

The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: analysis of potential systematics

Ashley J. Ross,^{1*} Will J. Percival,¹ Ariel G. Sánchez,² Lado Samushia,^{1,3} Shirley Ho,^{4,5} Eyal Kazin,⁶ Marc Manera,¹ Beth Reid,^{4,†} Martin White,^{4,7,8} Rita Tojeiro,¹ Cameron K. McBride,⁹ Xiaoying Xu,¹⁰ David A. Wake,¹¹ Michael A. Strauss,¹² Francesco Montesano,² Molly E. C. Swanson,⁹ Stephen Bailey,⁴ Adam S. Bolton,¹³ Antonio Montero Dorta,¹⁴ Daniel J. Eisenstein,⁹ Hong Guo,¹⁵ Jean-Christophe Hamilton,¹⁶ Robert C. Nichol,¹ Nikhil Padmanabhan,¹⁷ Francisco Prada,^{14,18,19} David J. Schlegel,⁴ Mariana Vargas Magaña,¹⁶ Idit Zehavi,¹⁵ Michael Blanton,²⁰ Dmitry Bizyaev,²¹ Howard Brewington,²¹ Antonio J. Cuesta,¹⁷ Elena Malanushenko,²¹ Viktor Malanushenko,¹⁹ Daniel Oravetz,²¹ John Parejko,¹⁷ Kaike Pan,²¹ Donald P. Schneider,^{22,23} Alaina Shelden,²¹ Audrey Simmons,²¹ Stephanie Snedden²¹ and Gong-bo Zhao¹

¹*Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth PO1 3FX*

²*Max-Planck-Institut für extraterrestrische Physik, Postfach 1312, Giessenbachstr., 85748 Garching, Germany*

³*National Abastumani Astrophysical Observatory, Ilia State University, 2A Kazbegi Avenue, GE-1060 Tbilisi, Georgia*

⁴*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA*

⁵*Department of Physics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

⁶*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, Victoria 3122, Australia*

⁷*Department of Physics, University of California, 366 LeConte Hall, Berkeley, CA 94720, USA*

⁸*Department of Astronomy, University of California at Berkeley, 601 Campbell Hall, Berkeley, CA 94720, USA*

⁹*Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, MS #20, Cambridge, MA 02138, USA*

¹⁰*Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721, USA*

¹¹*Yale Center for Astronomy and Astrophysics, Yale University, New Haven, CT 06511, USA*

¹²*Department of Astrophysical Sciences, Princeton University, Ivy Lane, Princeton, NJ 08544, USA*

¹³*Department of Physics and Astronomy, University of Utah, UT 84112, USA*

¹⁴*Instituto de Astrofísica de Andalucía (CSIC), E-18080 Granada, Spain*

¹⁵*Department of Astronomy, Case Western Reserve University, Cleveland, OH 44106, USA*

¹⁶*APC, Université Paris-Diderot-Paris 7, CNRS/IN2P3, CEA, Observatoire de Paris, 10, rue A. Domon and L. Duquet, Paris, France*

¹⁷*Department of Physics, Yale University, 260 Whitney Ave, New Haven, CT 06520, USA*

¹⁸*Campus of International Excellence UAM+CSIC, Cantoblanco, E-28049 Madrid, Spain*

¹⁹*Instituto de Física Teórica (UAM/CSIC), Universidad Autónoma de Madrid, Cantoblanco, E-28049 Madrid, Spain*

²⁰*Center for Cosmology and Particle Physics, New York University, New York, NY 10003, USA*

²¹*Apache Point Observatory, PO Box 59, Sunspot, NM 88349-0059, USA*

²²*Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA*

²³*Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA*

Accepted 2012 May 3. Received 2012 April 28; in original form 2012 March 29

ABSTRACT

We analyse the density field of galaxies observed by the Sloan Digital Sky Survey (SDSS)-III Baryon Oscillation Spectroscopic Survey (BOSS) included in the SDSS Data Release Nine (DR9). DR9 includes spectroscopic redshifts for over 400 000 galaxies spread over a footprint of 3275 deg². We identify, characterize and mitigate the impact of sources of

*E-mail: Ashley.Ross@port.ac.uk

†Hubble fellow.

systematic uncertainty on large-scale clustering measurements, both for angular moments of the redshift-space correlation function, $\xi_\ell(s)$, and the spherically averaged power spectrum, $P(k)$, in order to ensure that robust cosmological constraints will be obtained from these data. A correlation between the projected density of stars and the higher redshift ($0.43 < z < 0.7$) galaxy sample (the approximately constant stellar mass threshold ‘CMASS’ sample) due to imaging systematics imparts a systematic error that is larger than the statistical error of the clustering measurements at scales $s > 120 h^{-1} \text{Mpc}$ or $k < 0.01 h \text{Mpc}^{-1}$. We find that these errors can be ameliorated by weighting galaxies based on their surface brightness and the local stellar density. The clustering of CMASS galaxies found in the Northern and Southern Galactic footprints of the survey generally agrees to within 2σ . We use mock galaxy catalogues that simulate the CMASS selection function to determine that randomly selecting galaxy redshifts in order to simulate the radial selection function of a random sample imparts the least systematic error on $\xi_\ell(s)$ measurements and that this systematic error is negligible for the spherically averaged correlation function, ξ_0 . We find a peak in ξ_0 at $s \sim 200 h^{-1} \text{Mpc}$, with a corresponding feature with period $\sim 0.03 h \text{Mpc}^{-1}$ in $P(k)$, and find features at least as strong in 4.8 per cent of the mock galaxy catalogues, concluding this feature is likely to be a consequence of cosmic variance. The methods we recommend for the calculation of clustering measurements using the CMASS sample are adopted in companion papers that locate the position of the baryon acoustic oscillation feature, constrain cosmological models using the full shape of ξ_0 and measure the rate of structure growth.

Key words: cosmology: observations – distance scale – large-scale structure of Universe.

1 INTRODUCTION

In the last decade, wide-field surveys such as the Two Degree Field Galaxy Redshift Survey (2dFGRS; Colless et al. 2003), the Sloan Digital Sky Survey (SDSS; York et al. 2000) and the WiggleZ Redshift Survey (Blake et al. 2010) have obtained accurate spectroscopic redshifts of well over one million galaxies, allowing maps of the three-dimensional structure of the Universe to be constructed out to $z = 0.9$. These maps encode a wealth of information on cosmology (e.g. Tegmark et al. 2004; Cole et al. 2005; Eisenstein et al. 2005; Percival et al. 2010; Reid et al. 2010; Blake et al. 2011; Montesano, Sanchez & Phleps 2011) and the nature of galaxies (e.g. Norberg et al. 2002; Gómez et al. 2003; Swanson et al. 2008a; Wake et al. 2008; Tojeiro & Percival 2010; Ross, Tojeiro & Percival 2011a; Zehavi et al. 2011).

The Baryon Oscillation Spectroscopic Survey (BOSS) is designed to obtain spectroscopy of 1.5 million galaxies to $z = 0.7$ over an imaging area of $10\,000 \text{ deg}^2$ (Eisenstein et al. 2011). White et al. (2011) investigated an early sample from this survey, confirming the survey was making a high-quality map of massive galaxies with bias ~ 2 . We utilize spectroscopic redshifts for over 400 000 BOSS galaxies that will be released as part of the SDSS Data Release Nine (DR9). These galaxies cover close to one-third of the final (planned) footprint, and currently comprise the largest effective volume (Tegmark & Peebles 1998) of any spectroscopic galaxy catalogue – 2.2 Gpc^3 [assuming a concordance Λ cold dark matter (Λ CDM) model]. These data should therefore allow the best to-date statistical uncertainty on the measured power spectrum, $P(k)$, and thus the best to-date cosmological measurements determined using a galaxy catalogue. As such, discovery and elimination of systematic uncertainty is of vital importance to realizing the survey goals. Potential systematic effects on the observed density of galaxies must be robustly tested and ameliorated in an unbiased way.

The purpose of this study is to identify and minimize the impact of sources of systematic uncertainty in the large-scale clustering of

BOSS galaxies, in order to ensure that robust cosmological constraints are obtained. Ross et al. (2011b) studied systematic effects on the projected density of BOSS galaxy targets, finding a strong relationship with stellar density and differences in the sample in occupying the Northern and Southern Galactic Caps (hereafter NGC and SGC). We follow up and extend this work by testing how these systematic variations effect spatial clustering measurements and by testing against systematic effects incurred when obtaining spectroscopic redshifts. We aim to answer the following questions:

- (i) How do variations in photometric calibration, e.g. between the NGC and SGC footprints, affect the selection of BOSS galaxies?
- (ii) How does the observed density of galaxies depend on observing conditions?
- (iii) What is the best way to simulate the radial selection function and how important are effects related to galaxy evolution?
- (iv) How do permutations of (i)–(iii) affect the clustering we measure?

Our results have already been used in studies of the clustering of BOSS DR9 galaxies. Anderson et al. (2012) localize the position of the baryon acoustic oscillation (BAO) feature to better than 2 per cent accuracy. Reid et al. (2012) measure redshift-space distortions (RSD) and Samushia et al. (in preparation) thereby constrain dark energy and modified gravity models. See also Tojeiro et al. (2012a) for a complementary method of measuring structure growth using DR9 galaxies. Nuza et al. (2012) found that the clustering of BOSS galaxies can be well approximated by using a subhalo abundance matching method applied to a dark matter simulation. Sanchez et al. (2012) obtain cosmological constraints by fitting the full shape of the correlation function. We hope that future studies heed, and improve upon, our analysis, which we feel is the most careful analysis of observational systematics to date.

The presentation of our analysis is organized as follows. In Section 2 we describe the BOSS DR9 sample of galaxies and its corresponding angular mask. In Section 3, we describe how we estimate

clustering statistics, their covariance, and compare to models. For both the covariance and the models we utilize the mock catalogues of galaxies (hereafter ‘mocks’) generated by Manera et al. (2012). In Section 4, we investigate and explain the differences we find in the densities of galaxies in the NGC and SGC, addressing question (i). In Section 5 we describe potential sources of systematic variation in the density of galaxies targeted for spectroscopy and the methods we employ to remove these variations in an unbiased way, addressing question (ii). In Section 6, we investigate the radial distribution of our galaxy sample, using mock catalogues to determine the least biased way to simulate the radial selection function of BOSS galaxies and to check that the clustering we measure is robust when the galaxies are split into two samples above/below redshift 0.52, thus addressing question (iii). Throughout Sections 4–6, we address question (iv) using $\xi_i(s)$ measurements at $s < 150 h^{-1}$ Mpc. In Section 7, we consider the clustering at scales $s > 150 h^{-1}$ Mpc, also utilizing measurements of anisotropic clustering and the power spectrum. We conclude in Section 8. Throughout, we assume a flat cosmology with $\Omega_m = 0.274$, $\Omega_b h^2 = 0.0224$, $h = 0.70$, $n_s = 0.95$ and $\sigma_8 = 0.8$ (identical to that used in White et al. 2011 and Anderson et al. 2012) unless otherwise noted.

2 DATA

The SDSS-III BOSS (Eisenstein et al. 2011) obtains targets using SDSS imaging data. In combination, the SDSS-I, SDSS-II and SDSS-III surveys obtained wide-field CCD photometry (Gunn et al. 1998, 2006) in five passbands (u, g, r, i, z ; e.g. Fukugita et al. 1996), amassing a total footprint of $14\,555 \text{ deg}^2$, internally calibrated using the ‘uber-calibration’ process described in Padmanabhan et al. (2008), and with a 50 per cent completeness limit of point sources at $r = 22.5$ (Aihara et al. 2011). After completing the imaging, BOSS has targeted 1.5 million massive galaxies, 150 000 quasars and over 75 000 ancillary targets for spectroscopic observations over an area of $10\,000 \text{ deg}^2$ (Eisenstein et al. 2011). BOSS observations began in fall 2009, and the last data will be acquired in 2014. The BOSS spectrographs ($R = 1300\text{--}3000$) are fed by 1000 optical fibres in a single pointing, each with a 2-arcsec aperture. Each observation is performed in a series of 15-min exposures and integrated until a fiducial minimum signal-to-noise ratio, chosen to ensure a high-redshift success rate, is reached. This ensures a sample with nearly isotropic redshift selection complete to 98 per cent. We test this isotropy in Section 2.3.

2.1 Target selection

BOSS targets two samples of galaxies. These are the ‘LOWZ’ (the lower redshift sample) and ‘CMASS’ (the higher redshift sample with an approximately constant stellar threshold) samples, as described by Eisenstein et al. (2011). We are careful throughout this paper to distinguish between target objects and true galaxies – the majority of LOWZ and CMASS targets are galaxies, but a small fraction are stars (3 per cent of CMASS) and high-redshift quasars (1 per cent, i.e. not objects sampling the intended density field). Anything we refer to as a galaxy has been spectroscopically confirmed as such. The SDSS measures magnitudes using both point spread function (PSF) convolved fits to DeVaucouleurs (we denote these with the subscript ‘dev’) and exponential profiles (we denote these with the subscript ‘exp’). Each of these magnitudes is used to determine ‘model’, which we denote by the subscript ‘mod’, and ‘cmodel’ magnitudes, denoted by the subscript ‘cmod’, which are used in target selection. Model magnitudes denote the best fit of the two profiles in the r band (see Stoughton et al. 2002 for further

information on model magnitudes). The cmodel magnitudes, first defined in Abazajian et al. (2004), represent the best-fitting linear combination of the exponential and DeVaucouleurs model fluxes. We will also use PSF magnitudes, which we denote by the subscript ‘psf’.

We select BOSS targets using the photometry of objects identified as galaxies by the SDSS pipeline. Most of the sample (100 per cent for LOWZ and 90.9 per cent for CMASS) was targeted using the SDSS DR8 photometry designated as ‘primary’. The remaining sample was targeted from images now designated in DR8 as secondary. These data were superseded by overlapping imaging runs of better quality but whose reductions were unavailable at the time of targeting. Photometric scatter across the multiple selection boundaries listed below implies that many objects targeted using primary photometry would not have been targeted using secondary photometry (and vice versa). However, this result is simply due to the known statistical distribution of measured magnitudes, quantified by the magnitude error. This effect should not cause any additional systematic error beyond that potentially induced by targeting from a sample with magnitude errors that vary with angular position, assuming one always uses the photometry used *at the time of targeting* in an analysis. Indeed, we find restricting our analyses to data targeted using DR8 primary photometry results in no significant change in any clustering statistic we measure.

Eisenstein et al. (2011) define the selection criteria for BOSS galaxy targets. We repeat them here for completeness and ease of reference. The CMASS selection is defined by¹

$$17.5 < i_{\text{cmod}} < 19.9, \quad (1)$$

$$r_{\text{mod}} - i_{\text{mod}} < 2, \quad (2)$$

$$d_{\perp} > 0.55, \quad (3)$$

$$i_{\text{fib2}} < 21.5, \quad (4)$$

$$i_{\text{cmod}} < 19.86 + 1.6(d_{\perp} - 0.8), \quad (5)$$

where all magnitudes are corrected for Galactic extinction (via the Schlegel, Finkbeiner & Davis 1998 dust maps), i_{fib2} is the i -band magnitude within a 2-arcsec aperture, and

$$d_{\perp} = r_{\text{mod}} - i_{\text{mod}} - (g_{\text{mod}} - r_{\text{mod}})/8.0. \quad (6)$$

These colour cuts are designed to obtain a sample of galaxies with approximately constant stellar mass with $z \gtrsim 0.43$ and include many galaxies that would be considered ‘blue’ by traditional SDSS (rest-frame) colour cuts (see e.g. Strateva et al. 2001). Indeed, Masters et al. (2011) find that 26 per cent of CMASS galaxies have a late-type (i.e. spiral disc) morphology. See Tojeiro et al. (2012b) for a detailed description of the CMASS population of galaxies.

For CMASS targets, stars are further separated from galaxies by only keeping objects with

$$i_{\text{psf}} - i_{\text{mod}} > 0.2 + 0.2(20.0 - i_{\text{mod}}), \quad (7)$$

$$z_{\text{psf}} - z_{\text{mod}} > 9.125 - 0.46z_{\text{mod}}, \quad (8)$$

¹ In the early part of the survey, various super-sets of this selection were used, e.g. the fibre magnitude limit has changed from $i_{\text{fib2}} < 21.7$ to $i_{\text{fib2}} < 21.5$. We only use data satisfying the above selection cuts in our analysis and recommend the same for any cosmological analysis using BOSS galaxy data, as this provides a more isotropic selection and discards less than 3 per cent of the total available DR9 CMASS redshifts.




Figure 1. The footprint of BOSS DR9 galaxies, projected into two dimensions using the McBryde-Thomas flat polar quartic projection, is shaded in blue and red. Areas with CMASS data only are shaded blue. The CMASS and LOWZ footprints cover 3275 and 2208 deg², respectively. The grey area represents the final (planned) BOSS footprint.

unless the object also passes the LOWZ cuts (only 0.5 per cent of objects passing the LOWZ selection cuts are stars), which are defined by

$$r_{\text{cmod}} < 13.5 + c_{\parallel}/0.3, \quad (9)$$

$$|c_{\perp}| < 0.2, \quad (10)$$

$$16 < r_{\text{cmod}} < 19.6, \quad (11)$$

$$r_{\text{psf}} - r_{\text{mod}} > 0.3, \quad (12)$$

where

$$c_{\parallel} = 0.7(g_{\text{mod}} - r_{\text{mod}}) + 1.2(r_{\text{mod}} - i_{\text{mod}} - 0.18), \quad (13)$$

and

$$c_{\perp} = r_{\text{mod}} - i_{\text{mod}} - (g_{\text{mod}} - r_{\text{mod}})/4.0 - 0.18. \quad (14)$$

Some objects satisfy both the LOWZ and CMASS selection criteria. We therefore apply a minimum (maximum) redshift of 0.43 to the CMASS (LOWZ) sample, after obtaining a redshift in order to have two mutually exclusive samples.

The earliest set of spectra obtained for LOWZ data used an overly restrictive algorithm designed to remove stellar contamination, which unfortunately removed a significant number of galaxies from the target sample. This algorithm was changed for later data and, to maximize the size of the sample with an isotropic selection algorithm, we reduce the area by excluding the regions observed with the restrictive algorithm. Thus, the coverage (after accounting for completeness, see Section 2.2) of the LOWZ sample we use (2208 deg²) is smaller than that of the CMASS sample (3275 deg²). Fig. 1 displays the angular footprint of the LOWZ sample in red and the area that contains only CMASS data in blue. The footprint contains 327 349 CMASS targets and 132 060 LOWZ targets. All of the data in these catalogues will be publicly released in the SDSS DR9.

2.2 Mask

The BOSS DR9 geometry is constructed from a series of spectroscopic observations, each of which is a 3° diameter circle on the sky, corresponding to one pointing of the telescope. Each of these

circles² contains a unique set of targets and its area represents a ‘tile’ (see Blanton et al. 2003; Dawson et al., in preparation). The total area covered by these tiles forms the basis of our angular mask. Some tiles are not fully covered by observations, usually due to a lack of imaging data in the targeted region, so these are additional boundaries that we include in the mask. We divide the total area into ‘sectors’, which are defined as the areas covered by a unique set of tiles – i.e. the regions where the spectroscopic observing conditions are the same. We use the software package MANGLE³ (see Hamilton 1993; Hamilton & Tegmark 2004; Swanson et al. 2008b) to use the tile positions to divide the area into the unique sectors that we use to define the mask.

We also apply veto masks that exclude areas surrounding bright stars and imaging fields deemed not photometric, both of these processes are described in Anderson et al. (2012). In these areas, we do not expect the observing conditions to allow uniform detection of BOSS galaxies. Additionally, we mask the 92-arcsec diameter region at the centre of every tile where fibres cannot be placed due to physical limitations. The veto mask is only applied to the data after targeting; that is, galaxies observed in non-photometric fields and near bright stars are excluded from our analysis. This provides a cleaner sample, albeit with a more complicated mask, than if we were to quantify the varying target selection due to these effects. These masks remove 4 per cent of the observed footprint.

The density of galaxy targets on the sky varies due to galaxy clustering. Therefore, given that there are a finite number of fibres for each tile, the percentage of targets receiving a fibre will vary. Additionally, fibres cannot be placed closer than 62 arcsec due to the size of the cladding around each fibre. We denote a ‘close pair’ as any object not assigned a fibre due to a collision with an object of the same target type (i.e. CMASS target with CMASS target), since collisions with objects of different types should show no spatial correlation.

In each sector, we compile statistics using the same definitions as Anderson et al. (2012). The angular completeness, C_{BOSS} , and the redshift completeness, C_{red} , are determined by first counting the number of objects in each sector that are

- (i) spectroscopically confirmed stars (N_{star}),
- (ii) galaxies with redshifts from good BOSS spectra (N_{gal}),
- (iii) galaxies with redshifts from SDSS-II spectra (N_{known}),
- (iv) objects with BOSS spectra from which stellar classification or redshift determination failed (N_{fail}),
- (v) objects with no spectra, in a close pair (N_{cp}),
- (vi) objects with no spectra, not in a close pair (N_{missed}).

These definitions represent a complete accounting for the possible outcomes of BOSS targets. Objects contributing to N_{missed} will either be observed in the future or are fibre collisions with objects of a different target type. For each sector, Anderson et al. (2012) then define the following:

$$N_{\text{targ}} = N_{\text{star}} + N_{\text{gal}} + N_{\text{fail}} + N_{\text{cp}} + N_{\text{missed}} + N_{\text{known}}, \quad (15)$$

where N_{targ} is the total number of target objects, and

$$N_{\text{obs}} = N_{\text{star}} + N_{\text{gal}} + N_{\text{fail}}, \quad (16)$$

² Each observation corresponds to a ‘plate’. Each set of targets has a unique tile, but multiple plates can observe the same tile (and thus repeat observation of the exact same set of targets).

³ <http://space.mit.edu/molloy/mangle/>




Figure 2. The distribution of CMASS targets for a selection of mask sectors. Circles represent the sky area covered by observing tiles and the number of overlapping tiles is indicated by the level of shading. Black dots indicate the positions of targets for which we obtained a ‘good’ redshift, as defined in Section 2.3. Blue squares denote targets for which we did not allocate a fibre for spectroscopic observations because the target is within 62 arcsec of another CMASS target (‘close pair’). Green circles denote targets for which we did allocate a fibre that are not close pairs. Red triangles denote targets for which we allocated a fibre, but did not obtain a good redshift.

where N_{obs} is the total number of objects within the sector with a BOSS spectrum. C_{BOSS} is then

$$C_{\text{BOSS}} = \frac{N_{\text{obs}} + N_{\text{cp}}}{N_{\text{targ}} - N_{\text{known}}}, \quad (17)$$

and finally C_{red} is

$$C_{\text{red}} = \frac{N_{\text{gal}}}{N_{\text{obs}} - N_{\text{star}}}. \quad (18)$$

The C_{BOSS} completeness varies from sector to sector due to fibre collisions with objects with a different type and the fact that many objects will be observed in future observations. We subsample the known redshift sample (complete by definition) so that its completeness matches C_{BOSS} . We discard from our analysis any sectors where $C_{\text{BOSS}} < 0.7$ or $C_{\text{red}} < 0.8$. These cuts remove 8 per cent of the total footprint covered by BOSS DR9, but only 3.5 per cent of galaxy redshifts. Making the completeness cuts more restrictive does not significantly affect any clustering statistic we measure.

Fig. 2 displays CMASS targets for a selected observed area. Areas covered by more than one tile are shaded such that the outline of each tile is clearly visible. These overlapping regions cover 41 per cent of the total DR9 footprint. Targets with good redshifts are plotted as small black points. Targets not allocated a fibre are green. Within a given sector, these should be random with respect to the position of other CMASS targets, and these are therefore accounted for by C_{BOSS} . (Although they are more likely to occur in sectors where future observations are planned, they are still random *within these sectors*.) Targets not allocated a fibre due to close pair collisions are blue. This happens most frequently (but not exclusively) in regions covered by only one tile. Targets that were allocated a fibre but

whose observation did not result in a good redshift measurement are red. In general, these occur more frequently near the tile boundaries. We discuss these cases further in Section 2.3.

We create random (unclustered) catalogues by isotropically populating the sky, then selecting only those positions lying inside sectors with $C_{\text{BOSS}} > 0.7$ and $C_{\text{red}} > 0.8$ and outside of the veto mask. We then cull the random positions in every sector based on their C_{BOSS} (i.e. if $C_{\text{BOSS}} = 0.9$, we randomly remove 10 per cent of the random points). This process yields a random catalogue that mimics the angular distribution of our galaxy catalogues, save for fibre collisions with galaxies of the same type, redshift completeness and systematic effects in the imaging. We correct for these remaining effects using a series of weights, as described in Section 3.2. Our default approach is to assign to each random position the redshift of randomly selected galaxy and we test this approach in Section 6.1.

2.3 Redshift failures

We define a ‘good’ redshift as any galaxy that does not have any ‘zWarning’ flags (as defined in Adelman-McCarthy et al. 2008) determined by the spectroscopic pipeline. This flag indicates that the redshifts are unreliable typically because there are multiple acceptable redshift solutions (usually due to low signal-to-noise ratio) or that the spectrum is defective. An analysis of repeat observations of BOSS targets and visual inspection reveals that galaxies with $\text{zWarning} = 0$ are reliable [accurate to < 0.001 in $\Delta z/(1+z)$] at the 99.7 per cent level whereas those with $\text{zWarning} > 0$ are reliable just 67 per cent of the time. For CMASS targets, good redshifts are

obtained for 98.2 per cent of targets; for LOWZ, it is 99.6 per cent. Although this completeness is quite high, one may worry that the failures may relate to observational systematics or have a preferred location on an observing tile.

BOSS fibres are numbered such that a given fibre corresponds to a particular position on the CCD. The spectrograph optics PSF degrades near the edges of the CCDs, thus lowering the quality of the extracted spectra and reducing the likelihood of obtaining good quality redshifts from spectra near the CCD edges. (See Smee et al., in preparation for more details on the performance of the BOSS spectrograph.) This correlation between redshift quality and fibre number translates into a spatial dependence on the sky, given that fibres are not assigned randomly. In order to test this effect, we translate all of the fibre positions of galaxies targeted by BOSS to positions relative to the centre of the tile. This allows us to determine the redshift failure rate as a function of position on the tile (and thus whether redshift failures may impart angular fluctuations in the density of observed galaxies). The result of this test is displayed in Fig. 3. The failed redshifts are not only more likely to be on the edge of a tile, but appear concentrated near the minimum and maximum right ascension of each tile. We apply weights (see Section 3.2) to correct for this spatial dependence, but find there to be a negligible effect on the measured clustering (see Fig. 5).

Fig. 4 shows the galaxy spatial number density for the CMASS sample, as a function of redshift. We also plot the normalized (so that it has the same integral) number density against redshift for the galaxies nearest to a redshift failure, $n_{nzf}(z)$. Were there a strong trend with redshift, for example that we were only missing redshifts for high-redshift galaxies, we should expect that the nearest neighbours to the redshift failures (which should be selected with similar properties, such as fibre ID, seeing and extinction) should predominantly be at low redshift. In fact, we see no such trend – if anything we find evidence to the contrary. We estimate the uncertainty on $n_{nzf}(z)$ in each z bin by assuming a Poissonian distribution for the number counts in each bin and we determine the χ^2 when the (nor-




Figure 3. The percentage of failed CMASS redshifts as a function of the position on the tile, averaged over 817 DR9 tiles. The lightest regions are 0 per cent and the darkest regions are 12 per cent. $\Delta\bar{\alpha}$ is the distance along the right ascension direction and $\Delta\bar{\delta}$ is the distance along the declination direction (both transformed so that the true angular separations are represented).




Figure 4. Galaxy spatial comoving number density assuming a flat Λ CDM cosmology with $\Omega_m = 0.274$, for CMASS galaxies. The solid line is calculated for all galaxies, while the dashed line only includes those galaxies nearest to a redshift failure, renormalized to match the total density of the full sample. The error bars assume Poissonian distribution for the number counts in each bin.

malized) $n_{nzf}(z)$ is compared to that of the full $n(z)$. We find $\chi^2 = 34.4$ for $0.43 < z < 0.7$ (27 bins; 15 per cent of consistent samples would have a higher χ^2) and $\chi^2 = 23.7$ for $0.5 < z < 0.7$ (20 bins; 26 per cent of the consistent sample would have a higher χ^2). We therefore find no evidence that the spatially dependent component (which is the component that we are interested in, as it may create a spurious clustering signal) of the redshift-failure probability is dependent on redshift.

3 ANALYSIS TECHNIQUES

3.1 Clustering estimators

We use the Landy & Szalay (1993) estimator to calculate the anisotropic redshift-space correlation function, $\xi(s, \mu)$, where s is the redshift-space separation in h^{-1} Mpc and μ is the angle to the line of sight:

$$\xi(s, \mu) = \frac{DD(s, \mu) - 2DR(s, \mu)}{RR(s, \mu)} + 1, \quad (19)$$

where D represents the data sample (i.e. BOSS galaxies) and R represents the random sample (occupying the angular footprint and with the same redshift distribution as the data sample) and the pair counts are normalized to the total number.

In linear theory, the first three moments of $\xi(s, \mu)$, expanded in Legendre polynomials, contain all of the information:

$$\xi_\ell(s) = \frac{(2\ell + 1)}{2} \int_{-1}^1 d\mu P_\ell(\mu) \xi(s, \mu). \quad (20)$$

We therefore weight pairs by P_ℓ (yielding separate pair counts for $\ell = 0, 2, 4$ for each of DD , DR and RR). Labelling the P_ℓ weighted pair counts with subscript ℓ , we determine $\xi_\ell(s)$ via

$$\frac{2\xi_\ell(s)}{2\ell + 1} = \frac{DD_\ell(s) - 2DR_\ell(s) + RR_\ell(s)}{RR_0(s)}. \quad (21)$$

We count pairs in a bin $7 h^{-1}$ Mpc wide in s and focus our efforts on understanding ξ_0 and ξ_2 [rather than the full $\xi(s, \mu)$] as these two measurements are expected to contain almost all of the information. In general, one must be careful, as our procedure implicitly assumes that pairs are isotropically selected in μ , and this is not true for a general survey geometry. In this study, we do not need to be concerned, as we will always be comparing our results to those obtained via mock catalogues (which accurately match the survey geometry; see Section 3.3 and Manera et al. 2012), but it must be accounted for in studies that compare more general models to the data. For a discussion of the procedures one may use, in general, when a survey does not provide an isotropically distributed sample of pairs, see Samushia, Percival & Raccanelli (2012).

We also use HEALPIX (Górski et al. 2005) maps, using $N_{\text{side}} = 256$, to calculate projected auto-/cross-correlation functions, $\xi_p(r_{\text{eff}})$ of galaxies and potential systematics (such as Galactic extinction). We split the sample into redshift shells of width $\Delta z = 0.01$ and define the overdensity, δ , in redshift shell and pixel i as

$$\delta_{i,z} = x_{i,z}/x_{\text{ave},z} - 1, \quad (22)$$

where x_i is the value of the quantity in question in pixel i and x_{ave} is the average of the quantity over all pixels. We can thus calculate ξ_p , as a function of the effective scale, r_{eff} , using pixellized maps via

$$\xi_p(r_{\text{eff}}) = \frac{\sum_{i,j,z_1,z_2} \delta_{i,z_1}^1 \delta_{j,z_2}^2 \Theta_{i,j,z_1,z_2}(r) N_1(z_1) N_2(z_2) w_i w_j}{\sum_{i,j,z_1,z_2} \Theta_{i,j,z_1,z_2}(r) N_1(z_1) N_2(z_2) w_i w_j}, \quad (23)$$

where the indices i, j represent the angular positions of pixels i and j and z_1, z_2 represent redshift slices, $\Theta_{i,j,z_1,z_2}(r)$ is 1 if the distance between the pixels (as determined by i, j, z_1, z_2) is within the bin defined by $r_{\text{eff}} \pm \delta r_{\text{eff}}$ and 0 otherwise, $N(z_i)$ is the number of galaxies in shell z_i , and w_i is the weight of the pixel (see the following section), which we determine using the random catalogue. It is straightforward to insert a purely angular map (e.g. the Galactic extinction) to determine how its angular cross-correlation with the galaxy field is translated to the physical scale r_{eff} . One simply holds its overdensity field constant with redshift and assigns it a flat $n(z)$ and proceeds through the sum defined above.

We also calculate power spectra, $P(k)$, using the standard Fourier technique of Feldman, Kaiser & Peacock (1994), as described by Reid et al. (2010). Here the window function is accounted for as a convolution of the model. This implies that power spectra can only be easily compared, without correcting for varying window functions, when they have the same selection and weighting. We therefore predominantly use the correlation function to present results, but we will show $P(k)$ measurements in Section 7, as these measurements isolate the largest wavelength density perturbations (separate bins in s are highly covariant).

3.2 Weights

We use weights to account for spatial variations in redshift failures, fibre collisions and imaging systematics, i.e. those effects that are not quantified via the C_{BOSS} completeness as described above. Given the total weight, w_{tot} , for each galaxy,

$$DD(s, \mu) = \sum_i \sum_j w_{\text{tot},i} w_{\text{tot},j} \Theta_{ij}(s, \mu), \quad (24)$$

where $\Theta_{ij}(s, \mu)$ is 1 if the separation between the two galaxies and the angle they make to the line of sight is within the particular bin, and 0 otherwise. These weights correct the galaxy densities to provide a more isotropic selection. They should therefore not be applied to a random catalogue, if it is based on an isotropic selection.

We will find that there is a systematic relationship between the density of targets (see Section 5) and the density of stars, and we therefore require a weight, w_{star} . We describe how we determine w_{star} in Section 5. We account for systematics that affect the process of going from the target catalogue to a redshift catalogue with weights for both redshift failures, w_{zf} , and fibre collisions with targets of the same type ('close pairs'), w_{fc} . We start by assigning each w_{zf} and w_{fc} unit weight. We have found that the probability of a redshift failure depends on the position of the fibre on the tile centre (see Fig. 3). To account for this, we find the nearest neighbour on the sky to the object with the failed redshift, and increase its w_{zf} by one. Such a weighting preserves the large-scale angular auto-correlation function. Furthermore, this action should not bias the redshift-space correlation function, as we find no evidence of a redshift dependence on the probability of a redshift failure (see Fig. 4). Similarly, we determine w_{fc} by up-weighting the colliding neighbour by one. Guo, Zehavi & Zheng (2011) have developed an optimized method to account for fibre collisions at all scales, but at large scales ($s \gtrsim 10 h^{-1}$ Mpc), our method should produce identical results. We combine the redshift failure and fibre-collision weights into a single weight equal to $w_{fc} + w_{zf} - 1$. Thus, the total weight we apply to each galaxy is

$$w_{\text{tot}} = w_{\text{star}}(w_{fc} + w_{zf} - 1). \quad (25)$$

The weight for any galaxy can be arbitrarily large, e.g. a galaxy will be given a weight of three if it is nearest to a target with a redshift failure and also causes a fibre collision with another galaxy (and has $w_{\text{star}} = 1$).

Fig. 5 displays the difference we find in the measured ξ_0 and ξ_2 of CMASS galaxies when applying the redshift failure (red) and fibre-collision (blue) weights and their combination (black). At scales greater than $50 h^{-1}$ Mpc, the effect is negligible compared to the expected statistical uncertainty for the sample (displayed in dashed black lines). Redshift failures have no significant effect on the measured clustering.

The application of the fibre-collision weights increases the clustering amplitudes at scales less than $80 h^{-1}$ Mpc. This effect is expected, as fibre collisions will be more likely to occur in highly clustered samples. At $20 h^{-1}$ Mpc, the difference is nearly 1σ . The weights for redshift failures impart a slight decrease in the measured amplitudes, at a level consistently less than 20 per cent of the statistical uncertainty determined using the mock galaxy catalogues (see Section 3.3).

We also want to optimally weight galaxies based on their number density, as suggested by Feldman et al. (1994). We refer to these weights as 'FKP' weights. To do so, we use a constant $P_{\text{FKP}} = 20000 h^3 \text{Mpc}^{-3}$ [as do Anderson et al. 2012; this is roughly the amplitude of the CMASS $P(k)$ at $k = 0.1 h \text{Mpc}^{-1}$] and weight by

$$w_P = 1/(1 + n(z)P_{\text{FKP}}), \quad (26)$$

where $n(z)$ is the number density of galaxies at redshift z , determined using our assumed cosmology. The purpose of these weights is to optimally weight areas with different number densities, not to correct observed number densities for a systematic effect. Thus, they are applied to both the random objects and the real galaxies; the final pair counts can be expressed as

$$DD_\ell(s) = \sum_{D_i} \sum_{D_j} w_{P,D_i} w_{\text{tot},D_i} w_{P,D_j} w_{\text{tot},D_j} \Theta_{D_i D_j}(s) P_\ell(\mu), \quad (27)$$

$$DR_\ell(s) = \sum_{D_i} \sum_{R_j} w_{P,D_i} w_{\text{tot},D_i} w_{P,R_j} \Theta_{D_i R_j}(s) P_\ell(\mu), \quad (28)$$




Figure 5. The effect of including weights for redshift failures (w_{zf} , red) and un-observed close pairs due to fibre collisions (w_{fc} , blue) and their combination (black) on ξ_0 and ξ_2 . The dashed lines display the 1σ statistical uncertainty expected from mock catalogues.

$$RR_\ell(s) = \sum_{Ri} \sum_{Rj} w_{P,Ri} w_{P,Rj} \Theta_{RiRj}(s) P_\ell(\mu). \quad (29)$$

3.3 Covariance matrices

We use mock galaxy catalogues with realistic clustering to determine covariance matrices for the distributions of galaxies and their clustering measurements. The mock catalogues have been produced with the method explained in Manera et al. (2012). This method is inspired by the perturbation theory haloes (PTHalos) paper of Scoccimarro & Sheth (2002), and have been calibrated using 40 N -body galaxy and halo mock catalogues generated using the LASDAMAS simulations⁴ (McBride et al., in preparation).

Using second-order perturbation theory (2LPT), Manera et al. (2012) generate 600 matter fields at redshift 0.55 drawn using periodic boxes of size $L = 2400 h^{-1} \text{ Mpc}$ (one matter field from each box). Each uses a flat cosmology defined by $\Omega_m = 0.274$, $\Omega_b h^2 = 0.0224$, $h = 0.70$, $n_s = 0.95$ and $\sigma_8 = 0.8$. This is the same cosmology as of White et al. (2011) and is close to 7-year *Wilkinson Microwave Anisotropy Probe* (WMAP7) parameters (Larson et al. 2011). Haloes from 2LPT runs are identified using a friends-of-friends algorithm and are then mass-calibrated using the Tinker et al. (2010) mass function. These haloes are then populated with galaxies using the halo-occupation-distribution (HOD) parametrization defined by Zheng, Coil & Zehavi (2007). We determine the HOD parameters by fitting the CMASS $\xi_0(s)$ measurement using data in the range of $30 < s < 80 h^{-1} \text{ Mpc}$.

For each mock realization, the periodic boxes are reshaped to match the final BOSS geometries for the NGC or the SGC footprints (see Fig 1; the box sizes are not large enough to accommodate both NGC and SGC simultaneously). Redshift distortions are then applied based on the 2LPT velocity field, combined with a model for the intrahalo velocity dispersion. The final DR9 angular masks are then applied, and galaxies are sampled from the full simulations to match the CMASS radial selection function (displayed in Fig. 4),




Figure 6. Top panels: the average of $\xi_\ell(s)$ we determine using 600 mocks, for the NGC (black) and SGC (red) footprints and with (solid lines) and without (dashed lines) FKP weights. The BAO peak can clearly be seen at $s \sim 100 h^{-1} \text{ Mpc}$ in ξ_0 . Bottom panels: the standard deviation of the 600 mocks, using the same scheme as the top panels.

thus yielding 600 mock galaxy catalogues mimicking the clustering of BOSS CMASS galaxies. For further details of the methods, see Manera et al. (2012).

For each of the 600 mock galaxy catalogues, in both the NGC and SGC, we calculate ξ_ℓ with and without FKP weights (see Section 3.2). The mean results of these measurements (top panels) and their standard deviations (bottom panels) are displayed in Fig. 6. The FKP weights decrease the standard deviation, typically by 10 per cent. We note that from here on we do not test $\xi_4(s)$ measurements, as in Reid et al. (2012) the information added by incorporating $\xi_4(s)$ affords only marginal improvements.

The covariance of $\xi_\ell(s)$ at separations s_1, s_2 and moments ℓ, ℓ' is determined using the standard mathematical definition, i.e.

$$C_{\ell,\ell'}(s_1, s_2) = \frac{1}{599} \sum_{i=1}^{600} (\xi_\ell^i(s_1) - \xi_\ell^{\text{ave}}(s_1)) (\xi_{\ell'}^i(s_2) - \xi_{\ell'}^{\text{ave}}(s_2)). \quad (30)$$

⁴ <http://lss.phy.vanderbilt.edu/lasdama/>




Figure 7. The normalized covariance matrix of and between ξ_ℓ determined using 600 mock galaxy catalogues simulating the BOSS DR9 selection function.

To calculate the covariance matrix over the total footprint, we assume that the two regions are independent and thus $C_{\text{total}}^{-1} = C_{\text{North}}^{-1} + C_{\text{South}}^{-1}$. We use the full covariance matrix to calculate χ^2 in the usual manner and apply tests on how our treatment of the data affects derived parameters. Any time we quote a χ^2 value, it is calculated using the covariance matrix determined from the mocks. The normalized covariance matrix, determined from the mocks, of and between ξ_ℓ is displayed in Fig. 7.

3.4 Modelling

We use the mean of clustering measurements determined using the mock catalogues to define the fiducial models we test. We derive parameters relating to the amplitude of the real-space galaxy density field, \tilde{b} , the amplitude of the galaxy velocity field, \tilde{f} , and the factor by which the distances assumed by our cosmological model are incorrect, $\tilde{\alpha}$ (as this relates to the position of the BAO feature, see Anderson et al. 2012). We expect that if we obtain robust results on these parameters, one should be able to derive robust constraints on any derived parameter, over the same range in s .

We assume a linear model for redshift-space clustering (Kaiser 1987) and assume linear biasing. In this model, given the real-space correlation function of matter, ξ_M ,

$$\xi_0(b, f) = \xi_M(b^2 + 2/3bf + 0.2f^2), \quad (31)$$

$$\xi_2(b, f) = \left(\frac{4}{3}bf + \frac{4}{7}f^2\right) [\xi_M - \xi'_M], \quad (32)$$

$$\xi_4(f) = \frac{8}{35}f^2 \left[\xi_M + \frac{5}{2}\xi'_M - \frac{7}{2}\xi''_M\right], \quad (33)$$

and

$$\xi' \equiv 3s^{-3} \int_0^s \xi(r')(r')^2 dr', \quad (34)$$

$$\xi'' \equiv 5s^{-5} \int_0^s \xi(r')(r')^4 dr' \quad (35)$$

(Hamilton 1992) where b is the real-space linear bias of the galaxy sample and f is the rate of change of the linear growth rate: $f \equiv d \log D / d \log a$, where D is the linear growth rate and a is the scale

factor of the Universe. We therefore determine our model ξ_ℓ^{mod} by rescaling the average mock ξ_ℓ^{mock} according to

$$\xi_0^{\text{mod}}(\tilde{b}, \tilde{f}) = \xi_0^{\text{mock}}(\tilde{b}^2 + 2/3\tilde{b}\tilde{f} + 0.2\tilde{f}^2)/(4.657), \quad (36)$$

$$\xi_2^{\text{mod}}(\tilde{b}, \tilde{f}) = \xi_2^{\text{mock}}(4/3\tilde{b}\tilde{f} + 4/7\tilde{f}^2)/(2.188), \quad (37)$$

$$\xi_4^{\text{mod}}(\tilde{f}) = \xi_4^{\text{mock}}\tilde{f}^2/0.548, \quad (38)$$

The values in the denominators account for a bias of the mocks of $b = 1.9$ (fits to the model of Reid & White 2011 suggest this is accurate to within 2 per cent and we fix it at 1.9 for simplicity) and the fact that $f(z = 0.55) = 0.74$ for the cosmology used by the mocks (flat $\Omega_m = 0.274$; e.g. the denominator for equation 36 is given by $1.9^2 + 2/3 \times 1.9 \times 0.74 + 0.2 \times 0.74^2 = 4.657$). This model thus assumes the same scalings in amplitude as linear RSD theory, but uses the shape of the mean mock galaxy $\xi_\ell(s)$, which include non-linear RSD features. We note that linear theory is not appropriate to obtain accurate estimates of b and f (see Reid & White 2011); it is for this reason that we quote our measurements as e.g. \tilde{b} rather than b . However, we expect that if our treatment of the data yields robust estimates of \tilde{b} and \tilde{f} , the measurements will also be robust when more accurate models are tested in Samushia et al. (in preparation) and Reid et al. (2012).

We also test the robustness of the data to a simple dilation in scale, using models where we vary a ‘stretch parameter’, $\tilde{\alpha}$. In this case

$$\xi_0^{\text{mod}}(\tilde{b}, \tilde{\alpha}, s) = \xi_0^{\text{mock}}(\tilde{\alpha}s)(\tilde{b}^2 + 0.48\tilde{b} + 0.1036)/(4.657) \quad (39)$$

(where we are now fixing \tilde{f}). We determine $\xi_0^{\text{mock}}(\tilde{\alpha}s)$ via power-law interpolation of the fiducial mean mock result at scales $s < 80 h^{-1}$ Mpc and linear interpolation at larger scales. We note that this stretch parameter will contain information on, at least, $\Omega_m h^2$ (from the overall peak of the power spectrum) and the position of the BAO peak. We therefore believe that if our treatment of the data yields robust $\tilde{\alpha}$ values, robust measurements of both the BAO position and $\Omega_m h^2$ will be obtained when more sophisticated models are applied, as in Anderson et al. (2012) and Sanchez et al. (2012). We study the recovered values of \tilde{b} , \tilde{f} and $\tilde{\alpha}$ according to Galactic hemisphere (Section 4.2), angular weights (Section 5.3) and redshift (Section 6.2) and these results are summarized in Table 1, found in Section 8.

3.5 Default survey window

The following three sections define and justify our recommendations for how to treat the survey window in regard to the NGC and SGC footprints, photometric systematics and the radial selection function. In each section, these recommendations are used, unless otherwise noted. The default is to treat the NCG and SGC as having separate selection functions and optimally combining their individual pair counts, apply weights to each galaxy based on linear relationships between $i_{\text{fib}2}$ magnitude and stellar density, and apply redshifts to the random sample by randomly selecting redshifts from the galaxy sample.

4 DEPENDENCE ON GALACTIC HEMISPHERE

The SDSS imaging was carried out in two large contiguous areas in the NGC and SGC (see Fig. 1). The mean sky background and airmass are both higher for the SGC imaging, which results in larger uncertainties on the measured magnitudes of these data, as the mean uncertainty for i -band CMASS targets is 0.076 in the NGC

Table 1. The parameters and χ_{\min}^2 we derive from the clustering of BOSS DR9 CMASS galaxies, for different treatments and subsamples of the data: \tilde{b} is a measure of the amplitude of the measured clustering, \tilde{f} is a measure of the amplitude of ξ_2 and $\tilde{\alpha}$ measures the preferred dilation in scale, relative to the average of the mock ξ_0 .

Estimator	Hemisphere	z range	w_{star} weights?	$\tilde{b}, \chi^2/\text{dof}$	$\tilde{f}, \chi^2/\text{dof}$	$\tilde{\alpha}$
$\xi(s)$	Both	$0.43 < z < 0.7$	Yes	$1.936 \pm 0.035, 22.7/17$	$0.711 \pm 0.044, 11.8/17$	1.020 ± 0.019
$P(k)$	Both	$0.43 < z < 0.7$	Yes	$1.983 \pm 0.035, 19.2/11$	–, –	–
$\xi(s)$	NGC	$0.43 < z < 0.7$	Yes	$1.904 \pm 0.039, 24.3/17$	$0.691 \pm 0.052, 10.9/17$	0.994 ± 0.023
$P(k)$	NGC	$0.43 < z < 0.7$	Yes	$1.991 \pm 0.039, 18.5/11$	–, –	–
$\xi(s)$	SGC	$0.43 < z < 0.7$	Yes	$2.06 \pm 0.07, 18.8/17$	$0.79 \pm 0.09, 11.8/17$	1.083 ± 0.029
$P(k)$	SGC	$0.43 < z < 0.7$	Yes	$1.93 \pm 0.07, 11.9/11$	–, –	–
$\xi(s)$	Both	$0.43 < z < 0.52$	Yes	$1.85 \pm 0.06, 15.1/17$	$0.59 \pm 0.08, 6.9/17$	1.016 ± 0.038
$\xi(s)$	Both	$0.52 < z < 0.7$	Yes	$2.02 \pm 0.04, 22.9/17$	$0.75 \pm 0.05, 14.1/17$	1.013 ± 0.021
$\xi(s)$	Both	$0.43 < z < 0.7$	No	$1.949_{-0.035}^{+0.034}, 34.2/17$	$0.710 \pm 0.044, 12.7/17$	1.007 ± 0.019
$P(k)$	Both	$0.43 < z < 0.7$	No	$2.001 \pm 0.035, 38.8/11$	–, –	–

and 0.101 in the SGC. However, we show in Sections 5 and 6 that the projected number density and redshift distributions of BOSS targets do not depend on either sky background or airmass and we therefore find no evidence that the observing conditions should produce systematic differences between the properties of targets selected in the SGC and NGC. Of greater concern is the fact that the two regions are tied together with relatively few scans to measure the relative photometric calibration (Padmanabhan et al. 2008). This suggests the possibility of a significant photometric offset between these two regions.

Schlafly et al. (2010) and Schlafly & Finkbeiner (2011) have found systematic variations in the colours of the population of SDSS stars as a function of their position. These offsets reflect a combination of variations in stellar populations across the Galaxy, calibration errors in the SDSS photometry (at the 1 per cent level) and errors in the corrections for Galactic extinction. In particular, they find that there is a systematic offset in the measured photometry between the SGC and NGC (the amplitude of this offset is within the expected 1 per cent rms of DR8 photometric calibration errors). The CMASS cut is sensitive to the d_{\perp} colour, due to both the hard cut (equation 3) and the sliding cut (equation 5). The LOWZ sample is sensitive to the c_{\parallel} colour, due to its sliding cut (equation 9). Schlafly & Finkbeiner (2011) find a 0.015-mag offset in c_{\parallel} and a 0.0064-mag offset in d_{\perp} between the NGC and SGC (based on their ‘spectrum based’ method; see their table 6). Ross et al. (2011b) found that the 2 per cent difference in the number density of CMASS targets between the NGC and SGC hemispheres was consistent with this offset in d_{\perp} . In what follows, we repeat and improve upon the analysis performed in Ross et al. (2011b) using only spectroscopically confirmed galaxies (albeit over a footprint one-third the size).

4.1 Number densities

Fig. 8 displays the distribution in galaxies versus the colour/magnitude information used to select them. We show the relations for spectroscopically identified galaxies and apply the redshift failure and close pair weights described in Section 3.2 when determining number densities. The left-hand panel shows the distribution of LOWZ galaxies against the value of the sliding cut, with the relationship for the NGC plotted in black and the relationship for the SGC plotted in red. At the cut (at 4.05 mag) the slope of the number density relationship is roughly 1.4×10^6 galaxies $\text{sr}^{-1} \text{mag}^{-1}$ change in c_{\parallel} . Thus, a 0.015-mag offset in c_{\parallel} (as implied by Schlafly & Finkbeiner 2011) should cause a change of $\sim 2 \times 10^4$ galaxies

sr^{-1} . If we apply this offset, we get the distribution in blue in the left-hand panel of Fig. 8. The curve appears much more consistent with the distribution in the NGC (black) than the fiducial SGC relationship (red).

In the DR9 CMASS sample, we find a 3.2 per cent higher number density in the SGC than in the NGC ($2.832 \times 10^5 \text{sr}^{-1}$ compared to $2.744 \times 10^5 \text{sr}^{-1}$). This is 1.2 per cent higher than that found by Ross et al. (2011b) due to a combination of the different footprints and the fact that we use spectroscopically identified galaxies. The middle panel displays the distribution of CMASS galaxies against their value of d_{\perp} , and the right-hand panel against the sliding cut. For the $d_{\perp} > 0.55$ cut, we should expect 7×10^5 galaxies $\text{sr}^{-1} \text{mag}^{-1}$ change in d_{\perp} , and for the sliding cut we should expect an extra 5×10^5 change in the number of galaxies $\text{sr}^{-1} \text{mag}^{-1}$ change in d_{\perp} . Thus, we should expect a change of roughly 1.2×10^6 targeted galaxies $\text{sr}^{-1} \text{mag}^{-1}$ offset in d_{\perp} (the two cuts are not strongly covariant). Therefore, given the 0.0064-mag offset in d_{\perp} between the NGC and SGC, we should expect 7700 galaxies sr^{-1} in the SGC. This is 2.7 per cent of the number density of the CMASS sample. When we apply this offset to the target selection, we find that there only remains a 0.2 per cent excess in the number density of objects in the SGC compared to the NGC.

Fig. 9 displays the number density as a function of redshift, $n(z)$, for CMASS galaxies in the NGC (black) and SGC (blue), using bins of width $\Delta z = 0.01$. The number density is 10 per cent smaller around the peak of the distribution in the SGC, and this becomes more dramatic when we apply the Schlafly & Finkbeiner (2011) offsets to the target selection. However, the number density is 20 per cent larger at $z = 0.6$. The error bars are determined from the variations we find in $n(z)$ of mocks cut to the same angular footprint as our SGC data set [$n(z)$ for each individual mock varies due to the cosmic variance inherent in large-scale structure]. Using these mocks, we can also determine the covariance between the $n(z)$ bins, thus allowing us to calculate χ^2 between the SGC and the NGC. When the offsets are applied, we find $\chi^2 = 36$ (for 27 bins). The χ^2 is higher (39) when the offsets are not applied to the target selection. We find that 55 of the 600 SGC mock $n(z)$ have χ^2 that is greater than 36 (when compared to the average of the 600 mocks) and 34 have a χ^2 that is greater than 39 (this is roughly in line with the probabilities of 11 and 6 per cent one obtains for these χ^2 values and 27 degrees of freedom, dof). Thus, applying the offsets makes the redshift distributions of the NGC and SGC samples more consistent, but the differences between them slightly unusual.

In summary, for both the CMASS and LOWZ samples, we find that the difference in their number densities is consistent with the




Figure 8. The distribution of DR9 spectroscopically identified galaxies as a function of colour combinations used in their selection, for objects in the NGC (black) and SGC (blue), and in SGC after applying the Schlafly & Finkbeiner (2011) offsets to the target selection (red). Dashed vertical lines display the location of the cut applied to BOSS targeting.




Figure 9. The galaxy spatial number density assuming a flat Λ CDM cosmology with $\Omega_m = 0.274$ of CMASS objects in NGC and SGC. The red line displays the result when we apply the Schlafly & Finkbeiner (2011) offsets to the target selection in the SGC. The error bars are determined using 600 mock catalogues cut to the angular footprint of the SGC.

level of the colour offset between the NGC and SGC as determined by Schlafly & Finkbeiner (2011) using their spectrum method. Furthermore, the offset is understood – it is within the expected rms of DR8 calibration errors and found between the two regions that have the least available data for relative calibration. We can apply these colour offsets to the selection of galaxies in the SGC in an attempt to make a homogenous sample, as doing so only makes the cuts more restrictive. However, there is some uncertainty inherent in the level of the offset, and, based on the mocks, the expected variance in the number density between the NGC and SGC is 2 per cent. Furthermore, the $n(z)$ distributions remain slightly inconsistent. We

therefore believe that the most conservative approach is to treat the two samples as having separate selection functions, due (at least in part) to the fact that there are offsets in the photometry between the two regions. Thus, we analyse all galaxies observed in the SGC separately, accepting that they comprise a denser sample than those in the NGC.

4.2 Measurements of clustering

Fig. 10 displays the measured $\xi(s)$ in the NGC (red) and SGC (blue). These measurements include the weights that correct for stars, fibre collisions and redshift failures, defined by equation (25), and also the FKP weights. The area covered by the SGC data is only one quarter that of the NGC, and therefore the uncertainty in the SGC $\xi(s)$ is about twice as large as the NGC. At almost all scales, the ξ_0 measurements appear consistent with each other; the only notable exception is a significant dip in the measurements at $170 h^{-1}$ Mpc. Both of the $\xi_0(s)$ measurements display a prominent increase in clustering at around $100 h^{-1}$ Mpc, suggesting significant BAO peaks, though the peak does appear at a smaller scale in the SGC measurement. Interestingly, both the NGC and SGC measurements appear to also have a peak in $\xi_0(s)$ at around $s = 215 h^{-1}$ Mpc, but we note that the uncertainty on the measurements at these scales is much larger than that around the BAO scale. The ξ_2 measurements appear slightly less consistent, especially between 75 and $95 h^{-1}$ Mpc, where the measurements are clearly inconsistent within the 1σ error bars.

The black points in Fig. 10 display combined NGC and SGC measurements, which are produced by summing the DD, DR and RR pairs, which is appropriate when FKP weights are used. The number of randoms in each region use the same normalization with respect to the number density in each region (we use just over 15 times the number of galaxies). Thus, the relative normalization of randoms to galaxies between the two hemispheres is matched as if the samples were treated individually, and the results are optimally combined.

We test the consistency of the measurements by summing the covariance matrices of the NGC and SGC and determining χ^2 in the




Figure 10. Top panels: the measured redshift-space correlation functions, ξ_0 and ξ_2 , of CMASS data in the NGC (red) and SGC (blue), and their pair-weighted average (black triangles), using FKP weights and the w_{star} weights. The error bars are the standard deviations of the ξ_ℓ in the mocks drawn from the SGC footprint. For both the NGC and SGC measurements, the BAO feature is apparent at $s \sim 100 h^{-1}$ Mpc. Bottom panels: the difference between the measured $\xi_{0,2}$ of the NGC (red) and SGC (blue) CMASS samples and the mean of their respective mocks, after scaling the mocks for a best-fitting bias. The error bars are the standard deviations of the ξ_ℓ in the mocks drawn from the respective SGC and NGC footprints.

standard fashion. For $s < 250 h^{-1}$ Mpc (35 data points), we find $\chi^2 = 45.4$ for ξ_0 and 32.0 for ξ_2 , so despite the apparent differences, ξ_2 is actually more consistent than ξ_0 . For ξ_0 , the χ^2 is slightly high – 11 per cent of consistent samples drawn from a Gaussian distribution will have a higher χ^2 . Reducing the range of the fit to $25 < s < 150 h^{-1}$ Mpc (the primary range we study), χ^2 is 23.2 (18 data points), which shows consistency (18 per cent of consistent samples will have a larger χ^2). Sanchez et al. (2012) find the same χ^2/dof (1.3) when they fit their $\xi(s)$ measurements, which use an alternative binning, at $40 < s < 200 h^{-1}$ Mpc. Scaling to the NGC sample, we find that the best-fitting relative bias, b_{rel} , of the SGC sample is 1.057 ± 0.038 ($\chi_{\text{min}}^2 = 20.9$), when fitted to $25 < s < 150 h^{-1}$ Mpc. Relaxing the minimum bound to $s = 10 h^{-1}$ Mpc (adding two data

points), we find $b_{\text{rel}} = 0.983 \pm 0.015$ with $\chi_{\text{min}}^2 = 20.5$, $b_{\text{rel}} = 1$ is just outside the 1σ bounds, suggesting that the clustering in the two regions is consistent to within 1.1σ .

We further test the consistency of the NGC and SGC measurements by finding the best-fitting bias for the NGC and SGC samples by scaling the mocks, as described in Section 3.3, in the range $25 < s < 150 h^{-1}$ Mpc. We find the best-fitting $\tilde{b} = 1.904 \pm 0.039$, with $\chi_{\text{min}}^2 = 24.3$ (18 measurements) for the NGC data and 2.06 ± 0.07 , with $\chi_{\text{min}}^2 = 18.8$ for the SGC. The difference is nearly 2σ . The best-fitting bias of the combined sample, $\tilde{b} = 1.936 \pm 0.035$, is very close to the weighted average of the two samples ($\tilde{b} = 1.943 \pm 0.035$). Increasing the minimum scale to $s = 30 h^{-1}$ Mpc, we find a significant change in the best-fitting bias for the NGC ($\tilde{b} = 1.87 \pm 0.05$),

but we find negligible change for the SGC ($\bar{b} = 2.08 \pm 0.11$). For the combined sample, the bias decrease is even larger, as $\bar{b} = 1.886 \pm 0.048$ when we fit $30 < s < 150 h^{-1}$ Mpc. The values of \bar{b} (and other parameters we measure throughout this section) are summarized in Table 1 (found in Section 8).

Fixing the bias at the best-fitting value from ξ_0 , we find the best-fitting value of \tilde{f} from the ξ_2 measurements by scaling the mock ξ_2 . For the NGC, we find $\tilde{f} = 0.691 \pm 0.052$ (with $\chi^2_{\min} = 10.9$) and for the SGC data $\tilde{f} = 0.79 \pm 0.09$ (with $\chi^2_{\min} = 11.8$). For the combined sample, $\tilde{f} = 0.711 \pm 0.044$ ($\chi^2_{\min} = 10.2$) – very similar to the weighted average of 0.716 ± 0.045 . This suggests that the information content in $\xi_2(s)$ measurements related to the velocity field is consistent between the two regions.

As described in Section 3.4, we can use the mocks to fit for a bias and stretch parameter $\tilde{\alpha}$. We stress that these $\tilde{\alpha}$ values should reflect both changes we expect in the best-fitting $\Omega_m h^2$ and distance constraints one may obtain from the BAO feature, i.e. it only reflects the level of disagreement we should expect in derived cosmological parameters using the NGC/SGC footprint, and whether the level of disagreement is consistent with what we expect to find given cosmic variance, but it contains no information on specific parameters. For the NGC footprint, we find $\chi^2_{\min} = 22.1$ at $\tilde{\alpha} = 0.990$ and $\tilde{b} = 1.888$; marginalizing over the bias, $\tilde{\alpha} = 0.994 \pm 0.023$. For the SGC footprint, $\chi^2_{\min} = 12.3$ at $\tilde{\alpha} = 1.090$ and $\tilde{b} = 2.319$; $\tilde{\alpha} = 1.083 \pm 0.029$ when we marginalize over bias. For the combined sample, we find $\chi^2_{\min} = 21.5$ at $\tilde{\alpha} = 1.019$ and $\tilde{b} = 1.982$; marginalizing over the bias, $\tilde{\alpha} = 1.020 \pm 0.019$. This is smaller than the weighted average, 1.028 ± 0.018 , of the two $\tilde{\alpha}$ measurements, but it is still greater than a 1σ shift from the result obtained using only the NGC data. This implies that one may find differences of 1σ on recovered cosmological parameters when comparing results from only the NGC data to the combined sample. Indeed, Sanchez et al. (2012) find this level of variation.

The difference between the NGC and SGC $\tilde{\alpha}$ values we measure is 2.5σ . We find negligible changes in the values of $\tilde{\alpha}$ we obtain from the measured SGC ξ_0 if we apply the Schlafly & Finkbeiner (2011) offset to the selection of SGC CMASS galaxies, use any of the separate weighting schemes described in Appendix A, or neglect to apply any weight at all; that is, we have not been able to identify any systematic that may cause the differences we observe. Any true difference would represent a violation of isotropy. Anderson et al. (2012) find that the tension between their BAO scale measurements is reduced to 1.4σ when reconstruction is applied to the CMASS galaxy density field (and is 2.5σ without reconstruction). As reconstruction generally improves the signal-to-noise ratio in BAO scale estimation, this reduction in the tension between the two measurements implies that the difference is indeed driven by noise.

In general, the level of disagreement between the NGC and SGC correlation functions is between 1 and 2σ . The differences in the bias when scaling the mocks (1.9σ) and when we fit for a relative bias ($b_{\text{rel}} = 1$ is 1.5σ from the best fit) are both less than 2σ . The $n(z)$ distributions disagree at a similar level of significance. The $n(z)$ discrepancy is likely related to the disagreement in the clustering. Finally, when fixing $\tilde{\alpha} = 1.017$ (the upper 1σ bound on $\tilde{\alpha}$ from the NGC sample), $\chi^2_{\min} = 16.4$ (meaning that χ^2/dof is less than one) when testing ξ_0 of the SGC sample (at $\tilde{b} = 2.11$). The SGC footprint is currently only $705 \text{ deg}^2 - 28$ per cent of its final (planned) size (2500 deg^2). If the differences are indeed in the noise, we expect that all results between the NGC and SGC will become more consistent as the BOSS survey continues and the sample grows.

5 ANGULAR VARIATIONS IN TARGET CATALOGUE

Ross et al. (2011b) found significant correlations between the number density of galaxies in the SDSS imaging data with a photometric selection similar to that of the CMASS sample, and various parameters. In particular, the number density of observed galaxies decreased significantly as a function of the stellar density. We repeat the tests performed by Ross et al. (2011b), now using spectroscopically confirmed galaxies. We are only using data within the DR9 mask (which is about one-third of the imaging area used in Ross et al. 2011b) and we now have access to mocks that allow us to quantify the statistical variations we should expect to find.

5.1 Galaxy density versus potential systematics

We determine the number density of the DR9 spectroscopically identified galaxies as a function of stellar density, seeing, Galactic extinction, airmass and sky background (all during the imaging observations). To perform these tests, we make HEALPIX (Górski et al. 2005) maps of DR9 galaxies and compare them to maps of the number of stars with $17.5 < i < 19.9$ or maps of the mean values of the potential observational systematic based on data from the SDSS DR8 Catalog Archive Server (CAS)⁵ within pixels at $N_{\text{side}} = 256$, which splits the sky into equal area pixels of 0.0525 deg^2 . Rather than pixellate the mask, in each pixel we determine the number of galaxies, n_{gal} , and the number of randoms, n_{ran} (multiplied by the factor $n_{\text{gal,tot}}/n_{\text{ran,tot}}$), and therefore map $n_{\text{gal}}/n_{\text{ran}}$.

Fig. 11 displays the relationships between the number of galaxies observed and potential observational systematics. We weight each galaxy for redshift completeness and close pairs as described in Section 3.2 and we combine the NGC and SGC data by using the same normalization of randoms to galaxies in each respective region. We find no significant differences in our analysis if we analyse the two regions separately. We apply the same tests on the 600 mocks (described in Section 3.3) and use the standard deviation as the errors in each bin displayed in Fig. 11. The relationships for CMASS galaxies are displayed in black. As in Ross et al. (2011b), we find a 10 per cent decrease in the number density of galaxies between areas with high and low stellar density.

As quantified in section 4.1 of Ross et al. (2011b), 3 per cent of the decrease in galaxy density results from the fact that, within 10 arcsec of stars, seeing reduces the ability to detect galaxies (with little dependence on the magnitude of the star at $17.5 < i_{\text{mod}} < 19.9$). The relationship is not found in DR7 data (A. Bauer, private communication). The most significant change between the DR7 and DR8 photometric pipelines was a refinement in the sky background subtraction algorithm, in order to improve the photometry of bright galaxies (Aihara et al. 2011). One effect of this change is to increase the low surface brightness extent of both galaxies and stars, causing more objects to be linked together. In regions of higher stellar density, this means that the deblender more often has to deal with complicated superpositions of many objects. To control processing time, the deblending code will separate out up to 25 overlapping objects in one parent, but no more. This happens more often with the new code, meaning that there are more missing galaxies in regions of high stellar density than before. This may explain the remaining 7 per cent effect.

We also find an anticorrelation with Galactic extinction – this is at least partly due to the fact that the Galactic extinction and

⁵ <http://skyserver.sdss3.org/dr8/en/>




Figure 11. The projected number density of CMASS galaxies as a function of the potential systematics: stellar density (n_{star}), Galactic extinction in the r -band (A_r), the i -band seeing (see_i), i -band background sky flux in nanomaggies arcsec $^{-2}$, denoted $nmagg$, ($'sky_i'$; $nmagg$ are related to the magnitude m , via $m = 22.5 - 2.5 \log(nmagg)$), and the airmass (air). The black lines display the results for the CMASS sample, and the green lines display the results for the LOWZ sample. The red lines display the residual CMASS relationships after applying weights that account for the linear relationships between galaxy density, stellar density and fibre magnitude (w_{star}). The expected errors are determined by finding the standard deviation of the relationships measured from individual mock CMASS catalogues.

stellar density are correlated. Indeed, we find that the correlation with extinction becomes insignificant once weights are applied (see Section 5.2) to correct for the relationship with stellar density. See Yahata et al. (2007) for a more detailed study on the ways in which the Galactic extinction, as determined by the Schlegel et al. (1998) dust maps, correlates with the observed density of galaxies.

We find a sharp decrease in the number density of galaxies in areas with poor seeing; this effect was explained in Ross et al. (2011b) as being due to the fact that the star/galaxy separation criteria defined by equations (7) and (8) remove more true galaxies in areas where the seeing is poor. This systematic relationship had little effect on the measured clustering in Ross et al. (2011b), as the pattern of seeing in the DR8 imaging is essentially random on large scales. For sky background and airmass, the level of fluctuations is close to what we should expect due to cosmic variance (as represented by the error bars).

The relationships for the LOWZ sample are displayed in Fig. 11 with green lines. In contrast to the CMASS sample, we do not find any systematic dependency with stellar density, Galactic extinction or seeing. This is likely due to the fact that LOWZ galaxies are, on average, considerably brighter than CMASS galaxies, and their detection should therefore be less affected by imaging systematics. Given that the volume of the LOWZ sample is considerably smaller than the CMASS sample, we should expect larger cosmic variance. Indeed, it appears that all of the variance in number density we find for the LOWZ sample can be attributed to cosmic variance.

In Ross et al. (2011b), the relationship between the number density of galaxies and stellar density was found to depend on the surface brightness of the galaxy. Given that the mean surface brightness of CMASS galaxies is lower at higher redshift, the relationship between the galaxy density and stellar density may depend on redshift. In Fig. 12, we show the relationship between galaxy density and stellar density when splitting the sample into three sets based on the i_{fib2} magnitude, as this magnitude uses a fixed aperture and is thus essentially a surface brightness measurement. The slope in the relationship clearly grows more negative at fainter i_{fib2} . The fact that the effect correlates strongly with surface brightness is further evidence that it is related to a systematic in the DR8 imaging, likely related to the sky subtraction routine, as opposed to a real large-scale density fluctuation perfectly aligned with the Galaxy.




Figure 12. Same as the black points in the left-hand panel of Fig. 11, except that we have broken the CMASS sample into three subsamples based on the labelled fibre magnitudes, i_{fib2} : $i_{fib2} < 20.75$ has 78 065 good redshifts, $20.75 < i_{fib2} < 21$ has 85 284 good redshifts and $i_{fib2} > 21$ has 112 304 good redshifts.

For each of the potential systematics displayed in Fig. 11, we determine the auto-correlation, ξ_p , and cross-correlation, $\xi_{p,x}$, with the CMASS sample as a function of the effective scale, r_{eff} (all of which are defined by equation 23 and the surrounding text). As described in Ho et al. (2012) and Ross et al. (2011b), the effect of any potential systematic on the measured correlation function can be estimated as $\xi_{p,x}(r_{eff})^2/\xi_p(r_{eff})$. Fig. 13 displays this ratio for the five potential systematics displayed in Fig. 11. We confirm with the spectroscopic sample the result found in the angular clustering analysis (Ross et al. 2011b): the presence of stars has the greatest systematic effect. The effect of Galactic extinction is second




Figure 13. The cross-correlation with the full CMASS sample, $\xi_{p,x}(r_{\text{eff}})$, squared, divided by the auto-correlation, ξ_p , for stars (orange), Galactic extinction (green), the sum of seeing, sky background and airmass (red), the sum of all five when applying the linear-fit weights to all five potential systematics (w_{MCMC} , blue, which we consider in Appendix A), the sum of all five when applying the linear-fit weights for only stellar density as a function of $i_{\text{fib}2}$ (w_{star} , light blue), and the mean sum of all five and its standard deviation on the mocks (black error bars). The dotted black line displays the expected statistical uncertainty, determined from the variance of the mock $\xi_0(s)$ measurements.

largest, but Ross et al. (2011b) found it to be almost entirely degenerate with the effect of stars. The sky background, airmass and seeing all have negligible effects. However, we have found more significant correlations with sky background when the sample is split further by colour. Overall, we should expect a difference of ~ 0.002 between fiducial $\xi(s)$ measurements and those with corrections for systematics.




Figure 14. The best-fitting coefficients to the relationship $n_{\text{gal}}/n_{\text{ran}} = A + Bn_{\text{star}}$ as a function of the $i_{\text{fib}2}$ magnitude of the galaxies. The blue points display the results when we use FKP weights (see equation 26) and the red points show the results when we do not use this weighting. The black lines display our fits to these coefficients, which we use to determine weights as a function of $i_{\text{fib}2}$ and n_{star} .

The mean of the sum and standard deviation of $\xi_{p,x}(r_{\text{eff}})^2/\xi_p(r_{\text{eff}})$ for all five potential systematics we consider over the mock catalogues is displayed with black error bars in Fig. 13. This is non-zero on average because the auto-correlation of each systematic is positive and the cross-correlations, which have zero mean, are squared. Thus, we should expect a non-zero mean on the sum of these contributions, even when there is no real systematic effect.

5.2 Angular weights

As shown in Fig. 13, the primary source of systematic error is due to the relationship with stellar density. To account for this systematic effect, we apply weights that counteract the systematic relationship. Fig. 12 suggests that the systematic relationships depend on the surface brightness of the galaxy. We thus use this information in order to determine ‘linear-fit stellar density weights’, which we denote w_{star} . We split the sample by $i_{\text{fib}2}$ and assume $n_{\text{gal}}/n_{\text{ran}} = A + Bn_{\text{star}}$ for each subsample, now using $N_{\text{side}} = 128$ for the resolution of the maps. The result of this approach is shown in Fig. 14. We find that for $i_{\text{fib}2} < 20.45$, the relationship is consistent with being constant. At fainter $i_{\text{fib}2}$, we find a linear relationship with $A(i_{\text{fib}2})$ and $B(i_{\text{fib}2})$, and we thus use this linear fit to determine the w_{star} weights (which ignores the rest of the potential systematics). The linear fit is given by

$$A = A_0 + A_1 i_{\text{fib}2}, \quad (40)$$

$$B/(\text{deg}^2) = B_0 + B_1 i_{\text{fib}2}, \quad (41)$$

where $A_0 = 3.96$, $A_1 = -0.14$, $B_0 = -1.18 \times 10^{-3}$ and $B_1 = 5.76 \times 10^{-5}$ (for the case where the FKP weights are applied). For $i_{\text{fib}2} < 20.45$, A and B are set to the $A(20.45)$ and $B(20.45)$ given by the above equations.

The residual relationships after applying the w_{star} weights are displayed in red in Fig. 11 – the relationship with Galactic extinction changes from having a slightly negative to a slightly positive slope and the seeing, sky background and airmass relationship remain similar to the un-weighted case. For all but seeing, the relationship appears consistent with the variations we expect due to cosmic variance (as shown by the error bars in Fig. 11). Furthermore, the sum of the five potential systematic contributions is consistent with the

mean mock sum when the w_{star} weights (as shown in Fig. 13) are applied to the CMASS data. We therefore believe that the w_{star} weights are appropriate to apply to the CMASS sample. Additionally, we believe the $i_{\text{fib}2}$ dependence of the w_{star} weights should (mostly) account for changes in redshift. In Appendix A, we compare the w_{star} weights to two other weighting schemes and find that the application of the w_{star} weights has the least potential to remove true fluctuations from the density field and does not bias the clustering measurements of our mock galaxy samples.

5.3 Effect of angular weights on CMASS clustering

The top panels of Fig. 15 display the resulting $\xi_0(s)$ and $\xi_2(s)$ measurements for CMASS galaxies for the fiducial sample (which in-

clude FKP weights and weights for close pairs and redshift failures, see Fig. 5), and when we include the w_{star} weights. As expected (Fig. 13), the weights for stellar density cause a nearly constant decrease of 0.002 in the measured ξ_0 . This change is greater than the statistical uncertainty at scales greater than $110 h^{-1} \text{Mpc}$. The w_{star} weights only have a slight effect on the ξ_2 measurements. The difference is always smaller than the statistical uncertainty, reaching $\sim 0.5\sigma$ at the largest scales.

As described in Section 3.3, we can scale the mocks to determine a best-fitting bias \bar{b} . The consistency of the different weighting schemes can be further tested by comparing the best-fitting bias and associated χ^2_{min} , which we determine in all cases using the covariance matrices calculated using the mocks. Applying the w_{star} weights to the CMASS sample when calculating $\xi(s)$ and fitting




Figure 15. The measured redshift-space correlation functions of CMASS galaxies using the fiducial catalogue (red triangles), and applying weights that correct for the linear relationships between galaxy density and stellar density and $i_{\text{fib}2}$ magnitude (blue circles). The bottom panels display the difference between the measured ξ_0 and ξ_2 and that mean calculated from the mock ξ_0 and ξ_2 . Black error bars represent the standard deviation of the mock ξ_0 and ξ_2 . We analyse the apparent feature at $s \sim 200 h^{-1} \text{Mpc}$ in Section 7.

at $25 < s < 150 h^{-1} \text{Mpc}$, the best-fitting $\tilde{b} = 1.936 \pm 0.035$, with $\chi_{\min}^2 = 22.7$ (18 measurements). When we do not apply the w_{star} weights to the CMASS sample, the best-fitting bias increases by $\sim 0.5\sigma$ to $\tilde{b} = 1.949^{+0.034}_{-0.035}$, and χ_{\min}^2 increases to 34.2. Assuming Gaussian statistics, only 1.2 per cent of consistent samples would have a $\chi^2 > 34.2$, while 20.2 per cent would have $\chi^2 > 22.7$. The values of \tilde{b} (and other parameters we measure throughout this section) are summarized in Table 1 (found in Section 8).

We find similar results when we fit the weighted and un-weighted ξ_0 measurements for both \tilde{b} and $\tilde{\alpha}$ (again in the range $25 < s < 150 h^{-1} \text{Mpc}$). The $\chi_{\min}^2 = 33.9$ at $\tilde{\alpha} = 1.009$ and $b = 1.971$ when the weights are not applied; marginalizing over the bias, we find $\tilde{\alpha} = 1.007 \pm 0.019$. When applying the w_{star} weights, we find $\chi_{\min}^2 = 21.5$ is at $\tilde{\alpha} = 1.019$ and $b = 1.982$; marginalizing over the bias, $\tilde{\alpha} = 1.020 \pm 0.019$. We note that this test only reflects the level of systematic change we should expect in derived cosmological parameters. The change in the $\tilde{\alpha}$ value therefore suggests that the application of the weights could cause a shift in the best-fitting cosmology of close to two-thirds σ when constraints are derived from the full shape of ξ_0 . Anderson et al. (2012) measure the same BAO peak position, to within 0.1 per cent, whether or not the w_{star} weights are applied, suggesting that the change in our $\tilde{\alpha}$ measurement reflects the change in the shape of $\xi_0(s)$.

Fig. 15 suggests that the weights have little effect on ξ_2 . This is confirmed by fixing \tilde{b} at the best-fitting value from the ξ_0 measurements and finding the best-fitting \tilde{f} value, by scaling the mocks. When the w_{star} weights are applied, we find $\tilde{f} = 0.711 \pm 0.044$ with $\chi_{\min}^2 = 11.8$ (18 data points); when the weights are not applied, we find $\tilde{f} = 0.710 \pm 0.044$ with $\chi_{\min}^2 = 12.7$.

6 RADIAL SELECTION FUNCTION

We may be concerned that any parameter that causes a systematic effect in the angular distribution of galaxies may also cause a change in the redshift distribution. To test this possibility, we split the sample in half, based on each of the same five potential systematics in turn, and determine the redshift distributions. The results are shown in Fig. 16 when using the weights fit to the linear relationships between galaxy density, stellar density and $i_{\text{fib}2}$ (w_{star}). None of the distributions is significantly outside of the errors we determine based on the standard deviation in the distributions of mock samples within the NGC or SGC footprints (though we only plot the results for the NGC; see Section 3.3).

We use FKP weights (Feldman et al. 1994), defined by equation (26), to optimally weight the data as a function of redshift. This changes $n(z)$ from the blue curves to the red in the upper right-hand panel of Fig. 16, with solid lines representing data from the NGC and dashed lines for the SGC. Including the FKP weights effectively equalizes the contribution of every redshift interval we consider in the ξ_ℓ calculation. This is illustrated by the fact that the mean redshift changes from $z = 0.55$ to 0.57 . The FKP weights also make the NGC and SGC selection functions more similar to one another.

We display ξ_ℓ measurements, with and without FKP weights, in Fig. 17. This serves as a check that these weights have not imparted any systematic errors and illustrates the advantage of applying FKP weights. For ξ_0 , the amplitudes are marginally higher at all scales when the FKP weights are applied. This result is due in part to the fact that the FKP procedure assigns larger weights to the higher redshift data, which is more likely to include more luminous galaxies which thus have higher bias. We find, as expected, that the application of the FKP weights reduces the uncertainty on derived




Figure 16. The redshift distributions of CMASS objects in the NGC, splitting the area in half on seeing, i -band sky background in nanomaggies arcsec $^{-2}$, the r -band Galactic extinction (A_r), airmass and the number density of stars (n_{star}). The errors are determined by finding the standard deviation of the mock $n(z)$. The upper right-hand panel displays the result with (red) and without (blue) FKP weights, using solid lines for the NGC sample and dashed lines for the SGC sample.




Figure 17. Top panels: the measured redshift-space auto-correlation functions of the combined CMASS sample measured with (blue) and without (red) FKP weights. Points represent measurements; error bars are determined using the variance of the mock calculations. Bottom panels: points display the difference between the measurements using and not using the FKP weights. Error bars represent the mean and the variance of the difference using and not using FKP weights for the mock calculations.

parameters by at least 10 per cent and that the derived \tilde{f} and $\tilde{\alpha}$ are consistent to within 0.5σ whether or not the FKP weights are applied (without FKP weights, we find $\tilde{\alpha} = 1.033 \pm 0.025$ and $\tilde{f} = 0.711 \pm 0.044$).

6.1 Testing models of the radial selection function

To model the expected galaxy distribution, we must assign redshifts to the random catalogues we used to calculate ξ_ℓ . This is difficult to achieve without using the data themselves, as we would need a full theoretical model for the galaxy population targeted. Without such a model, we are limited by the fact that we can only estimate the true $n(z)$ empirically. However, we can test the effects of this dependence using the mocks. The mocks were constructed assuming a fiducial $n(z)$ [which is $n(z)$ we measure from the data] and $n(z)$ of each individual mock will scatter around this input $n(z)$ due to cosmic variance. For each mock, we consider three methods to determine $n(z)$ applied to the random catalogue:

- (i) ‘Spline’, where a spline fit to the observed redshift distribution of galaxies, using bins of width $\Delta z = 0.01$, is used to determine $n(z)$, we sample from this to construct a random catalogue.
- (ii) ‘Shuffled’, where for each point in the random sample we assign the redshift of a randomly selected redshift from the galaxy sample.
- (iii) ‘True’, where we use the input $n(z)$. The true $n(z)$ is, of course, not available for any observed sample and we test the difference between clustering measured using the true $n(z)$ and either the spline or shuffled $n(z)$ in this section.

We create a random catalogue for each of the 600 individual mocks using both the spline method and the shuffled method and compare the results to those derived using the true underlying $n(z)$.




Figure 18. Top panel: average offsets and 1σ deviations from a true ξ_0 for different methods of assigning redshifts to random catalogues, determined using the 600 mock galaxy catalogues. The solid black line corresponds to the statistical error bars. Bottom panel: as above, for ξ_2 .

For the spline method, we use an N -node spline. We examine the cases where $N = 10, 20, 30$ and 50 . We expect the results using the N -node spline to approach the limit of the results of the shuffled catalogue when N is very large. Importantly, we can compare all results to the true case, thereby quantifying the bias and additional uncertainty imparted by the need to self-average as a function of redshift, and thus addressing the concerns raised in Sylos Labini et al. (2009).

The top panel of Fig. 18 shows the average bias of ξ_0 measurements and its standard deviation, determined using 600 realizations of ξ_0 computed from mocks using different random catalogues. For the measurements of the monopole the average bias for all methods of constructing a random catalogue is a small fraction of the statistical errors and the standard deviation of the bias is about a third of the statistical errors. The bias is smallest when using the shuffled random catalogue and appears negligible for ξ_0 .

The bottom panel of Fig. 18 shows the results of the same test on the measurements of ξ_2 . For the second Legendre moment of the correlation function, the systematic offset is larger than that for ξ_0 , but is still small compared to the statistical errors. The standard deviation of the bias is ~ 50 per cent of the statistical errors. For both monopole and quadrupole of the correlation function, the shuffled catalogue performs the best; however, the average bias for both ξ_0 and ξ_2 is also small for the N -spline methods.




Figure 19. Contours showing the 68, 95 and 99.7 per cent confidence level contours on bias and growth factor estimated from ξ_0 and ξ_2 using different methods to determine the radial distribution used for the random catalogues.

The potential systematic error induced by the treatment of the random catalogue appears largest for ξ_2 . To see the effect of $n(z)$ systematics, we find the best-fitting values of bias b and growth rate f when performing a joint fit to ξ_0 and ξ_2 for each mock catalogue, first by using the random catalogue with the true $n(z)$ and then repeating the same analysis using N -splined and shuffled random catalogues. Fig. 19 shows the 1σ , 2σ and 3σ contours on the joint measurement of b and f when using the true $n(z)$ (black), a shuffled $n(z)$ and a 10- (green), 20- (blue), 30- (cyan) node spline. The $n(z)$ systematics push the measured bias towards slightly higher values and the measured growth rate towards lower values but all contours are consistent within 1σ . The results derived with the shuffled catalogue are on average closer to the results derived with a true catalogue than results obtained using a spline fit. Similar fits are performed to the CMASS ξ_ℓ in Reid et al. (2012).

The tests outlined in this section suggest that there is some systematic uncertainty introduced by the method in which random points are assigned redshifts. In general, this causes both an increase in the statistical uncertainty and a systematic bias. The added statistical uncertainty is at most 5 per cent (as given by $\sqrt{1 + 0.33^2}$) of the fiducial uncertainty for $\xi_0(s)$ and 12 per cent of the fiducial uncertainty for $\xi_2(s)$. This added statistical uncertainty is accounted for by measuring the mock correlation functions used for the covariance matrix using the same method as we employ on the data. Fig. 20 shows that the difference between the CMASS ξ_ℓ measurements made using a 20-node spline and shuffled random catalogue is indeed at the level we expect from the mocks.

The systematic bias induced by the treatment of the randoms is negligible for $\xi_0(s)$, but is larger for $\xi_2(s)$. In both cases, using a shuffled random catalogue, on average, produces the least biased result. Therefore, we use the shuffled method to obtain redshifts for the random catalogues we use in $\xi(s)$ calculations.

6.2 Clustering split at $z = 0.52$

One may worry that the clustering at higher redshift may be more prone to systematic errors, given that, all else being equal, higher




Figure 20. Top panel: the red line displays difference between ξ_0 measurements made using the NGC CMASS sample when using a 20-node spline and when randomly selecting redshifts from the galaxy sample ('shuffled') to assign redshifts to the random catalogues. The error bars represent the mean and standard deviation of this difference in the 600 mock galaxy catalogues occupying the NGC footprint. Bottom panel: as above, for ξ_2 .

redshift objects should be fainter. Therefore, we split our CMASS data into two samples, one with $z < 0.52$ and the other with $z > 0.52$. This split is close to the peak of the redshift distribution and represents the redshift at which the CMASS sample transitions from being approximately volume-limited to magnitude-limited [as can be inferred by inspecting $n(z)$ in e.g. Fig. 4]. We also find that the w_{star} weights become much more important at $z > 0.52$, as the mean $i_{\text{fib}2}$ magnitude is 21.01 above $z = 0.52$ and 20.74 below. Thus, we expect differences in $\xi(s)$ measurements to be greatest when split at $z = 0.52$ (and indeed, differences in the measured ξ_ℓ are smaller when we split at e.g. $z = 0.55$).

The resulting $\xi(s)$ are displayed in the top panels of Fig. 21, with open blue circles representing $z < 0.52$ and red triangles representing $z > 0.52$. The amplitudes of the $z > 0.52$ measurements are significantly larger at all scales than the lower redshift ones. This may partially be due to the fact that galaxies at $z > 0.52$ are more luminous, and thus we may expect them to have a higher bias. We




Figure 21. Top panels: the measured redshift-space auto-correlation functions of the combined CMASS data split by redshift into samples with $z > 0.52$ (red) and $z < 0.52$ (blue). Points represent measurements; error bars are determined from the mocks split at the same redshifts. Bottom panels: the difference between the CMASS measurement and the mean of the mock measurements scaled to the best-fitting bias of the respective CMASS samples. The values of s for the $z < 0.52$ sample have been shifted horizontally by $1 h^{-1}\text{Mpc}$ for clarity.

fit these data to our mocks at $25 < s < 150 h^{-1}\text{Mpc}$ and account for the 6.4 per cent change in the linear growth factor between $z = 0.61$ and 0.48 (the mean redshifts of the respective samples). We indeed find a higher bias for the $z > 0.52$ sample, as $\bar{b} = 2.02 \pm 0.04$ for $z > 0.52$ ($\chi^2_{\text{min}} = 22.9$) and $\bar{b} = 1.85 \pm 0.06$ for $z < 0.52$ ($\chi^2_{\text{min}} = 15.1$). The values of \bar{b} (and other parameters we measure throughout this section) are summarized in Table 1 (found in Section 8).

The bottom left-hand panel of Fig. 21 displays the difference between the measurement and the mean of mocks (600 each for $z < 0.52$ and $z > 0.52$), after scaling the mocks to the best-fitting bias. Even after this is done, the amplitudes of the $z > 0.52$ ξ_0 measurements are larger across all scales than the $z < 0.52$ counterparts. This illustrates the level of covariance between s bins in the ξ_ℓ measurements (which allows best-fitting solutions where most of the data are either above or below the model). Both measurements are consistent with the mocks at scales $150 < s < 250 h^{-1}\text{Mpc}$ (14 data points), as $\chi^2 = 7.6$ for $z < 0.52$ and 17.7 for $z > 0.52$ (91 per cent of consistent samples will have $\chi^2 > 7.6$ and 22 per cent will have $\chi^2 > 17.7$).

The right-hand panels of Fig. 21 display the same content as the left panel, but for ξ_2 . The values of $|\xi_2|$ are consistently smaller for $z < 0.52$. We test the significance of this result by fixing the bias at the value determined from ξ_0 and scaling the mean of the ξ_2 mocks to find a best-fitting \bar{f} , via equation (37). For $z > 0.52$, $\bar{f} = 0.75 \pm 0.05$ ($\chi^2 = 14.1$), and for $z < 0.52$, $\bar{f} = 0.59 \pm 0.08$ ($\chi^2 = 6.9$). Accounting for the fact that we expect a 6 per cent decrease in f between $z = 0.61$ and 0.48 , this is a 1.3σ discrepancy and is not surprising given we intentionally split the sample at a redshift where we expected to find the largest differences.

We find that splitting the sample at $z = 0.52$ yields consistent α values when marginalizing over the bias: for $z < 0.52$ we find

$\bar{\alpha} = 1.016 \pm 0.038$ and for $z > 0.52$ we find $\bar{\alpha} = 1.013 \pm 0.021$. Both best-fitting $\bar{\alpha}$ values are smaller than the best fitting value to the combined sample ($\bar{\alpha} = 1.021 \pm 0.019$), implying that the cross-correlation between the two slices contains significant information relevant to $\bar{\alpha}$. The consistency of the results further implies that fits to standard ΛCDM cosmological parameters will yield consistent results. We find similar levels of consistency when splitting the individual NGC and SGC regions at $z = 0.52$ and performing the same tests.

7 CLUSTERING AT THE LARGEST SCALES

To this point we have focused on scales $s < 150 h^{-1}\text{Mpc}$. In this section, we focus on larger scales. Models of the galaxy correlation function in a Universe dominated by dark energy and cold dark matter cross zero at a scale just beyond the BAO peak and asymptote towards zero. This is true even for models with a high level of primordial non-Gaussianity (where the amplitudes around the BAO scale and zero-crossing scale increase). Thus, $\xi_0(s)$ measurements that differ from this behaviour indicate the presence of systematic effects in the galaxy density field, effects not accounted for in the standard paradigm or correlated noise.

Comparing our measured ξ_0 (using the w_{star} and FKP weights) to the mean of that of the mocks at $150 < s < 250 h^{-1}\text{Mpc}$ (14 measurements), we find $\chi^2 = 27.3$. This value is rather large, as only 2 per cent of consistent samples have a greater χ^2 value. We note that if the w_{star} weights are not applied, $\chi^2 = 57.3$ (only 3×10^{-5} per cent of consistent samples would have such a large χ^2). If we do not use the FKP weights, but still use the FKP covariance matrix (to make sure the measurement and not the covariance is driving χ^2), χ^2 decreases to 23.2. This is still large enough that only 5.7 per cent of consistent samples would have a larger χ^2 . This discrepancy is a product of the full sample, over all redshifts (and the associated lower variance), as neither sample when split at $z = 0.52$ returned an abnormally large χ^2 value. The best-fitting bias of the CMASS sample (when using FKP weights) is 2 per cent higher than that of the mocks (1.938 compared to 1.9) and, naively, $\chi^2 \propto b^4$. Performing such a scaling reduces χ^2 to 25.2 – still large enough that we should expect a larger χ^2 value for only 3.3 per cent of consistent samples.

The effect of any systematic associated with the angular mask, e.g. variations in stellar density or errors in the normalization of pair counts, is to add roughly a constant amplitude to ξ_0 . Therefore, we add a constant, A , to ξ_0 of the mocks and find the best-fitting value, fitting at $150 < s < 250 h^{-1}\text{Mpc}$. The χ^2 is minimized at $A = 0.0006$, but is only reduced to 25.8 (23.8 if scaling by the bias). Given that we have reduced the number of degrees of freedom to 13, the probabilities of being consistent remain the same (to within the quoted number of significant digits). That is, A (and, thus, any remaining purely angular systematic) is not detected with any significance.

7.1 Anisotropic clustering and feature at $200 h^{-1}\text{Mpc}$

The inconsistency we find between the measured clustering and the mean of mocks appears driven by an excess at $s \sim 200 h^{-1}\text{Mpc}$. No matter how we split the CMASS sample or change our analysis, the bump-like feature around $200 h^{-1}\text{Mpc}$ remains. It is strange that this feature is nearly as robust in the NGC and SGC samples alone and also in both the samples split at $z = 0.52$. However, we can find mocks with similar large-scale ξ_0 to the CMASS sample. The ξ_0 of three such realizations are plotted along with ξ_0 of




Figure 22. The measured monopole of the correlation function, ξ_0 , multiplied by s^2 for the combined CMASS samples and three mock catalogues with similar levels of clustering at $s \sim 200 h^{-1}$ Mpc.

the combined CMASS sample in Fig. 22. Clearly, cosmic variance allows the possibility of obtaining peaks in ξ_0 at around $200 h^{-1}$ Mpc that are qualitatively similar to those we observe.

We investigate further by examining the clustering of BOSS galaxies as a function of radial, r_{\parallel} , and transverse, r_{\perp} , distances, $\xi(r_{\perp}, r_{\parallel})$. Given the redshift-space separation s and the cosine of the angle to the line-of-sight μ , $r_{\parallel} = \mu s$ and $r_{\perp} = \sqrt{s^2 - r_{\parallel}^2}$. The left-hand panel of Fig. 23 displays the hyperbolic sine, \sinh , of 75 times $\xi(r_{\perp}, r_{\parallel})$. This transformation allows information on all scales to be displayed on a single figure. The effect of RSD is apparent, as amplitudes along the line of sight are clearly decreased relative to those at the same transverse separations.

There is a ring at around $100 h^{-1}$ Mpc, as expected for the BAO feature. The extra information in the r_{\perp}, r_{\parallel} dependence of the

BAO feature is examined in Blanton et al. (in preparation). There also appears to be an excess in a ring around $200 h^{-1}$ Mpc. The middle panel displays the mean $\sinh(75\xi[r_{\perp}, r_{\parallel}])$ in the mocks. At scales less than $100 h^{-1}$ Mpc, this has a similar appearance to the measurements. The right-hand panel displays the difference between the measurements and the mocks, divided by twice the standard deviation of $\xi_{\text{mock}}[r_{\perp}, r_{\parallel}]$. In general, there is excess at the largest scales, and is most pronounced at scales $\sim 200 h^{-1}$ Mpc. The excess at $r \sim 200 h^{-1}$ Mpc is largest for pairs between 10° and 50° from the line of sight. The fact that the excess appears at approximately constant r , rather than r_{\perp} or r_{\parallel} , suggests that the feature is not due to a systematic strictly related to the angular mask (as this would show up at constant r_{\perp}) or the process of obtaining spectroscopic redshifts (as this would appear at constant r_{\parallel}). Finally, the feature shows up at nearly identical physical scales when the sample is split at $z = 0.52$ (see Fig. 21), further implying that the feature is not associated with a fixed angular scale.

To further assess the significance of the feature at $\sim 200 h^{-1}$ Mpc, we design a statistic that reflects the degree to which there is an excess (or decrement) of signal at constant separation. Thus, we determine the difference between the measured correlation function and the mean of the mock correlation function divided by the standard deviation determined from the mocks, $\sigma(r_{\parallel}, r_{\perp})$, averaged over measurements within a constant separation bin. Thus, we measure

$$t(r, \Delta r) = \left| \frac{\sum \Theta(r_{\parallel}, r_{\perp})(\xi[r_{\parallel}, r_{\perp}] - \bar{\xi}_{\text{mock}}[r_{\parallel}, r_{\perp}]) / \sigma(r_{\parallel}, r_{\perp})}{\sum \Theta(r_{\parallel}, r_{\perp})} \right|, \quad (42)$$

where $\Theta(r_{\parallel}, r_{\perp}) = 1$ if $(r - \Delta r)^2 < r_{\parallel}^2 + r_{\perp}^2 < (r + \Delta r)^2$ and 0 otherwise. In essence, this statistic is just a binned version of the information plotted in the right-hand panel of Fig. 23.

Fig. 24 displays the results of performing this test on the measurements, setting $\Delta r = 10 h^{-1}$ Mpc. The most discrepant results are at $r = 30 h^{-1}$ Mpc, although this is likely suggestive of a different preferred cosmology from the one used by the mocks. At $r = 204 h^{-1}$ Mpc, we find a peak with amplitude 0.72. If we perform the same test on each of the 600 mocks, we find a peak with amplitude greater than 0.72 at $r > 120 h^{-1}$ Mpc in 84 of the mocks, which suggests that the observed size of the peak is common in the mocks. If we demand that the peak be at least as wide as we find in the CMASS data by searching for features where $t > 0.5$ over




Figure 23. Left panel: the measured correlation function, $\xi(r_{\perp}, r_{\parallel})$, plotted a function of the radial, r_{\parallel} , and transverse, r_{\perp} , distance for the CMASS sample. To compare all scales, we plot the hyperbolic sine of $(75\xi[r_{\perp}, r_{\parallel}])$. Middle panel: the mean $\sinh(75\xi[r_{\perp}, r_{\parallel}])$ of 600 mocks masked to simulate the NGC footprint of the CMASS data set. Right panel: $(\xi_{\text{CMASS}}[r_{\perp}, r_{\parallel}] - \bar{\xi}_{\text{mock}}[r_{\perp}, r_{\parallel}]) / 2\sigma$, where σ is the standard deviation of $\xi_{\text{mock}}[r_{\perp}, r_{\parallel}]$ (and \sinh scaling is no longer used; we divide by 2σ , rather than 1, for clarity).




Figure 24. The average value of $(\xi[r_{\parallel}, r_{\perp}] - \bar{\xi}_{\text{mock}}[r_{\parallel}, r_{\perp}]) / \sigma(r_{\perp})$ calculated in rings with $\Delta r = 10 h^{-1}$ Mpc and $(r - \Delta r)^2 < r_{\parallel}^2 + r_{\perp}^2 < (r + \Delta r)^2$.

$25 h^{-1}$ Mpc and $t_{\text{max}} > 0.72$ (as is the case for the CMASS data), only 29 mocks (4.8 per cent) are selected. This implies that the combination of the size and width of the feature centred at $200 h^{-1}$ Mpc is driving the large χ^2 values of the ξ_0 measurements with $s > 150 h^{-1}$ Mpc. The local maximum at $r = 100 h^{-1}$ Mpc, $t_{\text{max}} = 0.45$ and $t > 0.3$ over $25 h^{-1}$ Mpc is not significant, as 45 per cent (267) of the mocks have a feature at least as large and wide centred at $r > 80 h^{-1}$ Mpc.

Fig. 22 displays ξ_0 of three realizations where $t_{\text{max}} > 0.65$ and is close to $200 h^{-1}$ Mpc, showing that these cases are qualitatively similar to the CMASS ξ_0 . The SDSS I/II LRG $\xi_0(s)$ measurements have larger than expected amplitudes at large scales. However, a systematic study of the SDSS I/II LRG sample, similar to our own, has not been published and Xu et al. (2012) still suggest that its large-scale clustering amplitudes are within 2σ of those expected. The final BOSS data set will have three times more data than the DR9 sample, and thus future data releases will confirm if the feature at $200 h^{-1}$ Mpc is simply noise, a yet-to-be-determined systematic, or a real feature in the clustering of galaxies. We do not know of any model that predicts such a feature, as e.g. the predicted $\xi(s)$ given a strongly non-Gaussian primordial power spectrum display much smoother variation. This is studied further in Ross et al. (in preparation).

7.2 Power spectrum measurements

We investigate results using the spherically averaged power spectrum, $P(k)$, in order to isolate the large-scale density modes [i.e. low k ; the $P(k)$ measurements are less covariant between k bins] and as a consistency check on results derived from ξ_0 measurements (which should contain the same information). In Fig. 25, we display the measured $P(k)$, calculated as described in Section 3.1. The open circles show the measurement without using the weights accounting for stellar density, while the solid circles display the measurement when these w_{star} weights are applied. These weights only cause a significant difference for the smallest k (largest scales), unlike the situation for $\xi_0(s)$, where the difference was nearly independent of scale. The inset displays the same information, divided by the best-fitting smooth model found in Anderson et al. (2012).

The open and solid circles are indistinguishable from each other. Clearly, the w_{star} weights do not affect the $P(k)$ measurements at the scales related to the BAO feature.

We scale the mock $P(k)$ to find best-fitting \tilde{b} values in the same manner as performed throughout for $\xi(s)$. We determine the covariance matrix of $\ln[P(k)k^3]$, from the 600 mock catalogues and minimize χ^2 of $\ln[P(k)k^3]$. We use the logarithmic scaling to account for the fact that we expect the cosmic variance contribution to the $P(k)$ uncertainty to be proportional to $P^2(k)$ (see e.g. Feldman et al. 1994). This scaling does not significantly alter the best-fitting values we determine, but it does result in significantly smaller χ_{min}^2 values. We find $\tilde{b} = 1.983 \pm 0.035$ fitting $k < 0.05 h \text{Mpc}^{-1}$ with $\chi_{\text{min}}^2 = 19.2$ (11 dof) when the w_{star} weights are applied and $\tilde{b} = 2.001 \pm 0.035$ with $\chi_{\text{min}}^2 = 38.8$ when no weights are applied. The χ_{min}^2 increases dramatically, by a factor greater than 2, without the weights. This shows how dramatic an effect the weights have – only 0.006 per cent of consistent samples would have $\chi^2 > 38.8$, while 5.8 per cent would have $\chi^2 > 19.2$.

The effect of the w_{star} weights on the $P(k)$ measurement shows a significant scale dependence, unlike for $\xi(s)$, where the change was nearly constant. We can assume that any unaccounted-for systematic will have the same k dependence as the w_{star} weights and determine if adding a factor $A[P(k) - P_{\text{weight}}(k)]$ improves our \tilde{b} fit. We find that $\chi_{\text{min}}^2 = 13.96$ at $\tilde{b} = 1.974 \pm 0.035$ and $A = -0.41$. That is, a 41 per cent stronger systematic correction decreases χ^2 by 5.2. These results strongly suggest that proper treatment of the weights is vital in any attempt to obtain robust measurements that use $P(k)$ measurements at low k , e.g. measurements of primordial non-Gaussianity or the scale of matter radiation equality [from the overall peak in $P(k)$]. The degeneracy between the systematic correction and the constraints on primordial non-Gaussianity one can obtain is studied further in Ross et al. (in preparation).

The best-fitting bias obtained from the $P(k)$ measurement is nearly 1σ larger than that we obtain when $\xi(s)$ is fitted at scales $> 25 h^{-1}$ Mpc. Measurements of absolute bias values are notoriously difficult, and the recovered value is often driven by the minimum scale that is fitted (since this measurement has the least uncertainty; see e.g. Ross et al. 2011a). Indeed, it is difficult even using dark matter simulations, as Manera, Sheth & Scoccimarro (2010) found systematic differences (at a level similar to that we find here) in large-scale bias measurements when comparing results obtained from matter-halo cross-power spectrum and halo auto-correlation measurements. Thus, obtaining robust absolute bias measurements is a generic problem rather than an issue specific to the BOSS DR9 galaxy sample.

Fig. 26 displays $P(k)$ measured for the NGC (open triangles) and SGC (open squares) samples and the average of the mock $P(k)$ for each region separately. The windows of the respective regions create a large difference in the expected $P(k)$, as the shape of the mock $P(k)$ is considerably different at small k . Clearly, a direct comparison of the respective $P(k)$ is not appropriate. Scaling the mock $P(k)$ for the respective regions, we find $\tilde{b} = 1.991 \pm 0.039$ for the NGC with $\chi_{\text{min}}^2 = 18.5$, and $\tilde{b} = 1.93 \pm 0.07$ for the SGC ($\chi_{\text{min}}^2 = 11.9$). The difference between the respective \tilde{b} is opposite what we found from the $\xi_0(s)$ measurements, where the bias of the SGC sample was significantly larger. This suggests that smaller scale modes ($k > 0.05$) have significant influence on correlation function measurements of the SGC sample at $s > 25 h^{-1}$ Mpc. Unlike for $\xi(s)$, the bias values determined using $P(k)$ are consistent between the two regions.

Fig. 27 displays the ratio between the measured $P(k)$ and the mean mock $P(k)$ when scaling k linearly. The measured ξ_0 has a peak




Figure 25. The measured spherically averaged (in redshift-space) power spectrum, $P(k)$, of the full DR9 CMASS sample with only the fiducial angular weight applied (open circles) and with the linear-fit weights for stellar density, w_{star} , applied (solid circles). The average of the mock $P(k)$ is displayed with a dashed line. The best-fitting smooth model $P(k)$ determined by Anderson et al. (2012) is plotted with a solid line. The inset displays the same information, divided by the smooth fit, with the solid line displaying the best-fitting model including BAO.

at $204 h^{-1} \text{Mpc}$, which suggests we should find a periodic feature in $P(k)$ with wavelength $\sim 0.03 h \text{Mpc}^{-1}$ in k . We have plotted dotted lines with a spacing of $0.029 h \text{Mpc}^{-1}$, with the first at $k = 0.027 h \text{Mpc}^{-1}$. There appears to be local minima at k values near each dotted line. This appears most significant at $k \sim 0.03 h \text{Mpc}^{-1}$. We caution that the possibility this feature is due to an undiscovered systematic may need to be considered when performing analysis of the shape of the CMASS power spectrum.

8 CONCLUSIONS

We have investigated potential systematic effects on the three-dimensional clustering of the DR9 sample of BOSS galaxies and tested the robustness of the results when we split the data into the regions we expect may be the most different for observational (Northern/Southern Galactic Caps) and physical (split at $z = 0.52$) reasons. Our main findings are summarized as follows.

(i) Redshift failures occur at preferred locations on spectroscopic tiles (see Fig. 3), but the nearest redshifts within the same sector (which therefore share the same observing conditions) have an $n(z)$ like that of the overall sample (see Fig. 4). We therefore account for redshift failures by up-weighting the nearest targeted object within

the same sector. We find that this approach has a minor effect on the measured $\xi(s)$ (see the red lines in Fig. 5).

(ii) We account for target objects that lack spectra due to fibre collisions by up-weighting the nearest targeted object. At large scales ($s > 10 h^{-1} \text{Mpc}$), this should be equivalent to the more sophisticated method proposed in Guo et al. (2011). This scheme increases the $\xi(s)$ measurements in a manner consistent with a small increase in the galaxy bias (see the blue lines in Fig. 5).

(iii) The overall number densities of observed galaxies in the Southern Galactic hemisphere are higher for both the LOWZ (8 per cent) and CMASS (3.2 per cent) samples. If we apply the offsets in colour found by Schlafly & Finkbeiner (2011) to the selection of BOSS galaxies, the number densities become consistent within 2 per cent for LOWZ and 0.2 per cent for CMASS. After applying the Schlafly & Finkbeiner (2011) offsets, the $n(z)$ values are discrepant at a level that we find in 10 per cent of mock samples; that is, the difference is not significant.

(iv) The measured clustering in the NGC and SGC generally agrees to within 2σ , depending on the specific test that is performed. For $\xi(s)$, the bias disagrees at 1.9σ , but for $P(k)$, the discrepancy is only 0.3σ . Measurements of the amplitude of ξ_2 , \bar{f} , are consistent to within 1σ .




Figure 26. The $P(k)$ measurements for the NGC (open red triangles) and SGC samples (open blue squares), with the mean of the respective mock samples displayed with a solid line. The difference between the two lines illustrates the effect of the different windows of the NGC and SGC on the expected $P(k)$.




Figure 27. Solid lines display the ratio between the CMASS $P(k)$ measurement to the average of the mock $P(k)$. The vertical dotted lines denote a spacing of $0.029 h \text{ Mpc}^{-1}$ in k , with the first line at $k = 0.027 h \text{ Mpc}^{-1}$.

(v) The measured clustering in the NGC and SGC disagrees most at around the BAO scale; we find that this causes a difference in stretch parameter, $\tilde{\alpha}$ [which in this study encodes information on both the BAO scale and the shape of $\xi(s)$] of 2.5σ . We find no treatment of the data (e.g. alternative weighting) that significantly alters the level of discrepancy.

(vi) We weight CMASS galaxies based on linear relationships between the expected number density of galaxies as a function of

their $i_{\text{fib}2}$ magnitude and the local stellar density (w_{star} weights). We find no evidence that similar weights are necessary for (the brighter) LOWZ galaxies. By applying the method used to determine the weights to mocks (which have no need for correction) we find that these weights produce no bias on the mean measured $\xi(s)$, whereas more aggressive weighting schemes may (see Appendix A). Application of these weights reduces the measured cross-correlations of the CMASS galaxies with stellar density, Galactic extinction, seeing, sky background and airmass to the level we expect to find randomly (see Fig. 13). These weights account for redshift dependence through the $i_{\text{fib}2}$ relationship, but may no longer be sufficient when the sample is split by colour.

(vii) Applying the w_{star} weights produces a 0.7σ shift in the measured value of $\tilde{\alpha}$, most of which we believe is due to the change in shape of the correlation function. Applying the weights reduces χ^2 from 34.2 to 22.7 when $\xi_0(s)$ measurements are fitted in the range $25 < s < 150 h^{-1} \text{ Mpc}$ (18 data points). The w_{star} weights have little effect on ξ_2 measurements, as the best-fitting \tilde{f} does not change and the χ^2 is reduced only from 12.7 to 11.8 when fitting scales $25 < s < 150 h^{-1} \text{ Mpc}$.

(viii) We use the mocks to determine the least-biased way to simulate the radial selection function of CMASS galaxies, in order to produce a random catalogue. Randomly selecting the redshift of a galaxy in the sample (‘shuffled’) produces a smaller bias than performing a spline fit to the measured $n(z)$ (‘spline’). In all cases, the bias is negligibly small for ξ_0 , but is as high as 50 per cent of the statistical uncertainty for ξ_2 . We therefore advocate using shuffled random catalogues and note that any constraints obtained using ξ_2 measurements should take this bias into account. We find that the differences between ξ_0 and ξ_2 measurements we obtain using CMASS galaxies using the spline and shuffled methods are consistent with the level found in the mocks.

(ix) The $\xi_0(s)$ measurements, when split at $z = 0.52$, yield bias and α values that are consistent within 1σ , when fitted in the range $25 < s < 150 h^{-1} \text{ Mpc}$. The ξ_2 measurements are somewhat discrepant, as the best-fitting \tilde{f} values differ by 1.7σ .

(x) The $\xi_0(s)$ measurements, at scales $s > 150 h^{-1} \text{ Mpc}$, are inconsistent at a greater than 94 per cent level with the expected clustering. Allowing for a constant offset in large-scale clustering (as angular systematics tend to contribute) produces no improvement.

(xi) The inconsistency at large scales is dominated by a peak in $\xi_0(s)$ at $s \sim 200 h^{-1} \text{ Mpc}$. This feature appears in a ring when we measure $\xi(r_{\perp}, r_{\parallel})$, implying that it is not due to a systematic solely to either the mask or target catalogue (as would show up in r_{\perp}) or the process of obtaining spectroscopic redshifts (as would show up in r_{\parallel}). Similar features are found that are at least as significant in 4.8 per cent of the mocks that we test.

(xii) The w_{star} weights have a dramatic effect on $P(k)$ measurements and the effect depends strongly on k . The χ^2 decreases from 38.8 to 19.2 before and after w_{star} weights are applied to $P(k)$ measurements fitted at $k < 0.05$ (12 data points). We find that χ^2 reduces to 14.0 when allowing a further correction of the form $A(P[k]_{\text{noweight}} - P[k]_{\text{weight}})$. This implies that one must be careful when obtaining any information that depends on measurements of $P(k)$ at $k < 0.01$ and this issue is studied further in Ross et al. (in preparation).

The fundamental conclusions of this work are that, for BOSS DR9 CMASS galaxies, we recommend the application of weights to account for fibre collisions, redshift failures and the systematic effect of stars, and we find no further systematic dependencies. We therefore expect Anderson et al. (2012), Sanchez et al. (2012), Reid

et al. (2012) and Samushia et al. (in preparation) (all of which apply the same weights) to obtain robust cosmological constraints using the clustering of BOSS DR9 galaxies.

Anderson et al. (2012) find that the measured BAO position does not depend on the application of the w_{star} weights – the position changes by less than 0.1σ . Therefore, our results do not suggest that there is a significant unaccounted-for systematic error in previous BAO measurements. However, our study does suggest that any previous finding of a large-scale excess in clustering measurements may have been due to systematic errors, similar to those we correct for, and requires extra scrutiny. Indeed, Ross et al. (2011b) show that much of the excess in large-scale power presented in Thomas, Abdalla & Lahav (2011) (who used SDSS-II imaging data) is removed when the systematic effect of stars on the projected density field is accounted for.

The SDSS-III BOSS DR9 sample of galaxies represents only one-third of the final BOSS CMASS sample and one quarter of the final LOWZ sample. Further data sets will allow potential systematics to be tested to an even greater extent and reveal whether the feature at $200 h^{-1}$ Mpc is noise.

ACKNOWLEDGMENTS

AJR is grateful to the UK Science and Technology Facilities Council for financial support through the grant ST/I001204/1. WJP is grateful for support from the UK Science and Technology Facilities Research Council, the Leverhulme Trust and the European Research Council. MAS acknowledges the support of NSF grant AST-0707266. MECS was supported by the NSF under Award No. AST-0901965.

We thank the anonymous referee for comments that helped improve the quality of this work.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation and the U.S. Department of Energy Office of Science. The SDSS-III website is <http://www.sdss3.org/>.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, University of Cambridge, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington and Yale University.

REFERENCES

Abazajian K. et al., 2004, *AJ*, 128, 502
 Adelman-McCarthy J. K. et al., 2008, *ApJS*, 175, 297
 Aihara H. et al., 2011, *ApJS*, 193, 29
 Anderson L. et al., 2012, *MNRAS*, preprint (arXiv:1203.6594)
 Blake C. et al., 2010, *MNRAS*, 406, 803
 Blake C. et al., 2011, *MNRAS*, 1598

Blanton M. R., Lin H., Lupton R. H., Maley F. M., Young N., Zehavi I., Loveday J., 2003, *AJ*, 125, 2276
 Cole S. et al., 2005, *MNRAS*, 362, 505
 Colless M. et al., 2003, preprint (astro-ph/0306581)
 Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
 Eisenstein D. J. et al., 2011, *AJ*, 142, 72
 Feldman H. A., Kaiser N., Peacock J. A., 1994, *ApJ*, 426, 23
 Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, *AJ*, 111, 1748
 Gómez P. L. et al., 2003, *ApJ*, 584, 210
 Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
 Gunn J. E. et al., 1998, *AJ*, 116, 3040
 Gunn J. E. et al., 2006, *AJ*, 131, 2332
 Guo H., Zehavi I., Zheng Z., 2011, preprint (arXiv:1111.6598)
 Hamilton A. J. S., 1992, *ApJ*, 385, L5
 Hamilton A. J. S., 1993, *ApJ*, 417, 19
 Hamilton A. J. S., Tegmark M., 2004, *MNRAS*, 349, 115
 Ho S. et al., 2012, preprint (arXiv:1201.2137)
 Kaiser N., 1987, *MNRAS*, 227, 1
 Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
 Larson D. et al., 2011, *ApJS*, 192, 16
 Manera M., Sheth R. K., Scoccimarro R., 2010, *MNRAS*, 402, 589
 Manera M. et al., 2012, *MNRAS*, preprint (arXiv:1203.6609)
 Masters K. L. et al., 2011, *MNRAS*, 418, 1055
 Montesano F., Sánchez A. G., Phleps S., 2012, *MNRAS*, 421, 2656
 Norberg P. et al., 2002, *MNRAS*, 332, 827
 Nuza S. E. et al., 2012, *MNRAS*, preprint (arXiv:1202.6057)
 Padmanabhan N. et al., 2008, *ApJ*, 674, 1217
 Percival W. J. et al., 2010, *MNRAS*, 401, 2148
 Reid B. A., White M., 2011, *MNRAS*, 417, 1913
 Reid B. A. et al., 2010, *MNRAS*, 404, 60
 Reid B. et al., 2012, *MNRAS*, preprint (arXiv:1203.6641)
 Ross A. J. et al., 2011b, *MNRAS*, 417, 1350
 Ross A. J., Tojeiro R., Percival W. J., 2011a, *MNRAS*, 413, 2078
 Samushia L., Percival W. J., Raccanelli A., 2012, *MNRAS*, 420, 2102
 Sanchez A. et al., 2012, *MNRAS*, preprint (arXiv:1203.6616)
 Schlafly E. F., Finkbeiner D. P., 2011, *ApJ*, 737, 103
 Schlafly E. F., Finkbeiner D. P., Schlegel D. J., Jurić M., Ivezić Ž., Gibson R. R., Knapp G. R., Weaver B. A., 2010, *ApJ*, 725, 1175
 Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525
 Scoccimarro R., Sheth R. K., 2002, *MNRAS*, 329, 629
 Stoughton C. et al., 2002, *AJ*, 123, 485
 Strateva I. et al., 2001, *AJ*, 122, 1861
 Swanson M. E. C., Tegmark M., Blanton M., Zehavi I., 2008a, *MNRAS*, 385, 1635
 Swanson M. E. C., Tegmark M., Hamilton A. J. S., Hill J. C., 2008b, *MNRAS*, 387, 1391
 Sylos Labini F., Vasilyev N. L., Baryshev Y. V., López-Corredoira M., 2009, *A&A*, 505, 981
 Tegmark M., Peebles P. J. E., 1998, *ApJ*, 500, L79
 Tegmark M. et al., 2004, *Phys. Rev. D*, 69, 103501
 Thomas S. A., Abdalla F. B., Lahav O., 2011, *Phys. Rev. Lett.*, 106, 241301
 Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlöber S., 2010, *ApJ*, 724, 878
 Tojeiro R., Percival W. J., 2010, *MNRAS*, 405, 2534
 Tojeiro R. et al., 2012a, *MNRAS*, preprint (arXiv:1203.6565)
 Tojeiro R. et al., 2012b, preprint (arXiv:1202.6241)
 Wake D. A. et al., 2008, *MNRAS*, 387, 1045
 White M. et al., 2011, *ApJ*, 728, 126
 Xu X., Padmanabhan N., Eisenstein D. J., Mehta K. T., Cuesta A. J., 2012, preprint (arXiv:1202.0091)
 Yahata K., Yonehara A., Suto Y., Turner E. L., Broadhurst T., Finkbeiner D. P., 2007, *PASJ*, 59, 205
 York D. G. et al., 2000, *AJ*, 120, 1579
 Zehavi I. et al., 2011, *ApJ*, 736, 59
 Zheng Z., Coil A. L., Zehavi I., 2007, *ApJ*, 667, 760

APPENDIX A: ANGULAR WEIGHTING SCHEMES

We considered three different weighting schemes in order to account for the systematic correlations found in Section 5. These included:

(i) ‘*Iterative weights*’ which we denote w_{it} . This technique was applied in Ross et al. (2011b). It assumes that the effects of each systematic are separable, and proceeds by starting with one systematic and setting the weight in every HEALPIX pixel equal to the inverse of the quantity plotted in black in Fig. 11. One then moves on to the next systematic and recalculates the relationship between the number density of galaxies and the systematic, and then multiplies the weights by the inverse of the relationship. If the effects are indeed separable, the $n_{gal}(sys)$ relationships should all remain consistent with unity after all of the weights have been calculated.

To determine w_{it} , we proceed in the order stellar density, Galactic extinction, airmass, seeing and sky background. If each is truly separable, the order should not matter, and we do find negligible differences for any permutation of the order we have tested. The residual relationships between the galaxy number density and the potential systematic, when weighting by the full w_{it} , are displayed with magenta lines in Fig. A1. In every case, the relationship is almost fully removed. This implies that the weighting is too aggressive, as we should actually expect variations consistent with the size of the error bars in Fig. A1.

We can test the extent to which the w_{it} weights may remove true power from clustering measurements by applying weights to each mock sample (which of course contain no imaging systematics) following the methods we apply to the data. The black triangles in Fig. A2 display the average difference between the fiducial ξ measurements and when the full iterative weights, w_{it} , are calculated and applied to each mock. For the monopole, this decreases the expected result by about half the statistical uncertainty (displayed with the black dotted lines). There is also a non-zero bias for the ξ_2 measurements (top panel), but the difference is insignificant compared to the statistical uncertainty.

(ii) ‘*Linear-fit MCMC weights*’, which we denote w_{MCMC} . These weights are calculated by using a Monte Carlo Markov chain (MCMC) to simultaneously find the linear coefficients that best describe the total $n_g(n_{sys})$ relationships.

The w_{MCMC} weights are determined by finding the best-fitting solution to




Figure A2. The average difference between the fiducial redshift-space correlation function of the mocks, ξ , and that using weights for each mock using the full iterative method (black triangles w_{it}), and that using weights for each mock using only a linear fit to the relationship with stellar density (red circles, w_{star}). Error bars represent the standard deviation of the difference. The black dotted lines display the variance on ξ found in the mocks. We note that the mocks do not require weights – a deviation from zero implies that a bias is imparted by the weight scheme.

$$n_{gal}/n_{ran} = K + An_{star} + BA_r + Csee_i + Dsky_i + Eair \quad (A1)$$

where K , A , B , C , D and E are the coefficients we fit for and see is the seeing, sky is the sky background and air is the airmass. This is solved efficiently using an MCMC, as coefficients can be applied to the HEALPIX map simultaneously (thereby accounting for any covariances between the potential systematics). The value of w_{MCMC} is then the inverse of the best-fitting relationship. The residual relationships after applying the w_{MCMC} weights are displayed in blue in Fig. A1. These weights allow more variation than the w_{it} weights. However, the sum of $(\xi_{p,x}(r_{eff})^2/\xi_p(r_{eff}))$ over all five potential systematics for CMASS galaxies with the w_{MCMC} weights, displayed in blue in Fig. 13, is substantially smaller than we expect from the




Figure A1. Same as Fig. 11, except we now plot the residual relationships after applying iterative weights (magenta; w_{it}), the residual relationships after using an MCMC to simultaneously fit linear relationships in order to determine the weights (blue; w_{MCMC}) and the residual when the weights are split as a function of the fibre magnitude, but calculated only based on stellar density (red; w_{star}).




Figure A3. The relationship between the projected number density of CMASS galaxies versus the offset in Δd_{\perp} , as determined from the Schlafly & Finkbeiner (2011) offsets determined for SDSS imaging runs. The error bars reflect the variation found in mock galaxy catalogues occupying the same area. The solid line is the expected relationship, based on the scalings found in Section 4.

mocks (black error bars in Fig. 13). This result implies that the w_{MCMC} weights may also overcorrect the CMASS galaxy density field and remove true fluctuations.

(iii) ‘*Linear-fit stellar density weights*’, which we denote w_{star} . These weights are calculated by performing linear fits to the dependence of galaxy density with stellar density and $i_{\text{fib}2}$ magnitude, as described in Section 5.2 and adopted for the analysis we performed throughout this paper.

The effect of applying only weights for stars, fitted to the linear relationship between n_{gal} and n_{star} , on the mocks is displayed in red circles in Fig. A2. The difference is consistent with zero for both

ξ_0 and ξ_2 . This suggests that the w_{star} weights do not oversmooth the galaxy density field. Furthermore, the sum of the five potential systematic contributions is consistent with the mean mock sum when the w_{star} weights (as shown in Fig. 13) are applied to the CMASS data. These results are in contrast to our previous tests that suggested that both the w_{it} and w_{MCMC} remove true power. Therefore, we believe that the w_{star} weights are the most appropriate to apply to the CMASS sample.

Finally, we considered the fact that Schlafly & Finkbeiner (2011) determined colour offsets for every SDSS imaging run available at the time of their study. This included 95 per cent of the CMASS data in the NGC and 63 per cent in the SGC. Restricting our analysis to these regions, we can measure the relationship between the projected number density of CMASS galaxies and the offset in d_{\perp} , as given by Schlafly & Finkbeiner (2011). This is displayed in Fig. A3, where we have applied the w_{star} weights to the CMASS sample and the error bars reflect the variation found in mock samples over the same footprint. The solid line displays the relationship expected from scalings found in Section 4, determined to be $n_{\text{gal}}/n_{\text{ran}} = 1 + 4.217\Delta d_{\perp}$, and appears consistent with what we measure. We measured $\xi_0(s)$ in this region, applying a weight to account for the predicted scaling with Δd_{\perp} , and found negligible differences between the recovered measurements and those using the full footprint, with separate NGC and SGC selection functions, and the w_{star} weights (our recommended procedure). Within a single hemisphere, the effect of the offsets is similar to that of seeing – there is a systematic relationship, but the pattern of the imaging runs is effectively random and therefore the relationship does not impart spurious power at scales relevant to our analysis. Ho et al. (2012) reached a similar conclusion when analysing angular power spectra of the BOSS imaging data. The exception is when one considers the NGC and SGC together, as the mean offset between the two regions is large enough to produce a significant offset in the number densities of the two regions, and thus imparts a systematic error in the large-scale clustering if separate NGC and SGC selection functions are not applied.

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