale structure in which our photometric sample is embedded using the spectroscopic 2dFLenS. We outline the basic properties of the survey here; full details are presented by Blake et al. (2016).

The principal aim of 2dFLenS is to expand the area of overlap between spectroscopic galaxy surveys and gravitational lensing imaging surveys. This facilitates two key science goals. First, it allows a joint analysis of lensing and galaxy redshift samples including all CC statistics (e.g. Cai & Bernstein 2012; Gaztañaga et al. 2012), with different applications presented by Amon et al. (in preparation) and Joudaki et al. (in preparation). Secondly, it allows the calibration of photometric-redshift distributions using CC techniques – which we present in this paper – and direct calibration techniques (Wolf et al. 2016).

2dFLenS was conducted on the Anglo-Australian Telescope over 53 nights in the 14B, 15A and 15B semesters. The two main target classes, selected from the VST-ATLAS Survey (Shanks et al. 2015), comprised ~40 000 Luminous Red Galaxies (LRGs) across a range of redshifts $z < 0.9$ selected by SDSS-inspired cuts (Dawson et al. 2013), and a magnitude-limited complete sample of ~30 000 objects in the range $17 < r < 19.5$ to assist with direct photometric calibration of the SkyMapper Survey (Wolf et al. 2016). In our study, we analyse the LRG sample, whose normalized redshift distribution (after merging the multiple target classes) is illustrated in Fig. 1.

2dFLenS observations cover an area of 731 deg^2. The overlap area with the imaging is currently limited by the progress of KiDS, which is still collecting data at the time of writing. Currently, the overlap area between 2dFLenS and the *i*-band (*r*-band) KiDS imaging is 431 deg^2 (152 deg^2), as shown in Fig. 2.

We map out the photo-z redshift distribution through CCs by dividing our spectroscopic sample into independent redshift bins. In choosing the width of these bins we balanced considerations of noise

---

[4] The *r*-band images are used for galaxy shape measurements because these are the deepest observations obtained in the best seeing conditions. The *ugri* bands have 5σ limiting magnitudes ~24.3, 25.1, 24.9, 23.8, respectively, in a 2 arcsec aperture (Hildebrandt et al. 2016).

Figure 1. The redshift probability distribution of 2dFLenS galaxies, displayed as both a histogram and Gaussian kernel density estimate. The multiple peaks arise because separate 2dFLenS LRG samples with different colour and magnitude selection criteria have been merged.


Figure 2. The overlap area between 2dFLenS and KiDS. 2dFLenS pointings are displayed as black dots (which are centres of circular fields of radius 1 deg) and the KiDS-800 i-band (KiDS-450 r-band) coverage is shown as blue (red) coloured tiles. The total area of overlap between 2dFLenS and the i-band (r-band) imaging is 431 deg^2 (152 deg^2). Each KiDS pointing has dimension 1 deg $\times$ 1 deg.

in the measurements with the desire for high-redshift resolution, defining 18 redshift bins of width $\Delta z = 0.05$ in the redshift range $0 < z < 0.9$.

3.3 Simulations

We tested the robustness of our methodology by constructing synthetic galaxy catalogues composed of overlapping photometric and spectroscopic samples, which allowed us to compare the redshift distributions reconstructed by our algorithm to the known input distributions. We generated these mock catalogues using the Scinet LIght Cone Simulations (SLICS) series of N-body simulations (Harnois-Déraps & van Waerbeke 2015) that have been produced using the CUBEP^{3}M code (Harnois-Déraps et al. 2013) using a WMAP9+BAO+SN cosmological parameter set: matter density $\Omega_m = 0.2905$, baryon density $\Omega_b = 0.0473$, Hubble parameter $h = 0.6898$, spectral index $n_s = 0.969$ and normalization $\sigma_8 = 0.826$. The box-size of the simulations is $L = 505\,h^{-1}$ Mpc. The simulations follow the non-linear evolution of 1536^3 particles inside a 3072^3 grid cube. For each simulation, the density field is output at 18 redshift snapshots in the range $0 < z < 3$, which are used to construct a survey cone spanning 60 deg^2. A spherical overdensity halo finder was executed on the particle data during the simulation run, and the resulting halo catalogues were post-processed to self-consistently sample the light-cone geometry. We sampled our mocks from these halo catalogues, as described further in Section 6.

4 PARAMETERIZATION AND MODELLING

In this section, we describe the redshift distribution parametrization and clustering model we adopt for our analysis. Throughout, we will assume the fiducial cosmological parameters of the SLICS simulations, stated above. For readability, we will begin with a heuristic description and then move on to a more rigorous treatment.

The intent of this work is to present a novel technique for measuring the redshift probability distribution of a given photometric sample of galaxies, $P^{(p)}(z)$. First, we need a method to parametrize this probability distribution. We do this by dividing the redshift range of the sample into step-wise bins and constraining the number of galaxies $N_i^{(p)}$ within each bin i, assuming that their probability distribution within the bin is constant. So if $W_i(z)$ is a normalized top hat filter ($\int W_i(z)\,dz = 1$) and $N_T^{(p)} = \sum_i N_i^{(p)}$ is the total number of photometric galaxies, then the probability distribution within each bin is $P_i^{(p)}(z) = W_i(z)$ and the total distribution is

$$P^{(p)}(z) = \sum_i (N_i^{(p)}/N_T^{(p)})\, W_i(z)\,, \tag{1}$$

which is normalized such that $\int P^{(p)}(z)\,dz = 1$.

The quantity of interest for constraining $P^{(p)}(z)$ is the cross-power spectrum between two samples of galaxies: specifically, between the full photometric sample (p) and a given redshift bin of the spectroscopic sample (s_i). We label this angular galaxy cross-power spectrum $C_{ps_i}^{(g)}(\ell)$, as a function of multipole ℓ. One can estimate this quantity by decomposing the projected density field for each sample

into spherical harmonic coefficients $\{p(\ell), s_i(\ell)\}$ for the photometric and spectroscopic sample, respectively. The statistical estimate is then $\widehat{C}_{ps_i}(\ell) = \langle p(\ell) \, s_i(\ell) \rangle$. In the remainder of this section, we will describe how we can model this quantity using our parametrization for $P^{(p)}(z)$, and convert this model from a power spectrum to the measured correlation function.

Using the small-angle (or 'Limber') approximation (Limber 1954), the matter cross-power spectrum between two redshift bins i and j, with redshift distributions $P_i(z)$, is

$$C_{ij}^{(m)}(\ell) = \int_0^\infty \mathrm{d}z \, P_i(z) \, P_j(z) \, \frac{\mathcal{P}(k, z)}{r^2(z) r_\mathrm{H}(z)} \,, \qquad (2)$$

where $\mathcal{P}(k, z)$ is the matter power spectrum at wavenumber k, $r(z)$ is the comoving distance to redshift z, $r_\mathrm{H}(z) \equiv \mathrm{d}r/\mathrm{d}z$ and $k = (\ell + 1/2)/r$.

To extend equation (2) to model the galaxy–galaxy power spectrum we need to model the galaxy biases for both samples – which we label $b^{(p)}(z)$, $b^{(s)}(z)$ – and to include a shot noise component. We will assume a linear relationship between the galaxy and matter density field, i.e. $\delta_\mathrm{g} = b \, \delta_\mathrm{m}$ [5], and model the galaxy bias in step-wise bins as

$$b^{(A)}(z) = b_i^{(A)} \quad \text{for } |z - z_i| < \Delta_i/2. \qquad (3)$$

(A) can be either (p) or (s), indicating the photometric or spectroscopic sample, respectively, and Δ_i is the width of the ith redshift bin. The galaxy cross-power spectrum can then be written as

$$C_{ps_i}^{(g)}(\ell) = \int_0^\infty \mathrm{d}z \, P^{(p)}(z) \, b^{(p)}(z) \, P_i^{(s)}(z) \, b_i^{(s)}(z) \, \frac{\mathcal{P}(k, z)}{r^2(z) r_\mathrm{H}(z)} + \omega_{ps_i} \,, \qquad (4)$$

where $P_i^{(s)}(z)$ is the probability distribution of the spectroscopic sample in the ith bin and ω_{ps_i} models the shot noise component.

We now wish to expand this expression in terms of $N_i^{(p)}$. Applying our parametrizations for $P^{(p)}(z)$, $P_i^{(s)}(z)$, $b^{(p)}(z)$ and $b_i^{(s)}(z)$ from equations (1)–(3), equation (4) reduces to

$$C_{ps_i}^{(g)}(\ell) = (N_i^{(p)}/N_\mathrm{T}^{(p)}) \, b_i^{(p)} \, b_i^{(s)} \, C_{ii}^{(m)}(\ell) + \omega_{ps_i} \,. \qquad (5)$$

Following a similar derivation, we compute the autocorrelations between the full photometric sample ($C_{pp}^{(g)}(\ell)$) and between the bins of the spectroscopic sample ($C_{s_i s_i}^{(g)}(\ell)$) as

$$C_{pp}^{(g)}(\ell) = \sum_i \left(\frac{N_i^{(p)} b_i^{(p)}}{N_\mathrm{T}^{(p)}} \right)^2 C_{ii}^{(m)}(\ell) + \omega_{pp} \,, \qquad (6)$$

$$C_{s_i s_i}^{(g)}(\ell) = \left(b_i^{(s)} \right)^2 C_{ii}^{(m)}(\ell) + \omega_{s_i s_i} \,. \qquad (7)$$

Assuming the Limber approximation, which ignores long-wavelength modes, the covariance between non-overlapping bins is zero: thus, $C_{s_i s_j}^{(g)}(\ell) = 0$ when $i \neq j$. For further details see MW13.

---

[5] This assumption requires explanation. First, the scales of relevance for the quadratic estimator are ~ 10 arcmin. This angular scale corresponds to a set of physical scales where we might expect linear galaxy bias to break down and introduce a systematic error. However, as emphasized by MW13, the smoothness of the weighting function implies that the Fourier modes being traced are on more linear scales than expected from this simplistic conversion. The mock catalogues provide a way for us to quantify the significance of this error. For future work, this investigation could be extended by up-weighting linear scales (reducing the overall constraining power) or introducing a non-linear galaxy bias component (at the cost of further complicating the methodology).

To model the shot noise components for equations (5)–(7), we assume Poisson statistics, neglecting non-Poisson contributions (Baldauf et al. 2013). Following this assumption, the shot noise components are $\omega_{A_i A_i} = N_i^{(A)}/\text{area}$ [steradian] and $\omega_{ps_i} = f_\mathrm{over} \, N_i^{(s)}/\text{area}$ [steradian], where f_over is the overlap fraction between the photo-z and spec-z sample.

For observational considerations we will switch to configuration space when applying our methodology to data. Thus, we need to transform equations (5)–(7) into configuration space. To simplify the final expressions we first transform the constant number count case of equation (2):

$$w_{ii}^{(m)}(\theta) \equiv \sum_\ell \frac{2\ell + 1}{4\pi} \, P_\ell(\cos\theta) \, C_{ii}^{(m)}(\ell) \,. \qquad (8)$$

Now the angular galaxy auto- and CC functions can be written as

$$w_{ps_i}(\theta) = (N_i^{(p)}/N_\mathrm{T}^{(p)}) \, b_i^{(p)} \, b_i^{(s)} \, w_{ii}^{(m)}(\theta) \,, \qquad (9)$$

$$w_{s_i s_i}(\theta) = (b_i^{(s)})^2 \, w_{ii}^{(m)}(\theta) \,, \qquad (10)$$

$$w_{pp}(\theta) = \sum_i \left(\frac{N_i^{(p)} b_i^{(p)}}{N_\mathrm{T}^{(p)}} \right)^2 w_{ii}^{(m)}(\theta) \,. \qquad (11)$$

To compute the various correlation statistics we use the public software chomp[6] introduced by Morrison & Schneider (2013).[7] This calculation requires as input the matter power spectrum for each redshift bin, which we model using the halofit code (Smith et al. 2003): the halofit parameters adopted are those fit by Takahashi et al. (2012). chomp computes the redshift evolution in each bin from linear theory using $\mathcal{P}(k, z) = D(z)^2 \mathcal{P}(k)$, where $D(z)$ is the growth function. This approximation is valid because the redshift bins we adopt are narrow.

We note a number of systematic modelling issues which could be improved in future analysis:

(i) Non-linear effects. We measure the angular correlation function to scales ~ 1 arcmin. On such scales non-linear effects become significant and the halofit model we adopt may become inaccurate. We also assume linear galaxy bias.

(ii) Bias evolution. The CC observable is the combination $b^{(p)}(z) P^{(p)}(z)$, such that our inference of $P^{(p)}(z)$ must depend on the redshift evolution of galaxy bias.

(iii) Flat $N(z)$ approximation. To derive the above equations we have approximated the redshift distributions using a step-wise parametrization, such that the redshift distribution is constant within each bin. This approximation will break down if there are steep gradients in the redshift distribution.

5 THE QUADRATIC ESTIMATOR

In this section, we outline the construction and properties of the 'quadratic estimator' we employ to measure the redshift distribution using CCs. This work extends that presented by MW13.

---

[6] https://github.com/morriscb/chomp

[7] We checked the accuracy of this code by comparing its output with our own calculations. For both the angular power spectrum and correlation function, the calculations agree. We adopt chomp because of its useful class-based structure.

5.1 Introduction

By quadratic estimator we are referring to a statistical estimator of a quantity, say N, based on a quadratic combination of the available data, $\boldsymbol{x}$. For example, $\widehat{N} = \boldsymbol{x}^T \boldsymbol{x}$. The symbol $(\widehat{\cdots})$ indicates a statistically estimated quantity: a value derived directly from data rather than the true value. Quadratic estimation is particularly relevant for Gaussian random fields as all the information content is contained within quadratic combinations of the data (second-order statistics).

To construct a quadratic estimator we need to (i) define the data, (ii) specify the quantity we wish to estimate and (iii) construct a method to combine the data that gives an estimate of the desired quantity – preferably this estimator will take advantage of all of the information content within the data, thus minimizing the final variance of the inferred parameter (such an estimator is said to satisfy the Cramer–Rao inequality and be optimal). Considering each of these points in turn:

(i) The data: we start by considering the spherical harmonic coefficients of the projected density fields, which we write as $\widehat{p}(\ell, m)$ and $\widehat{s}(\ell, m)$ for the photo-z and spec-z samples, respectively, where $\boldsymbol{s} \equiv s_i$ represents the coefficients for the ith spec-z bin. These coefficients are computed as follows. First, we define $n(\Omega)$ as the projected galaxy density field, where Ω indicates the angular position. The spherical harmonic coefficients are computed by projecting the density field on a basis of spherical harmonics (Y_ℓ^m):

$$\widehat{p}(\ell, m) = \frac{1}{\bar{n}} \int \mathrm{d}\Omega \, n(\Omega) \, Y_\ell^m(\Omega) \,. \tag{12}$$

We combine these coefficients into a single data vector $\boldsymbol{x} = \left(\widehat{p}(\ell, m), \widehat{s}(\ell, m) \right)$.

(ii) The estimated quantity: the quantity we wish to determine is the number count distribution of the photometric sample in step-wise bins, labelled $\widehat{N}_i^{(p)}$ (following the parametrization defined above).

(iii) The estimator: we begin by writing the estimator in the most general form possible: $\widehat{N}_i = \boldsymbol{x}^T \boldsymbol{E}_i \boldsymbol{x} - c_i$, where $\boldsymbol{E}_i$ is a symmetric matrix and c_i is a constant (Tegmark 1997; Bond, Jaffe & Knox 1998; Dodelson 2003). These free parameters will be fixed by imposing various conditions on the estimator. Rather than making a single estimate, we can iterate until we are satisfied with the convergence. Setting $[\widehat{N}_i]_{\text{last}}$ as our initial guess, the updated estimator is

$$\widehat{N}_i = [\widehat{N}_i]_{\text{last}} + \boldsymbol{x}^T \boldsymbol{E}_i \boldsymbol{x} - c_i \,. \tag{13}$$

Requiring that the estimator is unbiased and optimal one can solve for both free parameters. Being unbiased implies that an ensemble average of the estimates converges to the input or true value, $\langle \widehat{N}_i \rangle = N_i^{\text{true}}$. Being optimal implies that the estimator minimizes the variance, viz., $\langle \widehat{N}_i^2 \rangle - \langle \widehat{N}_i \rangle^2$ is minimized. The final form of the estimator is (Tegmark 1997)

$$\widehat{N}_i = [\widehat{N}_i]_{\text{last}} + \frac{1}{2} \sum_j [\boldsymbol{F}^{-1}]_{ij} \left[\boldsymbol{x}^T \boldsymbol{Q}_j \boldsymbol{x} - \mathrm{Tr}(\boldsymbol{Q}_j \boldsymbol{A}) \right] \,, \tag{14}$$

where

$$\boldsymbol{Q}_j = \boldsymbol{A}^{-1} \boldsymbol{A}_{,j} \boldsymbol{A}^{-1} \,, \tag{15}$$

$\boldsymbol{A} \equiv \langle \boldsymbol{x} \boldsymbol{x}^T \rangle$ is the covariance matrix of the data, and its derivative is $\boldsymbol{A}_{,\alpha} = \partial \boldsymbol{A}/\partial N_\alpha$. We note that implicitly $\boldsymbol{A}$ is a function of both ℓ

and m, i.e. $\boldsymbol{A} = \boldsymbol{A}(\ell, m)$. Assuming many modes are included one can approximate $\boldsymbol{F}$ as the Fisher matrix[8]

$$F_{ij} = \frac{1}{2} \sum_{\ell,m} \mathrm{Tr} \left[\mathbf{A}^{-1} \mathbf{A}_{,i} \, \mathbf{A}^{-1} \mathbf{A}_{,j} \right] \,. \tag{16}$$

When this assumption is violated the Fisher matrix will be biased by sample variance. We do not expect this assumption to have a significant effect on our results.

Note, from the previous section (equations 5–7), $\boldsymbol{A}$ is known. It is the full covariance matrix between the spec-z and photo-z samples, including autocorrelations. In particular,

$$\begin{aligned} \boldsymbol{A} = \langle \boldsymbol{x} \boldsymbol{x}^T \rangle &= \left\langle \begin{pmatrix} \widehat{p}(\ell, m) & \widehat{s}(\ell, m) \end{pmatrix} \begin{pmatrix} \widehat{p}(\ell, m) \\ quad\widehat{s}(\ell, m) \end{pmatrix} \right\rangle \\ &= \begin{pmatrix} C_{pp}^{(g)}(\ell) & C_{ps}^{(g)}(\ell) \\ C_{sp}^{(g)}(\ell) & C_{ss}^{(g)}(\ell) \end{pmatrix} \,. \end{aligned} \tag{17}$$

Moreover, from equation (5), the derivative of the off-diagonal terms is $\partial A_{0i}/\partial N_i = b_i^{(p)} b_i^{(s)} N_i^{(s)} C_{s_i s_i}(\ell)$.

A more intuitive derivation of equation (14) can be found by applying the Newton–Raphson method to the Gaussian likelihood function of $\boldsymbol{x}$, $\mathcal{L}(\boldsymbol{x})$. The basic idea is to solve, iteratively, for roots of $\partial \ln \mathcal{L}/\partial N_i$: the roots indicate a maximum of the likelihood function. For this alternative derivation we refer the reader to Bond et al. (1998).

5.2 Revised form of the estimator

In this section and the subsequent one, we present important yet rather tedious mathematics; therefore, in the aid of readability we offer a quick summary.

The primary purpose of these sections is to present the analytic form of equation (14). We derive this expression by computing the tensor $\boldsymbol{Q}$ and the Fisher matrix $\boldsymbol{F}$ using equation (17) combined with the results from Section 4. We re-visit this derivation, which was initially presented in MW13, because we find a number of corrections to the final form of the quadratic estimator: we show that the most general expression for $\widehat{N}$ presented in MW13 is biased, such that, $\langle \widehat{N}_i \rangle \neq N_i^{\text{true}}$. However, we note that in the limit where shot noise dominates (labelled the 'Schur limit' by MW13) their expression for the estimator becomes unbiased. We present the updated result for the harmonic-space estimator in equation (20) and its extension to configuration space in equation (27). We note that none of the numerical calculations in MW13 are affected, since these require the Fisher matrix and not the form of the estimator.

Our revisions can be understood as follows. First, our general form of the quadratic estimator (i.e. equation 14) differs from that presented in MW13:

$$\widehat{N}_i = [\widehat{N}_i]_{\text{last}} + \underbrace{\frac{1}{2} \sum_j [\boldsymbol{F}^{-1}]_{ij} \left[\boldsymbol{x}^T \boldsymbol{Q}_j \boldsymbol{x} - \mathrm{Tr}(\boldsymbol{A}^{-1} \boldsymbol{A}_{,j}) \right]}_{\text{Eq. (16) from MW13}} \,. \tag{18}$$

Importantly, equation (18) is a simplified version of equation (14). The two expressions agree because of the relation $\mathrm{Tr}(\boldsymbol{Q}_j \boldsymbol{A}) = \mathrm{Tr}(\boldsymbol{A}^{-1} \boldsymbol{A}_{,j})$. In order to simplify the derivation, MW13 neglect all derivatives of A_{00} – note, we also make this approximation. And

---

[8] The Fisher matrix is an approximation of the curvature matrix, equal to its ensemble average. For details see Bond et al. (1998).

such terms ($A_{00,i}$) occur in the expression for $\boldsymbol{Q}$. Thus, by dropping these terms, one is making an approximation of $\boldsymbol{Q}$, $\boldsymbol{Q}^{\text{approx}}$. This approximation breaks the relation that equates equation (18) to equation (14). So, $\text{Tr}(\boldsymbol{Q}_j^{\text{approx}}\boldsymbol{A}) \neq \text{Tr}(\boldsymbol{A}^{-1}\boldsymbol{A}_{,j})$. Therefore, when neglecting the derivatives of A_{00}, equation (14) needs to be the starting point of the derivation. A number of non-trivial corrections to the estimator proposed by MW13 result from this starting point.

5.3 Harmonic-space estimator

To simplify the expressions that follow, using the notation of MW13 we define the 'Schur' parameter S as

$$S \equiv A_{00}\left(A_{00} - \sum_i \frac{A_{0i}^2}{A_{ii}}\right)^{-1} = \left(1 - \sum_i r_i^2\right)^{-1},\tag{19}$$

where the coefficients of $\boldsymbol{A}$ are defined in equation (17). Additionally, we define $r_i(\ell)$ as the CC coefficient between the photo-z sample and the ith redshift bin of the spec-z sample: $r_i(\ell) \equiv A_{0i}/(A_{00}A_{ii})^{1/2}$. Adopting these definitions, the full estimator can be written as

$$\widehat{N}_i^{(p)} = [\widehat{N}_i^{(p)}]_{\text{last}} + \sum_j (F^{-1})_{ij} \sum_{\ell,m} \left(\frac{S}{A_{00}A_{jj}}\right) \frac{\partial A_{0j}}{\partial p_j}$$

$$\left[\sum_k \left(\delta_{jk}^K + 2Sr_jr_k\sqrt{\frac{A_{jj}}{A_{kk}}}\right)(\widehat{p}\,\widehat{s}_k - A_{0k})\right.$$

$$-\sum_k \frac{A_{0k}}{A_{kk}}\left(\delta_{jk}^K + Sr_jr_k\sqrt{\frac{A_{jj}}{A_{kk}}}\right)(\widehat{s}_k\,\widehat{s}_k - A_{kk})$$

$$\left. -\frac{SA_{0j}}{A_{00}}(\widehat{p}\,\widehat{p} - A_{00})\right],\tag{20}$$

where δ_{ij}^K is the Kronecker delta. In this expression the (ℓ, m) dependence of the multipole coefficients is implicit, so, $\widehat{s}_k = \widehat{s}_k(\ell, m)$ and $\widehat{p} = \widehat{p}(\ell, m)$.

One can check this expression converges to the input theory as follows. First, we write the correction term as $\delta N_i \equiv \widehat{N}_i^{(p)} - [\widehat{N}_i^{(p)}]_{\text{last}}$. Now, assuming the input theory is correct, $\langle \widehat{p}\,\widehat{s}_i \rangle = A_{0i}$, $\langle \widehat{p}\,\widehat{p} \rangle = A_{00}$ and $\langle \widehat{s}_i\,\widehat{s}_j \rangle = A_{ij}$. Then, following some algebra, one can show equation (20) implies $\langle \delta N_i \rangle = 0$, thus, proving that the estimator will converge. For the equivalent equation from MW13, one can show $\langle \delta N_i \rangle \neq 0$.

The Fisher matrix (equation 16) remains unchanged from MW13, where

$$F_{ij} = \sum_{\ell,m} \frac{S}{A_{00}}\left(\frac{\delta_{ij}^K}{A_{ii}} + 2S\sqrt{\frac{r_i^2r_j^2}{A_{ii}A_{jj}}}\right)[A_{0i}]_{,i}\,[A_{0j}]_{,j}.\tag{21}$$

In the limit where shot noise dominates [i.e. where $r_i(\ell) \approx 0$ and $S \approx 1$] and neglecting autocorrelations, our result (equation 20) agrees with equation (36) from MW13.

5.4 Configuration-space estimator

When analysing observational data we will work exclusively in configuration space, which allows us to avoid difficulties with complex survey geometries (in the future this may not be necessary, see Alsing et al. 2016; Köhlinger et al. 2016). Thus, in order to match our theory with observations, we need to convert our estimator in

equation (20) from harmonic to configuration space. This conversion is simplified by the following relation (MW13):

$$\sum_{l,m} v_i(l)\left[\widehat{p}(l,m)\widehat{s}_i(l,m) - N_i^{(p)}\right]$$

$$= 8\pi^2 \int \text{d}x\, v_i(x)\,\widehat{w}_{ps_i}(x)$$

$$\approx 8\pi^2 \sum_\alpha \Delta\theta_\alpha\theta_\alpha\, v_i(\theta_\alpha)\,\widehat{w}_{ps_i}(\theta_\alpha),\tag{22}$$

where $\widehat{w}_{ps_i}(\theta)$ is the observed angular CC function, θ is the angular separation scale and $x = \hat{n} \cdot \hat{n}' \equiv \cos\theta$. Here, we have explicitly subtracted the shot noise component.[9] Equation (22) is valid for an arbitrary function $\boldsymbol{D}(l)$, which is related to $\boldsymbol{D}(\theta)$ by

$$D_i(\theta) = \sum_\ell \left(\frac{2\ell + 1}{4\pi}\right) D_i(\ell)\, P_\ell(\cos\theta),\tag{23}$$

where P_ℓ are the Legendre polynomials.

Our measurements of the angular correlation functions are made in bins of width $\Delta\theta_\alpha$, with central values θ_α. These values set the properties of the summation in equation (22). Note, because the kernel ($= \theta_\alpha v_i(\theta_\alpha)\widehat{w}_{ps_i}(\theta_\alpha)$) is not a slowly varying function, a narrow θ spacing ($\Delta\theta_\alpha$) is needed to accurately approximate this integral.

Now, to convert equation (20) into configuration space we first re-write the estimator in terms of four weighting functions defined as

$$D_i(l) \equiv \left(\frac{S}{A_{00}A_{ii}}\right)\frac{\partial A_{0i}}{\partial p_i},\tag{24}$$

$$E_{ij}(l) \equiv D_i(l) \times 2Sr_ir_j\sqrt{\frac{A_{ii}}{A_{jj}}}\tag{25}$$

and

$$H_i(l) \equiv D_i(l) \times \frac{A_{0i}}{A_{ii}},\quad G_i(l) \equiv D_i(l) \times S\frac{A_{0i}}{A_{00}}.\tag{26}$$

Finally, using equation (22) one finds that our estimator in configuration space takes the form

$$\widehat{N}_i^{(p)} = [\widehat{N}_i^{(p)}]_{\text{last}} + 8\pi^2 \sum_j (F^{-1})_{ij}$$

$$\sum_\alpha \Delta\theta_\alpha\,\theta_\alpha\left[D_j(\theta_\alpha)\left\{\widehat{w}_{ps_j}(\theta_\alpha) - w_{ps_j}(\theta_\alpha)\right\}\right.$$

$$+H_j(\theta_\alpha)\,\widehat{w}_{s_js_j}(\theta_\alpha) - G_j(\theta_\alpha)\,\widehat{w}_{pp}(\theta_\alpha)$$

$$\left.+\sum_k E_{jk}(\theta_\alpha)\left\{\widehat{w}_{ps_k}(\theta_\alpha) - w_{ps_k}(\theta_\alpha) + \frac{1}{2}\widehat{w}_{s_ks_k}(\theta_\alpha)\right\}\right],\tag{27}$$

where $\{\boldsymbol{D}, \boldsymbol{E}, \boldsymbol{G}, \boldsymbol{H}\}$ have been transformed to configuration space using equation (23). Ignoring both the bin-to-bin correlations, such that the Fisher Matrix is diagonal, and the autocorrelation terms, the second term in equation (27) becomes (MW13)

$$8\pi^2(F^{-1})_{ii} \sum_\alpha \Delta\theta_\alpha\,\theta_\alpha\left[D_i(\theta_\alpha)\left\{\widehat{w}_{ps_i}(\theta_\alpha) - w_{ps_i}(\theta_\alpha)\right\}\right].$$

The result is now much more intuitive. The $N(z)$ is reconstructed from a weighted minimization of $\{\widehat{w}_{ps_i}(\theta_\alpha) - w_{ps_i}(\theta_\alpha)\}$, where the

---

[9] We note that incorrectly modelling the shot noise component will introduce a bias into the final measurements.

Figure 3. The weights $D_i(\theta)$ defined in Section 5.4 (left-hand panel) and the combination $D_i(\theta) \times w_{ps_i}(\theta)$ (right-hand panel) for the mock catalogue tests. The weights and correlation functions are derived for the $z = 0.225$ (blue) and $z = 0.525$ (red) redshift bins, for illustration. The combination $D_i(\theta) \times w_{ps_i}(\theta)$ represents the angular scales used by the estimator to measure the redshift probability distribution, in accordance with equation (27).

weights are given by $D_i(\theta)$. Note that the scales that contribute most to the estimator are represented by the combination $D_i(\theta)w_{ps_i}(\theta)$.

To illustrate the angular sensitivity of the estimator, in Fig. 3 we plot these weights for the mocks introduced in the next section. We find that the weights peak at $\theta \sim 2$ arcmin. However, as emphasized by MW13, the breadth and smoothness of the weighting function implies that the Fourier modes being traced are on quasi-linear scales: sharp cuts in angle have a greater sensitivity to non-linearity than a smoother filter.

6 APPLICATION TO SIMULATIONS

In this section, we test our methodology using 20 sets of mock halo catalogues[10] created from the *N*-body simulations described in Section 3.3. For each 60 deg^2 simulation we generated a uniform redshift distribution of mock spectroscopic sources within the range $0.1 < z < 0.9$, adopting an angular density of 1000 sources deg^{-2}. In addition, we sampled mock photometric galaxies using a Gaussian redshift distribution with mean 0.5 and standard deviation 0.1 with a density of 1 source arcmin^{-2} (which roughly mimicks a typical tomographic bin in KiDS). For the purposes of this test we generated each sample by randomly sampling from the halo catalogue at each redshift, such that the bias factors of the photometric and spectroscopic samples are expected to be the same (this would not necessarily be true for a real data sample). We then performed a CC analysis dividing the spectroscopic sources into 16 redshift bins of width $\Delta z = 0.05$. For further details on the mock catalogues we refer the reader to Blake et al. (2016), section 6.1.

We note that each individual mock provides constraints comparable in precision to the observational data sets used in our analysis (each mock realization contains $\sim 60\,000$ spec-z galaxies, compared to $\sim 40\,000$ spec-z galaxies in 2dFLenS).

6.1 Auto- and cross-correlation measurements

We measured the angular autocorrelation and CC functions in our analysis using the Landy–Szalay estimator (Landy & Szalay 1993),

[10] We find 20 mock catalogues is sufficient for the level of error we wish to investigate.

generating random catalogues 10 times larger than the data distribution. For example, the CC between samples i and j is

$$w_{i,j}(\theta) = \frac{(D_i D_j)_\theta}{(R_i R_j)_\theta} \frac{N_{R,i} N_{R,j}}{N_{D,i} N_{D,j}} - \frac{(D_i R_j)_\theta}{(R_i R_j)_\theta} \frac{N_{R,i}}{N_{D,i}}$$
$$- \frac{(D_j R_i)_\theta}{(R_i R_j)_\theta} \frac{N_{R,j}}{N_{D,j}} + 1 \,, \tag{28}$$

where $(D_i D_j)_\theta$, $(D_i R_j)_\theta$ and $(R_i R_j)_\theta$ are the respective pair counts between the two data samples, the data and random samples, and the two random samples, as a function of the angular separation θ. The number counts of the data sample i and the random sample j are $N_{D,\,i}$ and $N_{R,\,j}$, respectively. The correlation functions are all measured using 30 equally logarithmically spaced angular bins between 0.01 and 1 deg. We estimate the errors in the measurements using jackknife re-sampling, although these errors are not used in the quadratic estimation process.

For each mock catalogue we measured the following statistics: the autocorrelation of the photometric galaxies, $w_{pp}(\theta)$; the 16 spectroscopic autocorrelations in each redshift bin, $w_{s_i s_i}(\theta)$ and the 16 photometric-spectroscopic CCs, $w_{ps_i}(\theta)$, where the index i runs across the 16 redshift bins of the spectroscopic sample.

6.2 Applying the estimator to the mocks

In this section, we apply the quadratic estimator to the mock catalogues and infer results for the redshift probability distribution $P^{(p)}(z)$ across the spectroscopic bins. For the purposes of this test we only use the angular CC functions as inputs to the quadratic estimator. Thus, we apply equation (27) and drop the autocorrelation terms. We note that in this limit the estimator remains unbiased, as the autocorrelation terms cancel. As discussed in the next section, we use the autocorrelations of the spec-z sample separately, to measure the redshift evolution of the galaxy bias.

Every iteration of the quadratic estimator returns a correction term to the probability distribution δP_i computed from the previous best-guess estimate of the redshift distribution. We define the estimator to be converged once the condition $\sum_i \delta P_i < 5 \times 10^{-3}$ is met. At this level of accuracy, the estimated correction is on average an order of magnitude smaller than the error for a given redshift bin, i.e. $\delta P_i / \sigma(P_i) \approx 0.1$.

Figure 4. Examination of the convergence of the estimator. The blue points show the measured angular CC functions for a single mock catalogue for all the spectroscopic redshift bins. Each panel corresponds to a separate redshift bin and $\omega^{(p,s_i)}$ indicates a CC between the photo-z sample and the ith bin in the spec-z sample. The red lines show the predictions of the CC functions using the true underlying $P(z)$ – a Gaussian with mean 0.5 and standard deviation 0.1. The blue lines show our predictions using the inferred $\hat{P}(z)$ from the quadratic estimator, which are seen to closely track the measurements. We note the different y-axis ranges in the panels.

After each iteration of the quadratic estimator, we also enforce the normalization condition $\sum_i P_i = 1$. This produces an overall amplitude shift that is minimal in most cases, although we do expect a small bias to be introduced when imposing this constraint because the normalization condition holds only for the underlying P_i, not the estimated $\hat{P}$.

6.3 Convergence of the estimator

We can assess the convergence of the estimator and the accuracy of the modelling by comparing the observed, reconstructed and model angular CC functions obtained from a single (representative) mock catalogue, as shown in Fig. 4. Each panel within this figure illustrates the CC function $w_{ps_i}(\theta)$ for one of the spectroscopic redshift bins, starting from the lowest redshift in the top left corner. We overplot the model prediction for the angular CCs as the red lines, derived using equation (10) and using the bias of the samples determined as discussed in Section 6.5 below. The blue lines in Fig. 4 show the reconstructed angular CC functions, which we obtain by using the

recovered photo-z redshift distribution, $\hat{P}_i^{(p)}$, to compute the angular CC via

$$\widehat{w}_{ps_i}(\theta) = b_i^{(s)} b_i^{(p)} P_i^{(s)} \hat{P}_i^{(p)} w_{s_i s_i}(\theta). \tag{29}$$

We can now assess the convergence of the estimator by comparing the reconstructed predictions (blue lines) to the mock measurements (blue points). Thus, from Fig. 4, we observe that the combination $\{\widehat{w}_{ps_i}(\theta_\alpha) - w_{ps_i}(\theta_\alpha)\}$ is being successfully minimized. One should keep in mind the effective θ-dependent weights being applied to this minimization, as shown by Fig. 3. Comparing CC functions is a useful validation of the estimator, as this test is less sensitive to inaccuracies in modelling the correlation functions and galaxy bias than comparing the inferred $P_i^{(p)}$ distribution to the input[11].

---

[11] For example, in a situation where the model overpredicts the amplitude of the mock correlation function, the estimator will respond by reducing $P_i^{(p)}$ relative to the true value. However, in this case $\widehat{w}_{ps_i}$ should still agree with the simulation prediction.

Figure 5. Test of the accuracy of our CC model. The blue points show the mean angular CC function measured from 20 mock catalogues for all the spectroscopic redshift bins, with the error taken from a single mock. The red lines show our model predictions, using the cosmological parameters of the simulation and the mean bias factors. Each panel corresponds to a separate redshift bin and $\omega^{(p,s_i)}$ indicates a CC between the photo-z sample and the ith bin in the spec-z sample. The bar, $\bar{\omega}$, is simply a reminder that we are averaging over mocks when determining these results.

6.4 Accuracy of the cross-correlation model

Systematic errors in the modelling of the CC statistics – for example, due to the breakdown of an assumption such as linear galaxy bias – could propagate into a bias in the inferred photometric redshift distribution. We assessed this potential source of systematic error by comparing the average of the mock CC function measurements to our model predictions, the results of which are shown in Fig. 5. The models are generated using the galaxy bias determined for the spectroscopic sample in Section 6.5 below, and in this figure we plot the average error for individual mocks. We conclude that our modelling of the angular CC function is sufficient at the level of statistical errors present in a single mock catalogue.

However, if we instead use the error in the mock mean (dividing by $\sqrt{20}$) we observe some tension between the simulation predictions and the analytic modelling of the correlation functions. These tests reveal that there are non-linear effects in the mocks not captured by our model (e.g. non-linear galaxy bias), and also that the jack-knife errors do not fully capture the scatter in the measurements. These issues could be mitigated by restricting our analyses

to larger scales, and by deducing the statistical errors using a dispersion in the mocks, rather than by jack-knife techniques. However, we leave such investigations to future work.

6.5 Measuring the galaxy bias factors

In order to test whether our quadratic estimation pipeline recovers the input source redshift distribution of the mocks, we also require the redshift evolution of the galaxy bias factors, which we have arranged (by sampling haloes in the same mass range) to be the same for the spectroscopic and photometric samples. We determined the redshift evolution of this bias using the autocorrelation function measurements of each spec-z sample, using chi-squared minimization to fit for $b_i^{(s)}$.

We note that the effects of noise will cause fluctuations between $b_i^{(s)}$ and $b_i^{(p)}$. In particular, given the significantly lower number density of spec-z galaxies relative to the photo-z sample, the bias measurements of the spec-z sample will be less accurate. We overcome this by averaging the $b_i^{(s)}$ values over the 20 mocks.

Figure 6. The inferred redshift probability distribution for nine independent mock catalogues. The green points show the reconstructed galaxy redshift distributions estimated by quadratic estimation ($\hat{P}$), with error bars derived from the relevant Fisher matrix. The green bands show the 95 per cent confidence intervals for a Gaussian process model trained using the green points. The black line shows the input redshift probability distribution, which is recovered with reasonable accuracy.

We propagate the noise in the measurement of $b_i^{(s)}$ into the inferred redshift probability distributions by empirically determining that the scatter in the bias measurements across mocks, $\sigma(b) \sim 0.1$, produces a scatter in the probability distribution $\sigma(P) \sim 0.002$.

6.6 Reconstructed mock redshift distributions

We now consider the results for the recovered redshift distributions of the mock source catalogues, measured in 16 step-wise redshift bins of width $\Delta z = 0.05$. We present a random subset of these results for nine mocks in Fig. 6. The green data points display the best-fitting $P_i^{(s)}$ values and the 1σ errors are derived from the Fisher matrix. The black line shows the Gaussian redshift distribution assumed in the simulations. We can see that the reconstruction has proved generally successful.

For applications to weak gravitational lensing, a flexible functional form for the redshift distribution is more convenient than a step-wise binning. We adopt Gaussian processes (GPs) as a method to infer such functional forms for redshift distributions. Briefly, GPs provide non-parametric Bayesian modelling for supervised learning problems. For details, we refer the reader to Seikel & Clarkson (2013). GPs also readily allow for the inclusion of a prior on the

smoothness of the reconstructed function, naturally expected from a survey selection function.[12]

We build GP models using the python module SCIKIT-learn (Pedregosa et al. 2011). As input one needs to define the functional form for the adopted (redshift) correlation function $C(z, z')$ and set the characteristic scale L, which intuitively determines the typical distance between peaks (i.e. the smoothness scale) of the function. We set the redshift correlation as a Gaussian function

$$C(z, z') = \sigma_{\mathrm{f}}^2 \exp\left[-\frac{(z - z')^2}{2L^2} \right], \qquad (30)$$

where σ_{f}^2 defines the variance of the function, and we set $L = 1$. The optimal choice of L and C will depend on the survey in question.

In Fig. 6, we display the 2σ confidence intervals for the GP models as the green band. We observe that the reconstructed GP distributions encompass the input mock redshift distribution for the vast majority of mocks and redshift bins. More quantitatively, using all 20 mock catalogues we measure an average χ^2 of 19.95 for 18

---

[12] We are not implying here that the final $P(z)$ will necessarily be a very smooth function of redshift, as photo-z cuts can cause sharp variations. Rather, we are implying that the amplitude between redshift bins will be highly correlated. Thus, this setup still allows sharp fluctuations, although they are less probable.

Figure 7. The average of the reconstructed redshift probability distributions over 20 mock catalogues. The blue points show the estimated redshift distribution $\hat{P}$ when we set the bias of the photo-z sample to the mean value of the spec-z bias values. The green points show the estimated $\hat{P}$ when we set $b^{(p)}(z) = 1.0 + \alpha(z - z_0)$, where $\alpha = 2.2$ and $z_0 = 0.45$. The black line shows the input redshift probability distribution, which is recovered more accurately by the second method.

degrees of freedom (ignoring bin-to-bin correlations) demonstrating that, at the level of statistical error of individual mock catalogues (and hence of the observational data sets used in our study), our measurements are statistically consistent with the input distribution.

We now consider a more accurate test of our methodology, using the average $P^{(p)}(z)$ values over all 20 mock catalogues. For this test, we consider two methods to estimate the bias of the photo-z sample, the results of which are shown in Fig. 7. For method 1, we adopt our default model, i.e. $b_i^{(p)} = \langle b_i^{(s)} \rangle$, where we average the spec-z bias factors over 20 mocks (blue points and band). For method 2, we use a bias evolution model with two free parameters defined as $b^{(p)}(z) = 1.0 + \alpha(z - z_0)$, where we have fitted $\alpha = 2.2$ and $z_0 = 0.45$ from the mocks (green points and band). The input Gaussian redshift distribution is displayed as the black line.

Inspecting Fig. 7, we find a significant discrepancy between our predictions based on method 1 and the input distribution. Qualitatively, we observe that for $z < 0.5$ the distribution tends to be underestimated, while for $z > 0.5$ the distribution tends to be overestimated. We interpret these discrepancies as a limitation of our modelling of the CC function and the influence of non-linear galaxy bias. We find that the predictions using the second bias model are more representative.

7 APPLICATION TO KiDS AND 2dFLeNS

In this section, we apply the quadratic estimation methodology to infer the product $b^{(p)}(z) P^{(p)}(z)$ for sources detected in the separate r-band and i-band catalogues of the KiDS imaging survey, using their CC with the 2dFLeNS spec-z catalogue. Because (unlike for the mock catalogues) we have no information on $b^{(p)}(z)$, we cannot break the degeneracy $b^{(p)}(z) P^{(p)}(z)$ without further assumptions. We cross-correlate the photometric samples with the 2dFLeNS catalogue in 18 redshift bins in the range $0 < z < 0.9$. The r-band and i-band samples have a different degree of overlap with 2dFLeNS, and we use a total of 13 740 and 25 443 spec-z galaxies for the respective CCs. Following the cosmic shear analysis of Hildebrandt et al. (2016), we divided the KiDS-450 r-band imaging data set into four tomographic bins based on the bpz redshift $z_{\rm B}$: $0.1 < z_{\rm B} \leq 0.3$, $0.3 < z_{\rm B} \leq 0.5$, $0.5 < z_{\rm B} \leq 0.7$ and $0.7 < z_{\rm B} \leq 0.9$. Photo-z information is not available for the KiDS-800 i-band data set, and we do not sub-divide it. In the correlation function measurement, we now weight each source by its optimal weight in the estimation of

shear statistics, such that we are constraining the weighted redshift probability distribution of the sources (Miller et al. 2013).

Since we are only utilizing spectroscopic data in the range $z < 0.9$, we cannot derive the full KiDS source redshift distributions for most of the samples. Rather, our motivation is to demonstrate an application of our methodology on a realistic data set.

We fit for the spec-z bias values in each redshift bin as outlined in the previous section, and show the results in Fig. 8. We perform a rough scaling of the expected error in the bias compared to the mocks, and propagate this error into the determination of $b^{(p)} P^{(p)}$ by adding a term $\sigma(P) = 0.02$ to the standard Fisher matrix errors.

It is possible for measurements of the angular CC function to be negative due to either noise, or effects such as incompleteness in the imaging catalogue around bright spectroscopic sources or other systematics. These points are unphysical in our model, which predicts the CC functions by scaling the autocorrelations, which are positive definite. To address this problem we add a positive definite prior which effectively shifts negative $b^{(p)} P^{(p)}$ values to zero.

In Fig. 9, we show the reconstructed $b^{(p)}(z) \times P^{(p)}(z)$ measurements for each tomographic bin of the KiDS-450 r-band data, in comparison with other determinations of this distribution presented by Hildebrandt et al. (2016). The green data points show the quadratic estimation with 1σ error bars and the shaded green band gives the 68 per cent confidence interval for the GP model. Only 14 redshift bins are shown per tomographic bin; it was not possible to recover an estimate for the first two and last two redshift bins because of the low number of spec-z galaxies in these bins, which caused instabilities in the estimator, hence poor convergence. The red and blue bands display the distributions and 68 per cent confidence ranges obtained by applying two other methods:

(i) Calibration with small-scale CCs, shown by the red bands and implemented by applying the recipes of Newman (2008) and Matthews & Newman (2010) to the cross-clustering of KiDS sources and deep spectroscopic samples from DEEP2 and COSMOS.

(ii) Weighted direct calibration (DIR), shown by the blue bands and based on direct determination of the source redshift distribution by cross-matching with a number of overlapping spectroscopic samples, with appropriate re-weighting for incompleteness.

We refer the reader to Hildebrandt et al. (2016) for more details about these methods. We converted the DIR estimation from $P^{(p)}(z)$

Figure 8. The bias evolution of the spectroscopic 2dFLenS sample for the overlaps with the KiDS *r*-band and *i*-band samples, measured from the galaxy autocorrelations. The errors are determined by scaling the results from mock catalogues.


Figure 9. Reconstruction of the combination $b^{(p)}P^{(p)}(z)$ by applying quadratic estimation to CCs between the KiDS-450 *r*-band catalogue and 2dFLenS, for four tomographic bins of the photometric catalogue ($0.1 < z_B \le 0.3$, $0.3 < z_B \le 0.5$, $0.5 < z_B \le 0.7$ and $0.7 < z_B \le 0.9$). The green data points display the results of quadratic estimation and the green bands show the 68 per cent confidence intervals for a Gaussian process model trained using these measurements. These results are compared with determinations by methods using small-scale CC (red bands) and weighted direct calibration (blue bands) (see Hildebrandt et al. 2016 for more details about these methods). For the purposes of comparison, all the distributions have been normalized such that $\int_{0.1}^{0.8} \mathrm{d}z\, b^{(p)}\, P^{(p)} = 1$.

to $b^{(p)}(z) \times P^{(p)}(z)$ using the function $b^{(p)}(z)$ implicitly assumed in the CC method, as outlined by Matthews & Newman (2010). For the purposes of this comparison, all the distributions have been normalized such that $\int_{0.1}^{0.8} \mathrm{d}z\, b^{(p)}\, P^{(p)} = 1$. These different methods produce redshift distributions that agree in a qualitative sense, although comparisons illustrate systematic errors affecting each technique.

The equivalent quadratic estimation for the KiDS-800 *i*-band imaging data, analysed in a single tomographic bin, is shown in

Fig. 10. In this figure, the blue points display the reconstructed redshift distribution with 1σ errors and the shaded blue region shows the 95 per cent confidence interval for the GP model. We observe that the amplitude of the CC signal increases with redshift, driven by a combination of the underlying redshift distribution $P^{(p)}(z)$ and the source galaxy bias $b^{(p)}(z)$. Assuming that the latter is a slowly varying function, this analysis suggests that the redshift distribution of the *i*-band sources is broad and peaked at $z \gtrsim 0.7$. These findings are qualitatively consistent with the DIR estimate presented by

Figure 10. Reconstruction of the combination $b^{(p)}P^{(p)}(z)$ by applying quadratic estimation to CCs between the KiDS-800 *i*-band catalogue and 2dFLenS. The blue points display the results of quadratic estimation and the blue band shows the 95 per cent confidence intervals for a Gaussian process model trained using these measurements.

Amon et al. (in preparation), although further comparison is beyond the scope of the current study.

These results demonstrate that, although the 2dFLenS data set does not have sufficient redshift coverage to derive the full source distributions by CC, we can successfully apply our methodology to the KiDS data set and obtain results in qualitative agreement with previous determinations.

8 CONCLUSION

With the issue of source redshift calibration becoming increasingly pronounced for weak gravitational lensing surveys, new and versatile approaches to this problem are needed. Calibration via CC with overlapping spectroscopic surveys represents one such approach. In this work, we have presented our efforts to extend the accuracy and applicability of such methods to both simulations and data. We summarize our main findings as follows:

(i) We have developed a new, minimum-variance and unbiased quadratic estimator that infers the redshift probability distributions of photometric samples of galaxies $P^{(p)}(z)$, in a degenerate combination with their galaxy bias $b^{(p)}(z)$, based on their angular CC with an overlapping spectroscopic sample. This derivation expands on work presented by McQuinn & White (2013).

(ii) We have tested our methodology on a series of mock galaxy catalogues. We found that at the level of statistical errors of current surveys the estimator is unbiased. However, if we stack mocks together to perform a more accurate test, we discover small but significant discrepancies – we attribute these effects to the breakdown of our clustering model due, for example, to non-linear galaxy bias.

(iii) We derive non-parametric, continuous functional forms of $b^{(p)}P^{(p)}(z)$ by building GP models from the step-wise constraints inferred from the quadratic estimator. Such continuous functions are useful for modelling the lensing signal, and allow the computation of continuous confidence intervals.

(iv) We have applied our methodology to infer $b^{(p)}P^{(p)}(z)$ functions for KiDS *r*-band and *i*-band imaging catalogues in the range $0.1 < z < 0.8$, through CC with the 2dFLenS spectroscopic redshift survey. Our distributions are in qualitative agreement with the results of other methods.

Our analysis could be extended in a number of ways: by improving the modelling of non-linear effects, enhancing our mock

catalogues to match our data sets more closely and re-formulating the estimator to measure the spectroscopic galaxy bias and redshift distribution simultaneously. However, the calibration of the bias of the photometric sample remains the most critical component. Possible approaches to this problem include the use of redshift-space distortions, lensing magnification and galaxy–galaxy lensing. We hope that our work motivates more research on these topics.

REFERENCES

Abbott T. et al., 2016, Phys. Rev. D, 94, 022001

Alsing J., Heavens A., Jaffe A. H., Kiessling A., Wandelt B., Hoffmann T., 2016, MNRAS, 455, 4452

Baldauf T., Seljak U., Smith R. E., Hamaus N., Desjacques V., 2013, Phys. Rev. D, 88, 083507

Benítez N., 2000, ApJ, 536, 571

Benjamin J. et al., 2013, MNRAS, 431, 1547

Bernstein G., Huterer D., 2010, MNRAS, 401, 1399

Blake C. et al., 2016, MNRAS, 462, 4240

Bond J. R., Jaffe A. H., Knox L., 1998, Phys. Rev. D, 57, 2117

Brammer G. B., van Dokkum P. G., Coppi P., 2008, ApJ, 686, 1503

Brouwer M. M. et al., 2016, MNRAS, 462, 4451

Cai Y.-C., Bernstein G., 2012, MNRAS, 422, 1045

Carrasco Kind M., Brunner R. J., 2013, MNRAS, 432, 1483

Cavuoti S. et al., 2015, MNRAS, 452, 3100

Choi A. et al., 2016, MNRAS, 463, 3737

Cunha C. E., Huterer D., Busha M. T., Wechsler R. H., 2012, MNRAS, 423, 909

Cunha C. E., Huterer D., Lin H., Busha M. T., Wechsler R. H., 2014, MNRAS, 444, 129

Dawson K. S. et al., 2013, AJ, 145, 10

de Jong J. T. A. et al., 2015, A&A, 582, A62

Dodelson S., 2003, Modern Cosmology, Academic Press, New York

Driver S. P. et al., 2011, MNRAS, 413, 971

Duncan C. A. J., Joachimi B., Heavens A. F., Heymans C., Hildebrandt H., 2014, MNRAS, 437, 2471

Erben T. et al., 2009, A&A, 493, 1197

Gaztañaga E., Eriksen M., Crocce M., Castander F. J., Fosalba P., Marti P., Miquel R., Cabré A., 2012, MNRAS, 422, 2904

Graff P., Feroz F., Hobson M. P., Lasenby A., 2014, MNRAS, 441, 1741

Harnois-Déraps J., van Waerbeke L., 2015, MNRAS, 450, 2857

Harnois-Déraps J., Pen U.-L., Iliev I. T., Merz H., Emberson J. D., Desjacques V., 2013, MNRAS, 436, 540

Hearin A. P., Zentner A. R., Ma Z., Huterer D., 2010, ApJ, 720, 1351

Hildebrandt H. et al., 2016, MNRAS, preprint (arXiv:1606.05338)

Ho S., Hirata C., Padmanabhan N., Seljak U., Bahcall N., 2008, Phys. Rev. D, 78, 043519

Hu W., 1999, ApJ, 522, L21

Huterer D., 2002, Phys. Rev. D, 65, 063001

Huterer D., Takada M., Bernstein G., Jain B., 2006, MNRAS, 366, 101

Kitching T. D., Taylor A. N., Heavens A. F., 2008, MNRAS, 389, 173

Köhlinger F., Viola M., Valkenburg W., Joachimi B., Hoekstra H., Kuijken K., 2016, MNRAS, 456, 1508

Kuijken K. et al., 2015, MNRAS, 454, 3500

Landy S. D., Szalay A. S., 1993, ApJ, 412, 64

Lima M., Cunha C. E., Oyaizu H., Frieman J., Lin H., Sheldon E. S., 2008, MNRAS, 390, 118

Limber D. N., 1954, ApJ, 119, 655

Ma Z., Hu W., Huterer D., 2006, ApJ, 636, 21

Matthews D. J., Newman J. A., 2010, ApJ, 721, 456

McQuinn M., White M., 2013, MNRAS, 433, 2857 (MW13)

Ménard B., Scranton R., Schmidt S., Morrison C., Jeong D., Budavari T., Rahman M., 2013, preprint (arXiv:1303.4722)

Miller L. et al., 2013, MNRAS, 429, 2858

Morrison C. B., Schneider M. D., 2013, J. Cosmol. Astropart. Phys., 11, 9

Morrison C. B., Hildebrandt H., Schmidt S. J., Baldry I. K., Bilicki M., Choi A., Erben T., Schneider P., 2016, MNRAS, preprint (arXiv:1609.09085)

Newman J. A., 2008, ApJ, 684, 88

Newman J. A. et al., 2013, ApJS, 208, 5

Newman J. A. et al., 2015, Astropart. Phys., 63, 81

Padmanabhan N. et al., 2007, MNRAS, 378, 852

Pedregosa F. et al., 2011, J. Mach. Learn. Res., 12, 2825

Rahman M., Ménard B., Scranton R., Schmidt S. J., Morrison C. B., 2015, MNRAS, 447, 3500

Rahman M., Ménard B., Scranton R., 2016, MNRAS, 457, 3912

Sadeh I., Abdalla F. B., Lahav O., 2016, PASP, 128, 104502

Sánchez C. et al., 2014, MNRAS, 445, 1482

Schmidt S. J., Ménard B., Scranton R., Morrison C., McBride C. K., 2013, MNRAS, 431, 3307

Schmidt S. J., Ménard B., Scranton R., Morrison C. B., Rahman M., Hopkins A. M., 2015, MNRAS, 446, 2696

Schulz A. E., 2010, ApJ, 724, 1305

Scottez V. et al., 2016, MNRAS, 462, 1683

Seikel M., Clarkson C., 2013, preprint (arXiv:1311.6678)

Shanks T. et al., 2015, MNRAS, 451, 4238

Sifón C. et al., 2015, MNRAS, 454, 3938

Smith R. E. et al., 2003, MNRAS, 341, 1311

Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, ApJ, 761, 152

Tegmark M., 1997, Phys. Rev. D, 55, 5895

van Uitert E. et al., 2016, MNRAS, 459, 3251

Viola M. et al., 2015, MNRAS, 452, 3529

Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2013, Phys. Rep., 530, 87

Wolf C. et al., 2016, MNRAS, in press

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