

The clustering of intermediate-redshift quasars as measured by the Baryon Oscillation Spectroscopic Survey

Martin White,^{1,2*} Adam D. Myers,^{3,4} Nicholas P. Ross,¹ David J. Schlegel,¹ Joseph F. Hennawi,⁴ Yue Shen,⁵ Ian McGreer,⁶ Michael A. Strauss,⁷ Adam S. Bolton,⁸ Jo Bovy,⁹ X. Fan,⁹ Jordi Miralda-Escude,^{10,11} N. Palanque-Delabrouille,¹² I. Paris,¹³ P. Petitjean,¹³ D. P. Schneider,^{14,15} M. Viel,¹⁶ David H. Weinberg,¹⁷ Ch. Yeche,¹² I. Zehavi,¹⁸ K. Pan,¹⁹ S. Snedden,¹⁹ D. Bizyaev,²⁰ H. Brewington,¹⁹ J. Brinkmann,¹⁹ V. Malanushenko,¹⁹ E. Malanushenko,¹⁹ D. Oravetz,¹⁹ A. Simmons,¹⁹ A. Sheldon¹⁹ and Benjamin A. Weaver²⁰

¹Physics Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

²Departments of Physics and Astronomy, 601 Campbell Hall, University of California Berkeley, CA 94720, USA

³Department of Physics and Astronomy, University of Wyoming, Laramie, WY 82071, USA

⁴Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

⁵Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, MS-51, Cambridge, MA 02138, USA

⁶Steward Observatory, University of Arizona, 933 N. Cherry Avenue, Tucson, AZ 85721, USA

⁷Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁸Department of Physics and Astronomy, The University of Utah, 115 S 1400 E, Salt Lake City, UT 84112, USA

⁹Institute for Advanced Study, Einstein Drive, Princeton, NJ 08540, USA

¹⁰Institució Catalana de Recerca i Estudis Avançats, 08010 Barcelona, Spain

¹¹Institut de Ciències del Cosmos (IEEC/UB), 08034 Barcelona, Spain

¹²CEA, Centre de Saclay, IRFU, 91191 Gif-sur-Yvette, France

¹³Université Paris 6 et CNRS, Institut d'Astrophysique de Paris, 98bis Blvd. Arago, 75014 Paris, France

¹⁴Department of Astronomy & Astrophysics, Pennsylvania State University, University Park, PA 16802, USA

¹⁵Institute for Gravitation & the Cosmos, Pennsylvania State University, University Park, PA 16802, USA

¹⁶INAF-Osservatorio Astronomico di Trieste, via G. B. Tiepolo 11, I-34131 Trieste, Italy

¹⁷Department of Astronomy and CCAPP, Ohio State University, Columbus, OH 43210, USA

¹⁸Department of Astronomy, Case Western Reserve University, Cleveland, OH 44106, USA

¹⁹Apache Point Observatory, PO Box 59, Sunspot, NM 88349-0059, USA

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

Accepted 2012 May 4. Received 2012 May 2; in original form 2012 March 23

ABSTRACT

We measure the quasar two-point correlation function over the redshift range $2.2 < z < 2.8$ using data from the Baryon Oscillation Spectroscopic Survey. We use a homogeneous subset of the data consisting of 27 129 quasars with spectroscopic redshifts – by far the largest such sample used for clustering measurements at these redshifts to date. The sample covers 3600 deg^2 , corresponding to a comoving volume of $9.7 (h^{-1} \text{ Gpc})^3$ assuming a fiducial Λ cold dark matter cosmology, and it has a median absolute i -band magnitude of -26 , k -corrected to $z = 2$. After accounting for redshift errors we find that the redshift-space correlation function is fitted well by a power law of slope -2 and amplitude $s_0 = (9.7 \pm 0.5) h^{-1} \text{ Mpc}$ over the range $3 < s < 25 h^{-1} \text{ Mpc}$. The projected correlation function, which integrates out the effects of peculiar velocities and redshift errors, is fitted well by a power law of slope -1 and $r_0 = (8.4 \pm 0.6) h^{-1} \text{ Mpc}$ over the range $4 < R < 16 h^{-1} \text{ Mpc}$. There is no evidence for strong luminosity or redshift dependence to the clustering amplitude, in part because of the limited dynamic range

*E-mail: mwhite@berkeley.edu

in our sample. Our results are consistent with, but more precise than, previous measurements at similar redshifts. Our measurement of the quasar clustering amplitude implies a bias factor of $b \simeq 3.5$ for our quasar sample. We compare the data to models to constrain the manner in which quasars occupy dark matter haloes at $z \sim 2.4$ and infer that such quasars inhabit haloes with a characteristic mass of $\langle M \rangle \simeq 10^{12} h^{-1} M_{\odot}$ with a duty cycle for the quasar activity of 1 per cent.

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

1 INTRODUCTION

Quasars are among the most luminous astrophysical objects, and are believed to be powered by accretion on to supermassive black holes (e.g. Salpeter 1964; Lynden-Bell 1969). They have become a key element in our current paradigm of galaxy evolution – essentially all spheroidal systems at present harbour massive black holes (Kormendy & Richstone 1995), the masses of which are correlated with many properties of their host systems. The emerging picture is that quasar activity and star formation are inextricably linked (e.g. Nandra et al. 2007; Silverman et al. 2008) in galaxies that contain a massive bulge (and thus a massive black hole) and a gas reservoir. The galaxy initially forms in a gas-rich, rotation-dominated system. Once the dark matter halo grows to a critical scale some event – most likely a major merger (Carlberg 1990; Haiman & Loeb 1998; Cattaneo, Haehnelt & Rees 1999; Kauffmann & Haehnelt 2000; Springel, Di Matteo & Hernquist 2005; Hopkins et al. 2006) or instability in a cold-stream fed disc (Ciotti & Ostriker 1997, 2001; Di Matteo et al. 2012) – triggers a period of rapid, obscured star formation and the generation of a stellar bulge. After some time the quasar becomes visible, and soon after the star formation is quenched on a short time-scale, perhaps via radiative or mechanical feedback from the central engine (e.g. Shankar 2009; Alexander & Hickox 2012; Natarajan 2012).

The clustering of quasars as a function of redshift and luminosity provides useful constraints on our understanding of galaxy evolution. The large-scale clustering amplitude increases with the mass of the dark matter haloes hosting the quasars. Comparison of the abundance of such haloes to that of quasars can provide constraints on the duty cycle and degree of scatter in the observable halo mass relation (Cole & Kaiser 1989; Haiman & Hui 2001; Martini & Weinberg 2001; White, Martini & Cohn 2008; Shankar, Weinberg & Shen 2010). However, quasars are extremely rare, so very large surveys are necessary to suppress the shot noise from Poisson fluctuations. Samples of quasars have only recently included enough objects to study their clustering with some precision (Porciani, Magliocchetti & Norberg 2004; Croom et al. 2005; Hennawi et al. 2006a; Myers et al. 2006, 2007a,b; Porciani & Norberg 2006; da Angela et al. 2008; Padmanabhan et al. 2009; Ross et al. 2009; Shen et al. 2009).

Naively, measuring the clustering of quasars between redshift 2 and 3 should be a simple task, as this is where the comoving number density of luminous quasars seemingly peaks (Weedman 1986; Hartwick & Schade 1990; Croom et al. 2005; Richards et al. 2006). However, selection effects complicate quasar targeting in this range. The colours of normal (unobscured) quasars around $z \sim 2.7$ resemble far more abundant stellar populations, particularly metal-poor A and F halo stars (e.g. Fan 1999; Richards et al. 2001a). This issue is enhanced at the faint limits of imaging surveys that achieve similar depth to the Sloan Digital Sky Survey (SDSS; York et al. 2000), where those compact galaxies that are dominated by

A and F stellar populations contaminate selection at the 10–20 per cent level as star–galaxy separation becomes difficult (e.g. Richards et al. 2009; Bovy et al. 2012). In addition, faint quasars with $z \sim 2.2$ – 2.6 can have similar colours to quasars at $z \sim 0.5$, which are contaminated by redder light from their host galaxy (e.g. Budavári et al. 2001; Richards et al. 2001b; Weinstein et al. 2004). In combination, the cuts that must be made to efficiently select quasars in this redshift range mean that optical surveys of quasars may miss a significant number of quasars. A wide-area survey at high targeting density is thus an attractive proposition for the study of quasar clustering at moderate redshift. The data we consider in this paper are drawn from just such a survey: the Baryon Oscillation Spectroscopic Survey (BOSS; Eisenstein et al. 2011).

Physical effects also conspire to make quasars difficult to sample at $z \sim 2.5$. Obviously, quasars simply appear fainter at greater distances, but in addition quasars seem to exhibit cosmic downsizing, with the population of less luminous quasars peaking at lower redshift (Croom et al. 2009). Thus, the most luminous quasars are both more abundant and the most visible members of the quasar population at $z \sim 2.5$. Quasar clustering measurements at high redshift from the original SDSS (Shen et al. 2007) therefore only sample the most luminous quasars, implying that deeper spectroscopy than used for the original SDSS quasar survey (Richards et al. 2002a) is necessary for sampling quasars across a large dynamic range in luminosity near $z \sim 2.5$. Indeed, the final BOSS quasar sample should be ~ 10 times larger, and almost 2 mag deeper, than the original SDSS spectroscopic quasar sample at $2.2 < z < 3.5$ (14 065 objects; Schneider et al. 2010).

The outline of the paper is as follows. In Section 2 we describe the quasar samples that we use – drawn from the SDSS and the BOSS. The clustering measurements are described in Section 3, including comparisons with earlier work. The implications of our results for quasars are explored in Section 4. We conclude in Section 5. Appendix A discusses the impact of redshift errors on our measurements, while Appendix B contains the technical details of the model fits used in this paper. Where necessary we shall adopt a Λ cold dark matter cosmological model with $\Omega_{\text{mat}} = 0.274$, $\Omega_{\Lambda} = 0.726$ and $\sigma_8 = 0.8$ as assumed in White et al. (2011), Anderson et al. (2012) and Reid et al. (2012). Unless the h dependence is explicitly specified or parametrized, we assume $h = 0.7$. Dark matter halo masses are quoted as M_{180b} , i.e. the mass interior to a radius within which the mean density is 180 times the background density of the Universe. Luminosities will be quoted in Watts, and magnitudes in the AB system.

2 DATA

The SDSS (York et al. 2000) mapped nearly a quarter of the sky using the dedicated Sloan Foundation 2.5-m telescope (Gunn et al. 2006) located at Apache Point Observatory in New Mexico. A

drift-scanning mosaic CCD camera (Gunn et al. 1998) imaged the sky in five photometric bandpasses (Fukugita et al. 1996; Smith et al. 2002; Doi et al. 2010) to a limiting magnitude of $r \simeq 22.5$. The imaging data were processed through a series of pipelines that perform astrometric calibration (Pier et al. 2003), photometric reduction (Lupton et al. 2001) and photometric calibration (Padmanabhan et al. 2008). We use quasars selected from this five-band SDSS photometry as described in detail in Bovy et al. (2011) and Ross et al. (2012).

Selecting quasars in the redshift range $z \simeq 2-3$, where the space density of the brightest quasars peaks (Richards et al. 2006; Croom et al. 2009), is made difficult by the large populations of metal-poor A and F stars, faint lower redshift quasars and compact galaxies which have similar colours to the objects of interest (e.g. Fan 1999; Richards et al. 2001a). In our case the problem is compounded by the fact that we wish to work close to the detection limits of the SDSS photometry, where errors on flux measurements cause objects to scatter substantially in colour space, and that the BOSS key science programmes did not include the study of the clustering of $z \sim 2.5$ quasars.

2.1 Clustering subsamples and the angular mask

The quasar component of BOSS is designed primarily as a Ly α forest survey, which does not require quasars to be selected in a uniform – or even a *recreatable* – manner across the sky. For the study of quasar clustering however, uniform selection is key. To satisfy these competing scientific requirements, the survey thus adopted a CORE+BONUS strategy, where the CORE objects correspond to a uniform sample selected by the *extreme deconvolution* (XD)¹ algorithm. XD is applied in BOSS to model the distributions of quasars and stars in flux space, and hence to separate quasar targets from stellar contaminants (XDQSO; Bovy et al. 2011). It is this CORE sample that we analyse in this paper.

Specifically, we take as quasar targets all point sources in SDSS imaging that have an XDQSO probability above a threshold of 0.424 to the magnitude limit of BOSS quasar target selection ($g \leq 22.0$ or $r \leq 21.85$; Ross et al. 2012). By quasar targets in this sense, we mean all quasars that would have been observed in a perfect survey. In reality, not all such targets are observed because not all fibres can be placed on an object during normal survey operations. More importantly, the XDQSO algorithm was not adopted as the final CORE algorithm for BOSS quasar target selection until the second year of operations (Ross et al. 2012). We will work with spectroscopy taken on or before 2012 January 1 – just over the first two years of BOSS data. So, on average, more potential targets with an XDQSO probability greater than 0.424 are unobserved (spectroscopically) in areas that were covered in the first year of BOSS (see Fig. 1).

We use the MANGLE software (Swanson et al. 2008) to track the angular completeness of the survey (the mask). The completeness on the sky is determined from the fraction of quasar targets in a sector for which we obtained a spectrum; a sector is an area of the sky covered by a unique set of spectroscopic tiles (see Blanton et al. 2003; Tegmark et al. 2004). For our analyses, we limit the survey to areas with targeting completeness greater than 75 per cent. By

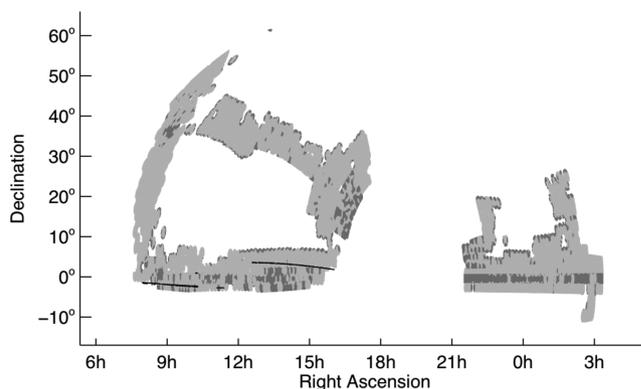


Figure 1. The angular distribution of our quasar sample, in J2000 equatorial coordinates and Aitoff projection. We have rotated the reference line by 90° so that the North and South Galactic survey regions appear contiguous in the left- and right-hand parts of the plot, respectively. Areas which we use in our analysis (light grey) have completeness to XDQSO targets of greater than 75 per cent. Other areas (dark grey) are mainly early survey regions where XDQSO was not used as the CORE targeting algorithm. The black areas depict imaging data in which the u -band chip was not operating, which are discarded from our analysis.

targeting completeness, we mean the ratio of the number of quasar targets that received a BOSS fibre to *all* quasar targets.

We do not correct for *spectroscopic* incompleteness – i.e. account for the fraction of observed targets which produce a spectrum of sufficient quality to measure a redshift. Quasars at $z > 2.2$ are identifiable in BOSS even at very low signal-to-noise ratio because the strong Ly α $\lambda 1215$ line always falls within the BOSS wavelength coverage (3600–10 000 Å; Eisenstein et al. 2011). In BOSS, almost all unidentifiable objects are likely to be stars, galaxies or low-redshift quasars, not the $z > 2.2$ quasars of interest in this paper. Correcting for spectroscopic incompleteness would induce a false large-scale clustering signal because the density of stellar contaminants varies over the sky.

We apply a veto mask to remove survey regions in which a quasar could never be observed – areas near bright stars and the centreposts of the spectroscopic plates (as described in White et al. 2011). We also remove fields where the conditions were not deemed photometric by the SDSS imaging pipeline (again see White et al. 2011). Finally, we remove areas that have bad u data in the SDSS imaging scans (section 3.3 of Abazajian et al. 2004, see Fig. 1). The resulting XDQSO CORE targets were matched to the list of objects for which the BOSS successfully obtained a spectrum, and throughout this paper we only consider regions where at least 75 per cent of the XDQSO CORE targets received a BOSS fibre for spectroscopic observation.

We study BOSS data taken on or before 2012 January 1. This limits our analysis to version 5_5_0 of the spectral reduction pipeline (spA11-v5_5_0) and to the areas plotted in Fig. 1. Note that these are slightly later reductions than made publicly available with SDSS Data Release 9 (DR9; which uses v5_4_45). Algorithmically v5_5_0 is the same as v5_4_45, but changes for calibrations of a newly installed CCD affect data past DR9.

2.2 Redshift assignment

The BOSS wavelength coverage is 3600–10 000 Å so for quasars above $z \sim 1$ the [O III] $\lambda\lambda 4958, 5007$ complex is shifted out of the BOSS window. Systemic redshifts for most BOSS quasars thus rely solely on information from broad emission lines in the rest-frame

¹ XD (Bovy, Hogg & Roweis 2009) is a method to describe the underlying distribution function of a series of points in parameter space (e.g. quasars in colour space) by modelling that distribution as a sum of Gaussians convolved with measurement errors.

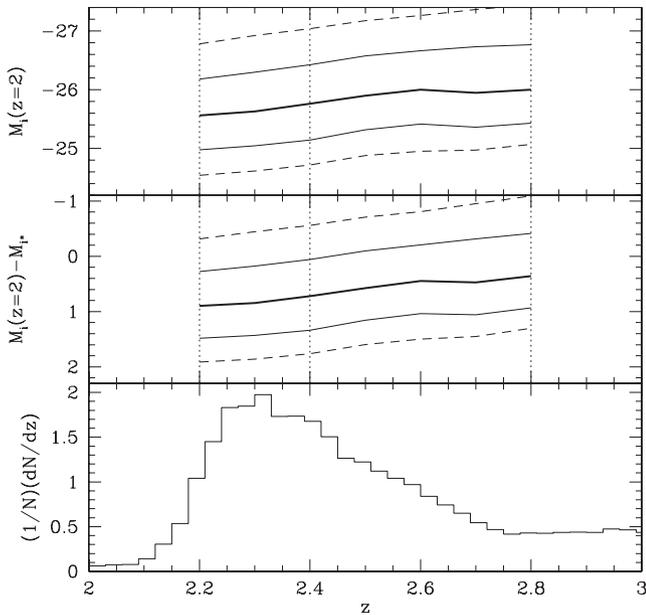


Figure 2. The absolute magnitude distribution and number of quasars versus redshift for our sample. Upper panel: the 10th, 25th, 50th, 75th and 90th percentiles of M_i versus redshift (see text). Middle panel: the same percentiles now in $M_i - M_{*,i}$ versus redshift. Lower panel: the (normalized) redshift distribution of quasars. The vertical dotted lines indicate the redshift ranges we consider in our study.

ultraviolet. The next lowest ionization line typically found in quasar spectra Mg II λ 2798, which is a good redshift indicator (Richards et al. 2002b), is shifted entirely beyond the BOSS spectral coverage near $z \sim 2.5$. To measure quasar redshifts above $z \sim 2.5$, we rely on combinations of the C III λ 1908, C IV λ 1549 and Ly α λ 1215 lines. In addition to being broad, the centroid of each line may be biased – C III] is often blended with Si III], Al III and Fe III complexes, C IV can be shifted from the systemic redshift by strong outflows, and Ly α is often affected by Ly α forest absorption and is blended with N V (e.g. Vanden Berk et al. 2001; Richards et al. 2002b, 2011).

To ameliorate these issues, when the BOSS pipeline² identifies a quasar spectrum as having an Mg II, C III] or C IV line that is within the BOSS wavelength coverage, we default to the redshift from that line offset using the prescription of Hewett & Wild (2010). For spectra with no such lines, we adopt the BOSS pipeline redshift. The BOSS collaboration is visually inspecting all quasar spectra to check the pipeline redshifts. When a pipeline redshift conflicts with the visual redshift (438 objects), we adopt the human-corrected redshift. As we mainly analyse clustering on scales that correspond to velocities that are larger than typical quasar broad lines, altering redshifts by a small amount does not strongly affect our results. Appendix A discusses redshift errors further.

2.3 Quasar luminosities and k -corrections

Fig. 2 plots the conditional magnitude distribution of our sample, compared to the characteristic luminosity of quasars at that redshift. We correct all magnitudes to $z = 2$ using the k -corrections derived by Richards et al. (2006). The characteristic luminosity – where the luminosity function (LF) changes slope – from Croom et al. (2004),

as modified by Croton (2009), is

$$M_{i,*}(z) = -21.61 - 2.5 (k_1 z + k_2 z^2) - 0.71, \quad (1)$$

where $k_1 = 1.39$ and $k_2 = -0.29$ for $z < 3$ and $k_1 = 1.22$ and $k_2 = -0.23$ for $z \geq 3$. We have converted from the b_J band used by Croom et al. (2004) to the i band (k -corrected to $z = 2$) using $M_i(z = 2) = M_{b_J} - 0.71$ (Richards et al. 2006). Note that in the range $2.2 < z < 2.8$, which will be the focus of this paper, we are able to probe 1–2 mag fainter than the characteristic magnitude.

Using SDSS broad-band colours to derive k -corrections for our sample is problematic. Most BOSS quasars at redshift $z \sim 3$ are near the flux limit of SDSS imaging, and so they have noisier colours than for previous SDSS prescriptions at brighter limits (e.g. Richards et al. 2006). Deriving full k -corrections at $z \sim 3$ as a function of flux, colour and redshift is beyond the scope of this paper (but see McGreer et al., in preparation). Precise k -corrections will require proper modelling of subtle changes due to, e.g. the Baldwin effect, the presence of complex iron emission crossing through the i band, and the movement of broad lines – which can be offset due to luminosity- and redshift-dependent winds and absorption features – across the SDSS filter set. For the redshifts on which we focus in this paper ($2.2 < z < 2.8$), only the C III λ 1908 line enters the i band. Fortunately, this complex does not shift much with luminosity, which reduces any flux dependence to the k -correction for our sample.

3 CLUSTERING

All of our clustering measurements are performed in configuration, rather than Fourier, space. For rare objects, where shot noise is an important or dominant piece of the error budget, the configuration space estimators have the advantage of more nearly independent errors. They also deal well with irregularly sampled geometries, such as we have for our sample. We shall compute both the real- and redshift-space correlation functions, using the Landy & Szalay (1993) estimator, with a density of random points 50 times the density of quasars.

3.1 The random catalogue

As discussed in Section 2.1 we use MANGLE (Swanson et al. 2008) to track the angular completeness of the survey. Angular positions for the random points, modulated by the angular completeness of the survey, were obtained from the MANGLE program ransack. To assign redshifts to the random catalogue we tried three methods, which yielded almost identical results. The first was to assign to each point a redshift drawn at random from the data. While this method would produce artificial structure in the redshift distribution of the random points for a small survey, due to sample variance, the wide angular coverage of the BOSS survey ensures that this method performs well. We also tried fitting splines to the histogram of the quasar redshifts and the cumulative histogram of the redshifts and using those splines to generate random redshifts. The results were insensitive to using the histogram or cumulative histogram, to the number of spline points and to the type of spline used. For the results presented below we use the first method.

3.2 Fibre collisions and small-scale clustering

We cannot obtain spectroscopic data for a few per cent of quasars due to fibre collisions – no two BOSS fibres can be placed closer than 62 arcsec on a specific plate. At $z \simeq 2.5$ the 62-arcsec exclusion

² The BOSS pipeline is described in Aihara et al. (2011).

corresponds to $1.26 h^{-1}$ Mpc (comoving). Where possible we obtain redshifts for the collided quasars in regions where plates overlap. We account for the remaining exclusions by restricting our analyses to relatively large scales and by upweighting quasar–quasar pairs with separations smaller than 62 arcsec. The upweighting is derived by comparing the angular correlation function of all targets with that of the quasars for which we obtained redshifts (Hawkins et al. 2003; Li et al. 2006; Ross et al. 2007; White et al. 2011). This ratio is close to unity above 62 arcsec but drops to about two-thirds below 62 arcsec. The number of pairs for which this correction is appreciable is quite small, and the impact of this correction is much less than 1σ even on the smallest scales. If fibre-collided quasars preferentially live in regions of higher than average density, the large-scale clustering would be affected by fibre collisions and would not be properly corrected by our weighting procedure. The efficiency of the tiling algorithm, the prioritization of quasar targets over galaxies and the depth of the survey combine to make fibre collisions a very small effect on our analysis (see also Ross et al. 2009).

3.3 Tests of systematics

We have performed numerous jackknife tests to check whether our results are robust to possible systematics. Specifically we have investigated whether our results are stable to cuts on targeting, at what point in the survey the plates were drilled and what targeting algorithm was used, sector completeness, Galactic latitude and hemisphere (Fig. 3), extinction in the g band, target areal density, stellar density, raw i -band magnitude (which is a proxy for signal-to-noise ratio), sky brightness, g -band seeing (Fig. 4) and selection threshold. In all cases but one we see no evidence for a statistically significant systematic effect. The exception is that there is weak evidence that the large-scale clustering of quasars in the South Galactic Cap is stronger than that in the North Galactic Cap. It will require more data to determine whether this is a statistical fluctuation or a significant difference – and, of course, we conducted 11, not one, different tests of systematics. When 11 (independent) trials are performed the likelihood of a 2σ detection is quite high (~ 40 per cent instead of ~ 5 per cent for a single trial). In future a quasar catalogue with good photometric redshifts could help with some of these is-

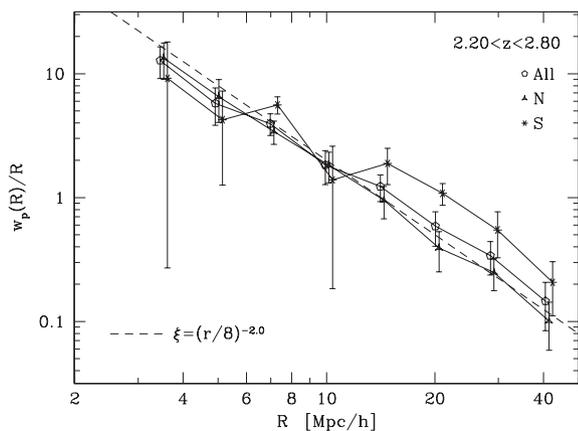


Figure 3. The projected correlation function split by hemisphere (or Galactic latitude), compared to the fiducial sample. The dashed line corresponds to the projected correlation function for a real-space correlation function with $r_0 = 8 h^{-1}$ Mpc and a power-law slope of -2 to guide the eye. Note the weakly significant excess power at large scales for the south-only sample (see text).

Figure 4. The projected correlation function split by whether the median seeing in g band in a sector is better than 1.25 arcsec (SB125) or worse than 1.25 arcsec (SW125) compared to the fiducial sample. The dashed line corresponds to the projected correlation function for a real-space correlation function with $r_0 = 8 h^{-1}$ Mpc and a power-law slope of -2 to guide the eye. There is no statistically significant difference between the two halves of the data. This is typical of the other jackknife tests we have performed.

sues. In addition, BOSS will continue to obtain quasar data until 2014, so we defer a more detailed investigation of geographical discrepancies to a future publication.

3.4 Clustering results

We have insufficient sensitivity to measure the angular dependence of the redshift-space clustering induced by redshift-space distortions for our highly biased quasars. Therefore, we only quote redshift-space results from the angle-averaged correlation function, which we denote $\xi(s)$ at redshift-space separation s . Real-space clustering is constrained by the projected correlation function

$$w_p(R) \equiv \int dZ \xi(R, Z), \quad (2)$$

avoiding the need to model redshift-space distortions and mitigating any effects of redshift errors. We truncate the integral over the line-of-sight separation, Z , to $\pm 50 h^{-1}$ Mpc. This value represents a trade-off between the goal of fully integrating out the effects of redshift-space distortions and the disadvantages of introducing noise from only weakly correlated structures along the line of sight and mixing a wide range of 3D scales into a single R bin. By $50 h^{-1}$ Mpc the effects of redshift-space distortions are negligible, and the truncation has only a modest effect on our largest scale point. However, this truncation must be kept in mind when precise modelling of the data at the largest R is important (see below).

We divide our quasar sample into bins of redshift over which the bias and mass correlation function are evolving strongly. Fortunately, on the scales of relevance the effects approximately cancel, i.e. the clustering amplitude stays approximately constant. The redshift-bin-averaged ξ can be approximated as a measurement of ξ evaluated at an effective redshift, z_{eff} :

$$z_{\text{eff}} = \frac{\int dz (dN/dz)^2 (H/\chi^2) z}{\int dz (dN/dz)^2 (H/\chi^2)}$$

so that

$$\xi(s, z_{\text{eff}}) \simeq \langle \xi(s) \rangle = \frac{\int dz (dN/dz)^2 (H/\chi^2) \xi(s, z)}{\int dz (dN/dz)^2 (H/\chi^2)}, \quad (3)$$

Table 1. A summary of the quasar samples we consider. The columns list the sample number and name, redshift and magnitude ranges, and the number of quasars. Magnitudes are k -corrected to $z = 2$.

Sample	Name	Redshift	Magnitude	N_{qso}
1	All	$2.2 < z < 2.8$	$[-50.0, -10.0]$	27 129
2	Bright	$2.2 < z < 2.8$	$[-50.0, -25.8]$	13 564
3	Dim	$2.2 < z < 2.8$	$[-25.8, -10.0]$	13 564
4	Fid	$2.2 < z < 2.8$	$[-27.0, -25.0]$	19 111
5	LoZ	$2.2 < z < 2.4$	$[-27.0, -25.0]$	8835
6	HiZ	$2.4 < z < 2.8$	$[-27.0, -25.0]$	9977

where dN/dz is the redshift distribution of the sample, H is the Hubble parameter at redshift z and χ is the comoving angular diameter distance to redshift z (Matarrese et al. 1997; White et al. 2008). Assuming passive evolution, constant halo mass or constant bias leads to differences between $\xi(z_{\text{eff}})$ and $\langle \xi \rangle$ that are far smaller than our observational errors.

Results for each of the samples in Table 1 are shown in Figs 5 and 6, and the values for the full sample, with $2.2 < z < 2.8$ and no cuts on magnitude, are given in Table 2. Both the real- and redshift-space clustering are fitted well by a model with an underlying power-law correlation function, once the effects of projection and redshift errors are taken into account. As it provides a good fit to the data, and for consistency with earlier work (e.g. Myers et al. 2006; Ross et al. 2009; Shen 2009), we shall show lines in the figures and provide fits in the tables assuming a power-law slope of r^{-2} for the real-space correlation function. The actual slope of the correlation function is poorly determined by the projected correlation function data. Power-law indices from -1 to -2.6 are viable. The best-fitting slope is quite shallow, near -1.2 . The value of $w_p(R)$ at $R \simeq 9 h^{-1}$ Mpc for the best-fitting model is almost independent of the assumed slope.

We estimate the covariance matrix of our measurements by bootstrap resampling (e.g. Efron & Gong 1983). We divide the survey

into angular regions specified by HEALPIX pixels (Gorski et al. 2005) with $N_{\text{side}} = 4$ (i.e. approximately 15° on a side). Pixels which contain fewer random points than two-thirds of the mean are merged with higher occupancy pixels to ensure that pixels have similar weight. We then estimate both the mean and covariance matrix by bootstrap resampling pair counts from a random draw of pixels (with replacement).

The bootstrap-determined correlation matrices for the full $2.2 < z < 2.8$ sample for w_p and ξ are shown in Figs 7 and 8, respectively. Note that the matrices are diagonal dominated as might be expected for shot-noise-limited measurements – extending to larger scales we see larger correlations between the bins, most notably in w_p . As we begin to restrict the quasar sample in redshift or luminosity, and the number of objects becomes smaller, the covariance matrices and their inverses can become increasingly noisy. We have several options at this point – with the most desirable being a reduction in the number of degrees of freedom (i.e. data compression) so that we can apply bootstrap resampling to obtain a converged covariance matrix, or error on a summary statistic (or statistics). The simplest approach would be to reduce the number of bins so that we have fewer, better constrained points. However, this is not optimal for our purposes. We describe below a different approach based on the sparsity of our sample and the nature of the clustering.

To begin we note that the correlation functions are fitted well ($\chi^2 = 4.5$ for 7 degrees of freedom for w_p and $\chi^2 = 12$ for 9 degrees of freedom for ξ) by power laws over the range where our constraints are tightest (as expected if quasars are hosted by massive haloes, see Fig. 9). We adopt a two-parameter model for both $\xi(s)$ and w_p of the form $\xi = (s_0/s)^\gamma$ and

$$\frac{w_p(R)}{R} = \frac{\sqrt{\pi} \Gamma[(\gamma - 1)/2]}{\Gamma[\gamma/2]} \left(\frac{r_0}{R}\right)^\gamma, \quad (4)$$

which corresponds to a 3D correlation of the form $\xi(r) = (r_0/r)^\gamma$ integrated to $\pm\infty$ along the line of sight. For $\gamma = 2$ the pre-factor is π and we shall fix $\gamma = 2$ throughout so that we have one remaining degree of freedom. We fit these models to the measured correlations over the range $3 < s < 25 h^{-1}$ Mpc and $4 < R < 16 h^{-1}$ Mpc.

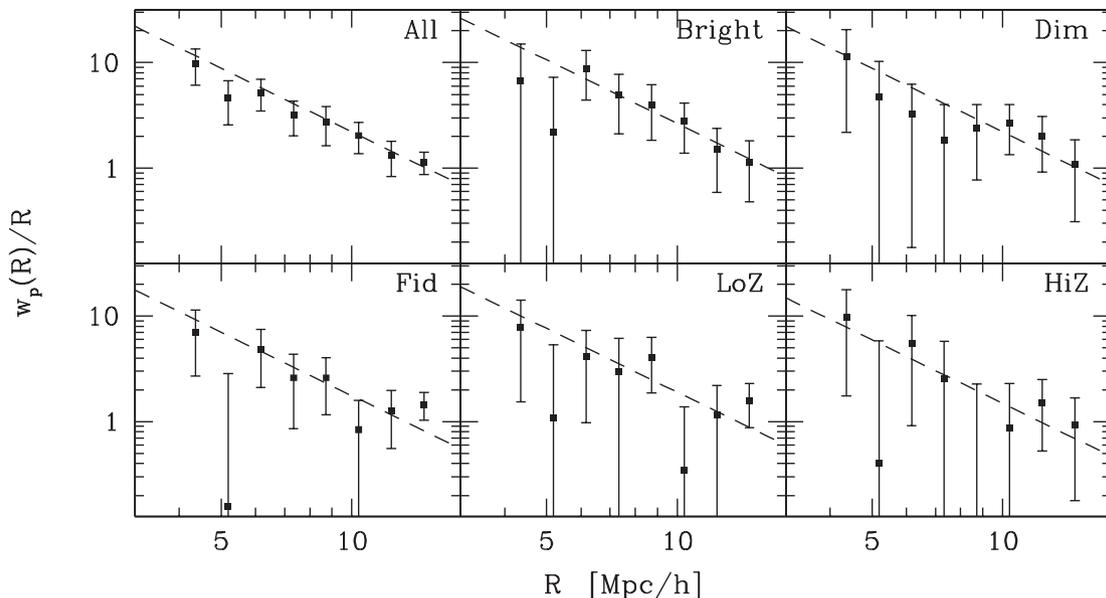


Figure 5. The projected correlation functions, $w_p(R)$, for the six samples considered in this paper (Table 1). The error bars are the square roots of the diagonal elements of the covariance matrices, as determined by bootstrap resampling (see text). The dashed lines show the best-fitting power laws with slope -2 (see Table 3).

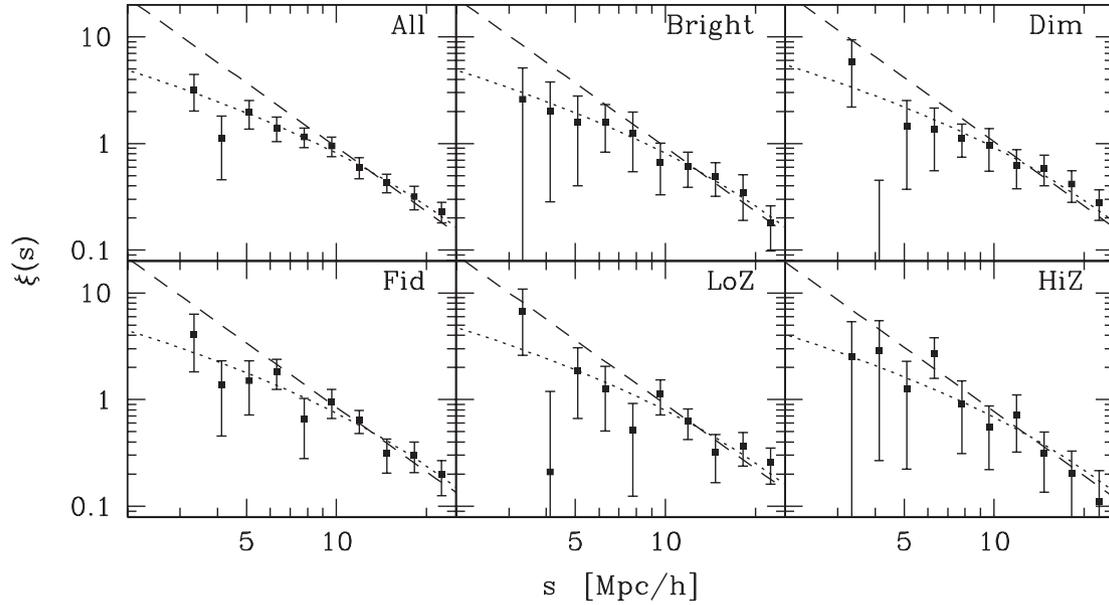


Figure 6. The redshift-space correlation function, $\xi(s)$, for the six samples considered in this paper. The error bars are the square roots of the diagonal elements of the covariance matrices, as determined by bootstrap resampling (see text). The dashed lines show the best-fitting power laws with slope -2 (see Table 3), while the dotted lines show the power law once redshift errors are taken into account (see Appendix A).

Table 2. The w_p data for sample 1, ‘All’, (the largest data set). The first three rows list the transverse separation, R , w_p and its error (all in h^{-1} Mpc). The remainder of the table presents the correlation coefficients as estimated from the covariance matrix computed using bootstrapping, as described in the text. These values are plotted as the grey-scale image in Fig. 7.

R	4.36	5.19	6.17	7.34	8.72	10.37	12.34	14.67
w_p	42.87	24.18	32.12	23.21	23.98	21.13	16.25	16.80
σ	17.08	10.89	10.77	8.36	9.85	7.12	6.16	4.08
4.36	1.000	0.459	0.156	0.125	0.187	0.154	-0.198	-0.146
5.19	-	1.000	0.210	0.094	0.261	0.032	-0.256	-0.019
6.17	-	-	1.000	0.341	-0.029	0.118	0.168	0.173
7.34	-	-	-	1.000	0.069	0.404	-0.113	-0.076
8.72	-	-	-	-	1.000	0.004	0.053	0.091
10.37	-	-	-	-	-	1.000	-0.107	-0.119
12.34	-	-	-	-	-	-	1.000	0.251
14.67	-	-	-	-	-	-	-	1.000

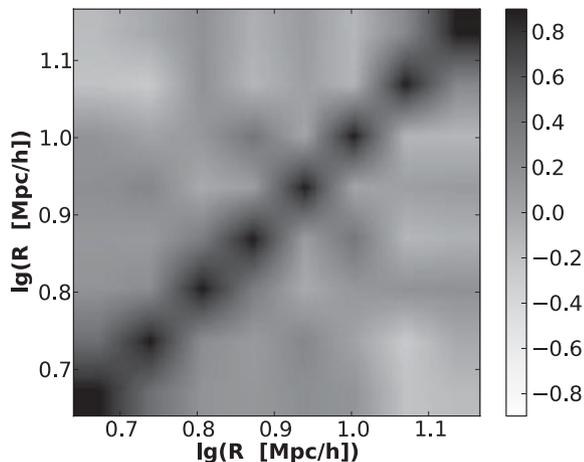


Figure 7. The correlation matrix for the projected correlation function, $w_p(R)$, of quasars with $2.2 < z < 2.8$ and no cuts on magnitude (i.e. sample 1).

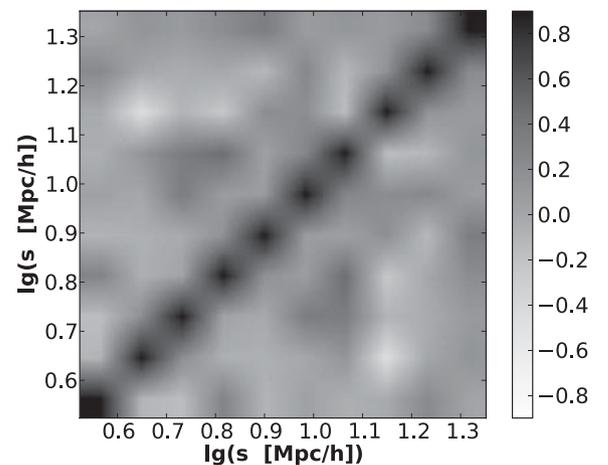


Figure 8. The correlation matrix for the redshift-space correlation function $\xi(s)$, of quasars with $2.2 < z < 2.8$ and no cuts on magnitude (i.e. sample 1).

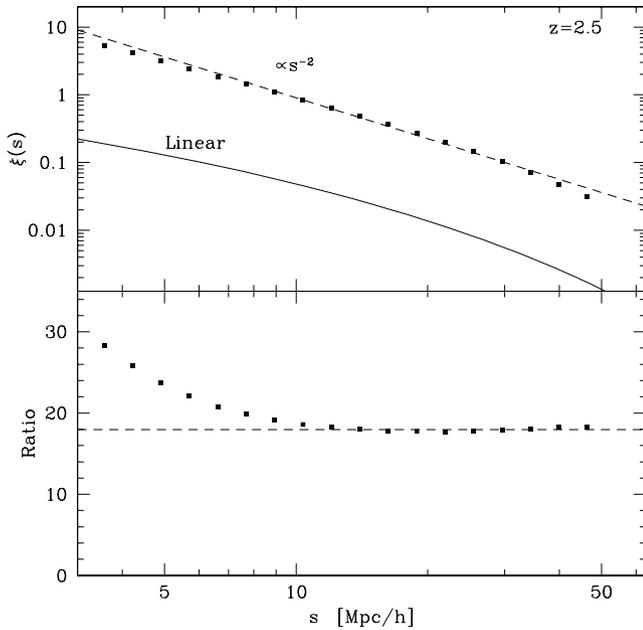


Figure 9. The clustering of a sample of haloes spanning an octave (i.e. a factor of 2) in mass centred on $2 \times 10^{12} h^{-1} M_{\odot}$. This sample has a narrow range of halo masses with a clear ‘characteristic’ mass while also containing enough haloes to enable a precision estimate of the correlation function. The solid line shows the linear-theory, real-space correlation function for our fiducial cosmology, while the dashed line is a power law with slope -2 . The points are computed from haloes at $z = 2.5$ from two simulations of 3000^3 particles ($m = 6 \times 10^{10} h^{-1} M_{\odot}$) each in a $2.75 h^{-1}$ Gpc box run using the code described in White (2002). The haloes are found using the friends-of-friends method (Davis et al. 1985) with a linking length of 0.168 times the mean inter-particle separation. In the upper panel the points show the mean of the angle-averaged, redshift-space correlation functions computed from the periodic boxes in the distant observer approximation. The lower panel shows the ratio of the two, which becomes constant at large scales. The horizontal dashed line is a fit to the asymptote in the range of $20 < s < 40 h^{-1}$ Mpc. Assuming $\xi(s) = [b^2 + (2/3)bf + f^2/5]\xi_{\text{lin}}$ (Kaiser 1987) with $f \equiv [\Omega_m(z)]^{0.56}$ gives $b \simeq 3.9$, in good agreement with the value inferred from the real-space clustering. The slope and amplitude of the power-law piece of $\xi(s)$, and how low in s it extends, vary with the particular halo subsample chosen, but the scale independence of the bias at large scales is generic.

The non-linear, scale-dependent bias of dark matter haloes makes their correlation functions close to a power law on Mpc scales (see Figs 9 and 10) but at larger scales the bias becomes scale independent and the correlation functions drop more quickly than the power-law extrapolation would suggest. In addition, on larger scales sample variance becomes increasingly important and the radial bins become increasingly correlated (especially for w_p). For these reasons we limit the fitting range as indicated (see also Croom et al. 2005). With more data and a numerical model for the covariance matrix we could extend the range of the measurement and tighten the constraints on quasar models.

For a power-law correlation function of fixed slope each point provides an estimate of the correlation length, s_0 , since $s_0^\gamma = s_i^\gamma \xi(s_i)$. The number of pairs in a bin of fixed $\Delta \log s$ is $N_{\text{pair}} \propto (1 + \xi)s^3$ and if shot noise dominates the bins are independent and the fractional error on $1 + \xi(s)$ in each bin scales as $N_{\text{pair}}^{-1/2}$. Thus, assuming logarithmic bins and that shot noise dominates we can

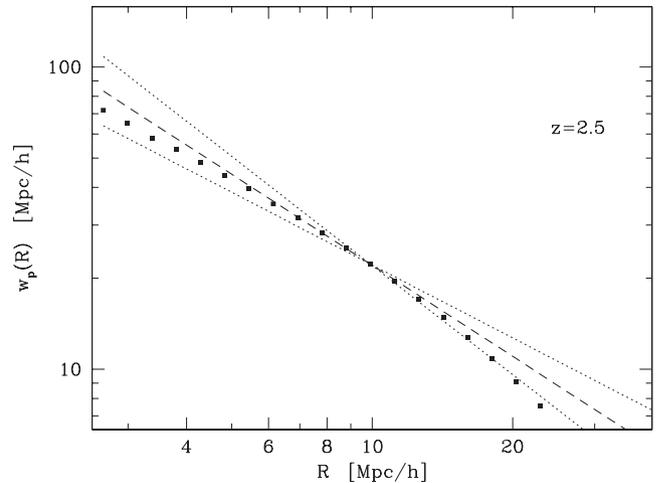


Figure 10. The projected correlation function for the same sample of haloes described in Fig. 9. The dotted and dashed lines correspond to power laws of slope -0.8 , -1 and -1.2 , arbitrarily normalized to the data at $R \simeq 10 h^{-1}$ Mpc. We imposed an upper limit of $\pm 50 h^{-1}$ Mpc on the line-of-sight separation in equation (2). Changing the mix of haloes in the sample can change the amplitude and slope of w_p , but the break is a generic feature.

average the estimates with inverse variance weights so that

$$s_0 = \left[\frac{\sum_i s_i^\gamma \xi(s_i) w_i}{\sum_i w_i} \right]^{1/\gamma}, \quad (5)$$

where the weights are

$$w_i^{-1} \propto \left[1 + \left(\frac{s'_0}{s_i} \right)^\gamma \right] s_i^{2\gamma-3}. \quad (6)$$

The values of w_i depend on an estimate, s'_0 , for s_0 , and reduce to

$$w_i \propto \frac{s_i}{s_i^2 + s_0^2} \quad \text{for } \gamma = 2. \quad (7)$$

Most of the weight in the fit is produced by $s_i \sim s'_0$, with the weight scaling as $s^3 - \gamma$ for small s and $s^{3-2\gamma}$ for large s . Since the bins are assumed to be logarithmically spaced, for $\gamma \approx 2$ this suppression is quite rapid in both directions, reflecting the paucity of pairs at small s and the weakness of the correlation signal at large s . If the correlation function continued as a power law to large s we could tighten our constraints by extending the fitting range, but it is at these larger scales where deviations from a power-law behaviour are most expected, and where the correlations between measurements in adjacent radial bins act to weaken the constraints.

Unfortunately, it is difficult to find a precise redshift for quasars in the range $2 < z < 3$ from optical spectroscopy, so our correlation function is smeared by redshift errors which reduce the small-scale clustering signal (the dotted lines in Fig. 6; Appendix A discusses the impact of redshift uncertainties further). We include redshift errors in our model through a multiplicative factor, $\xi(s) \rightarrow F(s)\xi(s)$, derived in Appendix A. The best-fitting s_0 can be derived from a generalization of equation (5) which replaces the weights with $w_i \rightarrow F^2(s_i)w_i$, multiplies the $\xi(s_i)$ in the weights by $F(s_i)$ and divides each term in the numerator of equation (5) by $F(s_i)$.

For a given estimate, s'_0 , the optimal estimate of s_0 can be written either as a weighted sum of the ξ_i bins or directly in terms of the pair counts themselves. One can iterate this estimator, in which case the error properties of s_0 become more complex and are best

handled by a bootstrap procedure. We generate s_0 for each bootstrap sample, with the iterative weighting scheme above starting from $s'_0 = 0$, and use the standard deviation for our estimate of uncertainty. This uncertainty does not include the additional contribution from our uncertainty in the redshift error correction. A change in redshift error (σ_z) of ± 25 per cent moves the best-fitting s_0 by 1σ , which can be considered as an additional systematic error on the fit. While a 25 per cent uncertainty on the redshift error is consistent with Appendix A, we find that reasonable fits to $\xi(s)$ can be obtained for a wide range of redshift error due to a degeneracy between the assumed redshift error and the amplitude (and slope) of the underlying correlation function. Because of this additional uncertainty, our strongest cosmological constraints will come from the projected correlation function, to which we now turn.

A similar estimator can be used for the real-space correlation length, r_0 , under the same assumptions. The weights in this case become $w_i^{-1} \propto (2Z_{\max} + w_{pi})R_i^{2\gamma-4}$ if the integration in equation (2) extends from $-Z_{\max}$ to Z_{\max} . For $\gamma \approx 2$ the weights are nearly constant for all the samples we consider for all R of interest. As for the case of the redshift errors and $\xi(s)$, we can account for the effects of finite Z_{\max} on equation (4) by modifying the weights to $w_i \rightarrow F^2(R_i)w_i$ and dividing each term in the numerator by $F(R_i)$, where³ $F(R) = (2/\pi) \arctan(Z_{\max}/R)$ for $\gamma = 2$.

The clustering strength derived from this procedure is a statistically valid summary of the data under the assumption that a power law provides a good fit to w_p . However, it does not need to be an optimal compression of the available information – if the data are sufficiently informative a better constraint on the clustering amplitude could be obtained by fitting all of the data. For the full sample (1 in Table 1), where we have a reasonably converged estimate for the covariance matrix, we can compare the different methods. In this case, the likelihood derived from the r_0 determined as above is very similar to that derived from the full covariance matrix (or the diagonal elements) indicating that in our case our estimate of r_0 does provide an almost exhaustive summary of the constraints available (Fig. 11). This will be even more the case for samples with more shot noise. Our simple estimator is the preferred approach for quoting clustering measurements and errors on sparse samples where estimating a full covariance matrix is not feasible.

In addition to our estimates of r_0 and s_0 we also provide another summary statistic (motivated by Croom et al. 2005; da Angela et al. 2008; Ross et al. 2009),

$$\bar{\xi} \equiv \frac{3}{s_{\max}^3 - s_{\min}^3} \int_{s_{\min}}^{s_{\max}} s^2 ds \xi(s), \quad (8)$$

with $s_{\min} = 5 h^{-1}$ Mpc and $s_{\max} = 20 h^{-1}$ Mpc. For $\xi(s) = (s_0/s)^2$ equation (8) becomes

$$\bar{\xi} = \frac{3s_0^2}{s_{\max}^2 + s_{\max}s_{\min} + s_{\min}^2} \quad (9)$$

$$\approx 3 \left(\frac{s_0}{s_{\max}} \right)^2 \left[1 - \frac{s_{\min}}{s_{\max}} + \dots \right], \quad (10)$$

where the last step assumes $s_{\min} \ll s_{\max}$. We adopt a lower limit, $s_{\min} \neq 0$, to mitigate the effect of redshift errors and scale-dependent bias – equation (10) shows that this differs from the s_{\min} case by

³ Assuming that Z_{\max} is large enough that fingers-of-god are correctly included and $b > 1$ so that the anisotropy due to redshift-space distortions is small.

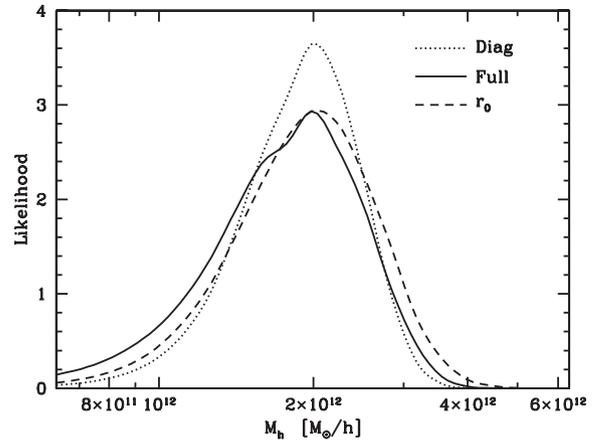


Figure 11. A comparison of fits to the real-space clustering data, using a variety of approximations. The model in each case is the clustering of a sample of haloes covering one octave in mass, centred on M_h . The likelihood of the central mass is computed by fitting to the w_p measured for sample 1 using the full covariance matrix determined by bootstrap (solid), a diagonal covariance matrix (dotted) or the value of r_0 determined as described in the text (dashed). For this model the central mass has an error of $0.5 \times 10^{12} h^{-1} M_{\odot}$.

25 per cent for a power law of index -2 . With our lower limit the bias inferred from modelling $\bar{\xi}$ using the Kaiser (1987) prescription

$$\bar{\xi}(s) \approx \left(b^2 + \frac{2bf}{3} + \frac{f^2}{5} \right) \bar{\xi}_{\text{real}} \quad (11)$$

agrees to 1 per cent with that inferred from the real-space clustering for the simulation results described in Figs 9 and 10. If ξ_{lin} is used in place of ξ_{real} in equation (11), the error is also 1 per cent for the case shown in Figs 9 and 10, though it becomes larger for more biased samples. We compute $\bar{\xi}$ from the data by assuming that $\xi(s)$ can be modelled as a constant within the 10 bins, spaced equally in log between s_{\min} and s_{\max} . The values are corrected upwards by 7 per cent to account for the effect of redshift errors. To allow easy comparison with earlier work we use the value of $\bar{\xi}$ to estimate the bias in Table 3.

Our results are listed in Table 3 and compared to previous work in Fig. 12. We do not detect a luminosity or redshift dependence of the clustering strength, although our sensitivity to this dependence is weak due to the limited dynamic range in both variables in our sample. When comparing to previous work, we do not plot the last four points quoted in Table 2 of the SDSS Data Release 5 (DR5; Adelman-McCarthy et al. 2007) quasar clustering analysis of Ross et al. (2009). Due to the flux-limited quasar selection by the original SDSS, the projected w_p measurement for these high- z bins are quite noisy (Shen et al. 2007; Ross et al. 2009), and thus unreliable for any direct comparison to our BOSS measurement. We have checked that even if we use the SDSS DR7 data this situation is not improved for our $2 < z < 3$ redshift range of interest.

Our results strongly favour the consensus that quasars inhabit rare and highly biased dark matter haloes on the exponential tail of the mass function. In the absence of merging we would expect the clustering of such haloes to evolve slowly with time. For our assumed cosmology our quasar samples have biases in the range of 3.4–4, consistent with early observations and the extrapolation of previous measurements by Croom et al. (2005). Estimates in the literature on the typical halo mass for a bright quasar at comparable

Table 3. A summary of our clustering results. The samples are the same as in Table 1. The effective redshift is computed using equation (3.4). The ranges, medians and mean magnitudes refer to absolute, i -band magnitude k -corrected to $z = 2$. The correlation lengths are measured in (comoving) h^{-1} Mpc. The 1σ errors on r_0 , s_0 and $\bar{\xi}$ are from bootstrap resampling, as described in the text. The error for s_0 does not include the additional uncertainty due to the redshift error from the pipeline. The bias is estimated from $\bar{\xi}$ using equations (8) and (11) with the fitting function of Smith et al. (2003) for the real-space correlation function of the mass.

Sample	Redshift	z_{eff}	Magnitude	Median	Mean	r_0	s_0	$\bar{\xi}$	Bias
1	$2.2 < z < 2.8$	2.39	$[-50.0, -10.0]$	-25.8	-25.9	8.4 ± 0.6	9.6 ± 0.5	0.52 ± 0.06	3.8 ± 0.3
2	$2.2 < z < 2.8$	2.41	$[-50.0, -25.8]$	-26.5	-26.6	9.2 ± 0.9	9.6 ± 1.0	0.54 ± 0.12	3.9 ± 0.5
3	$2.2 < z < 2.8$	2.36	$[-25.8, -10.0]$	-25.2	-25.1	8.4 ± 1.7	10.2 ± 0.8	0.65 ± 0.12	4.3 ± 0.4
4	$2.2 < z < 2.8$	2.39	$[-27.0, -25.0]$	-25.9	-25.9	7.5 ± 0.9	9.2 ± 0.7	0.48 ± 0.06	3.7 ± 0.3
5	$2.2 < z < 2.4$	2.28	$[-27.0, -25.0]$	-25.8	-25.9	7.8 ± 1.3	9.5 ± 0.9	0.51 ± 0.10	3.7 ± 0.4
6	$2.4 < z < 2.8$	2.51	$[-27.0, -25.0]$	-25.9	-25.9	6.9 ± 1.8	8.8 ± 1.1	0.43 ± 0.13	3.6 ± 0.6

redshifts vary wildly, in part due to methodological differences and the fitting functions assumed. So, we now turn to how quasars occupy dark matter haloes (see also Appendix B).

4 INTERPRETATION AND MODELLING

One of the main goals of studying quasar clustering is to provide information on the parent dark matter haloes hosting luminous quasars. The large-scale bias of the quasars provides information on the mean dark matter halo mass; the small-scale clustering provides information on satellite fractions and potentially radial profiles within haloes. Unfortunately, the space density of quasars cannot be used directly in such constraints because the quasar duty cycle (or activity time) is not known. This is a major difference with studies of e.g. galaxy clustering, and has serious implications for the constraints that can be derived.

Our median quasar has $M_i(z=2) = -26$ and so a bolometric luminosity of $L_{\text{bol}} = 2.5 \times 10^{39}$ W (Croom et al. 2005; Shen et al. 2009). If this object is radiating at the Eddington limit [$L_{\text{Edd}} = 10^{40.1} (M_{\text{bh}}/10^9 M_{\odot})$ W], then the median M_{bh} in our sample is $2 \times 10^8 M_{\odot}$. As we shall describe, our data are consistent with host haloes having a characteristic mass of $2 \times 10^{12} h^{-1} M_{\odot}$ (Fig. 11), in agreement with earlier work (Porciani et al. 2004; Croom et al. 2005; Lidz et al. 2006; Porciani & Norberg 2006). This value is also consistent with estimates from the $M_{\text{bh}}-M_{\text{h}}$ relation (Ferrarese 2002; Fine et al. 2006), $M_{\text{h}} \simeq (1-3) \times 10^{12} h^{-1} M_{\odot}$, with the differences arising from different assumptions about the halo profiles or data sets. This result suggests $L/L_{\text{Edd}} \sim 1$, consistent with the results of Croom et al. (2005).

However, we do not expect quasars to inhabit haloes of a single mass (though see Shanks et al. 2011). Constraining the possibly complex manner in which quasars occupy haloes from the relatively featureless correlation functions we have access to is difficult, but the modelling is made easier by a number of facts. Quasars are rare, their activity times are short and the fraction of binary quasars is very small⁴ (Hennawi et al. 2006a; Myers et al. 2007b). This suggests that most quasars live at the centre of their dark matter haloes and the majority of haloes host at most one active quasar. To place constraints on the range of haloes in which quasars may be active we consider two illustrative models described below (see Appendix B for further details and De Graf et al. 2011 for a recent discussion in the context of numerical hydrodynamic simulations).

⁴ When the virial radius of the hosting halo is much smaller than the mean inter-quasar separation, the fraction of quasar-hosting haloes which contain a second quasar scales as $V_{\text{halo}} \bar{n}_Q \xi_V \ll 1$, where ξ_V is the volume-averaged correlation function within the virial radius.

To make interpretation easier we use a quasar sample which is limited at both the bright and faint ends, i.e. $-27 < M_i < -25$ or $39.0 < \log_{10} L_{\text{bol}} < 39.8$ (see Table 1, sample 4, ‘Fid’). We use the r_0 measurements listed in Table 3 to constrain the models, though similar results are obtained using the w_p data and its covariance matrix. We also find that the best-fitting models below provide a good fit to the redshift-space clustering, assuming our fiducial model for redshift errors.

4.1 Lognormal model

In this model we assume that quasars of some luminosity, L , live in haloes with a lognormal⁵ distribution of masses, centred on a characteristic mass that scales with L (see also Appendix B). Each halo hosts at most one quasar with probability

$$P(M_{\text{h}}|L) \propto \exp \left[-\frac{(\ln M_{\text{h}} - \ln M_{\text{cen}}(L))^2}{2\sigma^2(L)} \right], \quad (12)$$

where the normalization is set by the observed space density of quasars but does not matter for the clustering. We expect that M_{cen} will be larger for more luminous quasars that are hosted by more massive galaxies.

For illustration we assume that the lognormal form holds for quasars in the luminosity bin $-27 < M_i < -25$, i.e. that the luminosity dependence of M_{cen} is weak. This is a reasonable approximation to many models (see Appendix B) and in particular the type described in the next section. We generate quasar samples from haloes in the $z = 2.4$ output of four cosmological simulations, each employing 1500^3 particles ($m = 2 \times 10^{10} h^{-1} M_{\odot}$) in a $1 h^{-1}$ Gpc box. As in Fig. 9, haloes are found in the simulations by the friends-of-friends algorithm with a linking length of 0.168 times the mean inter-particle separation. The haloes are selected with a lognormal probability centred on a series of M_{cen} . To test for sensitivity to the width of the distribution we consider $\sigma = 25, 50$ and 100 per cent. The former is more appropriate for higher redshift or brighter quasars (White et al. 2008; Shankar et al. 2010; De Graf et al. 2011), while the latter is roughly expected based on the amount of observed scatter in the local $M_{\text{BH}}-M_{\text{gal}}$ relation. We average the projected correlation function for each model – with the same binning and Z_{max} as the data – over the four simulations to reduce noise, and compute a goodness of fit. Using either the full bootstrap covariance matrix or just the diagonal entries, we find that the best-fitting model is

⁵ Quasars are known to have a relatively high bias and hence live in haloes on the steeply falling tail of the mass function. The differences between occupation models which include a high- M cut-off and those that do not are therefore relatively small.

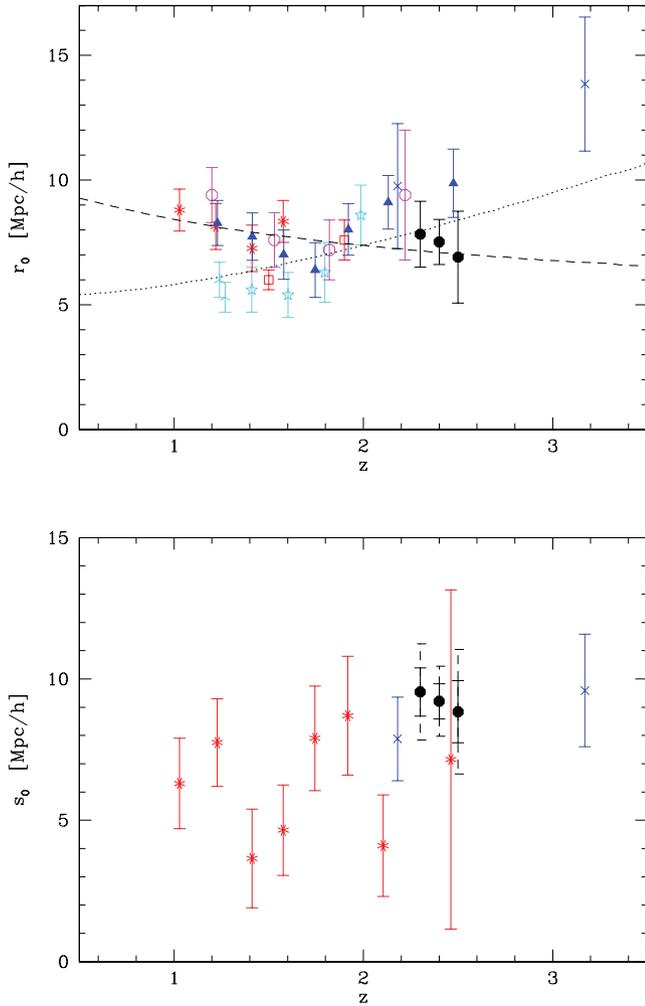


Figure 12. Estimates of the real- and redshift-space correlation lengths for our last three samples (4–6, solid circles) and from previous work. Open squares are from Porciani et al. (2004), solid triangles from Croom et al. (2004, converted from ξ measurements), open five-pointed stars from Porciani & Norberg (2006), open octagons from Myers et al. (2006), eight-pointed stars from Ross et al. (2009), crosses from Shen et al. (2009) and three-pointed stars from Hickox et al. (2011, obscured and unobscured). In the upper panel the dotted line indicates the evolution of r_0 for a sample of haloes of a single mass, $M = 2 \times 10^{12} h^{-1} M_\odot$, while the dashed line shows the evolution of a passively evolving sample with no mergers (Fry 1996). Both lines are meant for illustrative purposes, and there has been no attempt to fit to the data. In the lower panel the solid error bars plotted exclude the contribution from our uncertainty in the redshift error, which is significant (see text). The dashed error bars show the effect of doubling the errors for our measurements.

good for all choices of σ . Not surprisingly, we obtain consistent results if we fit r_0 to the average w_p of the simulations using the procedure of Section 3 and then compare to the value in Table 3. Our measurements suggest $\log_{10} M_{\text{cen}} \simeq 12.00, 12.05$ and $12.15 h^{-1} M_\odot$ for $\sigma = 25, 50$ and 100 per cent, corresponding to an average halo mass $\langle \log_{10} M \rangle \simeq 11.9\text{--}12.0 h^{-1} M_\odot$.

The majority of high-mass galaxies at high redshift are the central galaxy in their dark matter halo, so observational stellar mass functions can provide constraints on the stellar-to-halo-mass relation at $z \sim 2\text{--}3$ (see Behroozi, Conroy & Wechsler 2010; Moster et al. 2010, for recent examples). In combination with the constraints on M_{cen} from quasar clustering, and an assumption about the mean

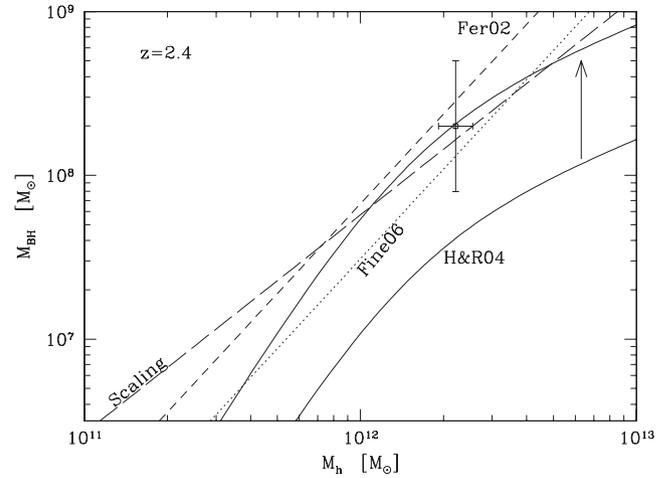


Figure 13. The halo–black hole mass relation. The point with errors shows the mean and 1σ spread in $\log M_h$ for our lognormal model (as in Section 4.1) with 100 per cent scatter. The vertical error bars convert the luminosity range to M_{BH} assuming $L = L_{\text{Edd}}$. The lines indicate $M_{\text{BH}}\text{--}M_h$ relations inferred from the literature and from our scaling model. The long dashed line gives the relation for the scaling model (Section 4.2) which best fits our data. The short dashed line is the result of Ferrarese (2002) and the dotted line with the same slope is the result of Fine et al. (2006). The lower solid line indicates the local scaling relation of Haring & Rix (2004) assuming that M_{bulge} equals the stellar mass inferred from the relation of Moster et al. (2010) at $z = 2.4$. The upper solid line assumes that M_{BH} is five times larger at fixed M_* than the local relation. Note the curvature of the line due to the inefficiency of galaxy formation at high and low halo mass.

Eddington ratio of our sample, we can infer the typical $M_{\text{BH}}\text{--}M_*$ relation for our quasars. Taking the lognormal model and adopting the conversion of Moster et al. (2010), the average stellar mass is $\log_{10}(M_*/M_\odot) = 10\text{--}10.2$ (with larger values corresponding to larger σ). If the accretion occurs at λL_{Edd} the median black hole mass is $\log_{10}(\lambda M_{\text{BH}}/M_\odot) = 8.3$ using the conversions of (Croom et al. 2005; Shen et al. 2009). We compare these numbers to a variety of published $M_{\text{BH}}\text{--}M_h$ relations in Fig. 13. Our results are in broad agreement with the high-redshift inferences but predict larger black holes (at fixed halo mass) than what would be inferred from the local relation of Haring & Rix (2004), even if we assume that all of the stellar mass associated with the halo central galaxy is in the bulge and that quasars radiate at Eddington ($\lambda = 1$). This result argues that M_{BH} should increase, at fixed M_* , by a factor of approximately $5 \lambda^{-1}$ between $z = 0$ and $z \simeq 2.4$. This change is consistent with the increase measured in lensed quasar hosts by Peng et al. (2006) and the model of Hopkins et al. (2007b). By comparison, the model of Croton et al. (2006, Fig. 1) predicts roughly an order-of-magnitude increase in M_{BH} at $M_* \sim 10^{10}$ between $z = 0$ and 3. On the other hand, the simulations of Sijacki et al. (2007, fig. 15) predict almost no evolution at the massive end. Merloni et al. (2010) infer evolution of $M_{\text{BH}}\text{--}M_*$ from the zCOSMOS survey with a best-fitting power law of $(1+z)^{0.68}$ – a factor of 2.3 between $z = 0$ and 2.4 – while Decarli et al. (2010) measure a best-fitting power law of $(1+z)^{0.28}$ – a factor of 1.4 between $z = 0$ and 2.4 – from a carefully constructed sample of 96 quasars drawn from the literature. Our result favours stronger evolution, but given the statistical and systematic uncertainties in all of the measurements, the uncertainties in stellar mass estimates and selection biases towards more massive black holes in flux-limited surveys all we can say is that it is encouraging that we see evolution in the same *sense*.

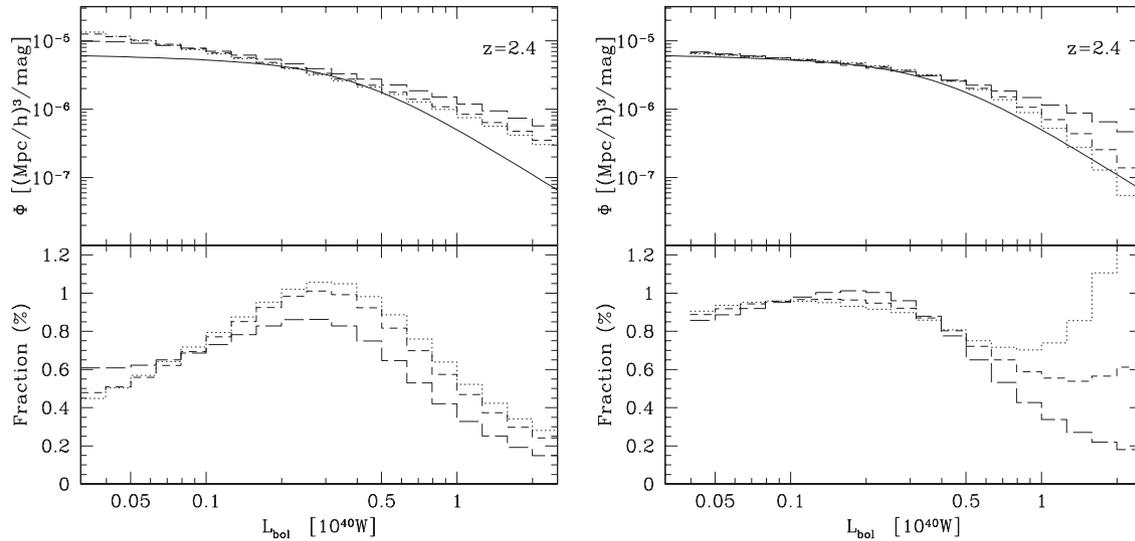


Figure 14. The duty cycle or active fraction of quasars. The solid line is the LF of quasars from Croom et al. (2004), as modified by Croton (2009), converting from b_j to M_i as $M_i(z=2) = M_{b_j} - 0.71$ (Richards et al. 2006) and from M_i to L_{bol} via $M_i(z=2) = 72.5 - 2.5 \log_{10}(L_{\text{bol}})$ (Shen et al. 2009). The left-hand plot shows, as the histograms in the upper panel, the predicted LFs (divided by 100) from the scaling relation model with $L_{\text{bol}} \sim v_{\text{peak}}^4$ and 25, 50 and 100 per cent scatter, assuming that all black holes are active. The lower panel shows the fraction that needs to be active at any given time in order to obtain the observed LF (solid line). This is the duty cycle. The right-hand plot shows the same parameters for a model with $L_{\text{bol}} \sim M_{\text{gal}}^{4/3}$ and in which we have randomly sampled haloes above $10^{12} h^{-1} M_{\odot}$ as described in the text. The possible quasars in our simulations are downsampled by these duty cycles to ensure a perfect match to the LF as described in the text.

Inverting this argument, we note that, at these high redshifts and masses, obtaining a reasonable $M_{\text{BH}}-M_{\text{h}}$ relation provides constraints on the possible occupancy distributions of quasars. In particular, because the halo mass function is much steeper than the galaxy stellar mass function, the typical stellar mass of a central galaxy drops steeply with decreasing halo mass – another way of stating that galaxy formation is inefficient in low-mass haloes. If black hole properties are set by the galactic potential rather than by halo properties we expect curvature in the $M_{\text{BH}}-M_{\text{h}}$ relation, which impacts how we interpret the duty cycle or the active quasar fraction.

4.2 Scaling-relation model

Another possibility is that the instantaneous luminosity of the quasar is drawn from a lognormal distribution with central value proportional to a power of the halo mass (or circular velocity) or the central galaxy mass (or dispersion). Physically such a model would arise if quasars radiate with a small range of Eddington ratios and are powered by black holes whose masses are tightly correlated with the mass or circular velocity of the galaxy or host halo (through e.g. black hole–bulge, bulge–galaxy and galaxy–halo correlations). The lognormal scatter is a combination of the dispersions in each of the relations connecting instantaneous luminosity to halo mass (see e.g. Croton 2009; Shen 2009, for recent examples of such models and Appendix B for further references).

We consider two examples here. First we relate the black hole properties to those of the host halo directly. We choose to use the peak circular velocity of the dark matter halo as our measure of halo size and take $\log L$ to be normally distributed around the (log of)

$$L_{\text{pk}} = L_0 \left(\frac{v_{\text{peak}}}{200 \text{ km s}^{-1}} \right)^4. \quad (13)$$

The normalization, L_0 , is set by matching the clustering amplitude at a given luminosity. In principle we can allow the power-law index

to vary.⁶ Unfortunately the range of luminosities we can probe observationally is relatively small and the differences in index are difficult to measure.

For a given L_0 and scatter the halo population defines a LF of possible quasars. The comparison of this to the observed LF of active quasars allows us to set the duty cycle, which will be luminosity (and hence halo mass) dependent. For each model we generate a mock catalogue drawn from the haloes of the simulations introduced in Section 4.1. We impose the duty cycle by randomly subsampling the possible quasars to ensure that the distribution matches the observed LF and then impose the magnitude limits to match the observed sample. We compute w_p and fit for r_0 as described previously.

Fig. 14 shows the duty cycle for the best-fitting model with $\log_{10} L_0 = 38.8$. The duty cycle peaks near 1 per cent at $\log_{10} L_{\text{bol}} \simeq 39.5$, corresponding to $M_i(z=2) \simeq -26.3$ or black hole masses of $(2-3) \times 10^8 M_{\odot}$. This is near the centre of our magnitude range and in our model corresponds to haloes of several times $10^{12} h^{-1} M_{\odot}$ or $v_{\text{peak}} \sim 300 \text{ km s}^{-1}$. Converting the duty cycle into an activity time is somewhat ill-defined. If we assume $t_Q = f_{\text{on}} t_{\text{H}}$, with t_{H} the Hubble time, we find $t_Q \sim 10^7 \text{ yr}$. These activity times are broadly consistent with those derived at $z \sim 0$, though since the Hubble time is significantly shorter the duty cycles are significantly higher. Also, note the luminosity/halo mass dependence of the duty cycle, which implies that in this model we need extra physics to describe the LF beyond simply major mergers with a fixed light curve.

A second option is to tie the black hole properties to the host galaxy, relating the galaxy properties to those of the halo by abundance matching. For the stellar masses and dispersions of interest here the velocity dispersion of the galaxy is proportional to the galaxy circular velocity, and we can take the black hole mass to

⁶ For example, an index of 4 would be appropriate to black holes whose growth is stopped by momentum-driven winds and 5 for those whose growth is stopped by luminosity-driven winds (Silk & Rees 1998).

scale as the fourth power of either quantity (as in the local universe; Tremaine et al. 2002). Then

$$L_{\text{pk}} = L_0 \left(\frac{\sigma_*}{200 \text{ km s}^{-1}} \right)^4 \propto M_*^p, \quad p \approx 1, \quad (14)$$

where the power-law index is approximately unity in both observations and numerical simulations (for representative examples, see Haring & Rix 2004; Hopkins et al. 2007b, and references therein). Fig. 14 shows that in this model the low-luminosity slope of the LF is in good agreement with the observations or the duty cycle has little luminosity dependence – much of the suppression of low-luminosity quasars can be accomplished by the same physics as is invoked to suppress star formation in lower mass galaxies. This cut-off in the occupancy to low halo masses tends to flatten the run of bias with luminosity (Fig. 15) in a manner similar to luminosity-dependent lifetime models, where there are also very few quasars in low-mass haloes.

At the high-luminosity end the suppression could be due to increasing inefficiency in feeding a black hole as cold-mode accretion becomes less effective (e.g. Sijacki et al. 2007) or the fact that it is harder to have a major galaxy merger (which would simultaneously drive gas to the centre and deepen the potential, allowing luminous quasar activity) when the stellar mass is increasing slowly with halo mass or due to curvature in the $M_{\text{BH}}-M_{\text{gal}}$ relation (e.g. Graham 2012). To illustrate the general point we have randomly subsampled the haloes above $10^{12} h^{-1} M_{\odot}$ by a fraction $2 \times 10^{12}/(M + 10^{12})$ to obtain the duty cycles shown in the lower right-hand panel of Fig. 14. While the agreement is by no means perfect, the general trends are in quite good agreement with observations, i.e. the luminosity dependence of the duty cycle is relatively weak.

The determining factors for the flatness of the bias–luminosity relation (Fig. 15) is the slope of the halo mass observable (i.e. luminosity) relation and the degree of scatter in that relation. For very high clustering amplitudes (as measured at high z) one obtains an upper limit on the scatter, which can be quite constraining for mod-

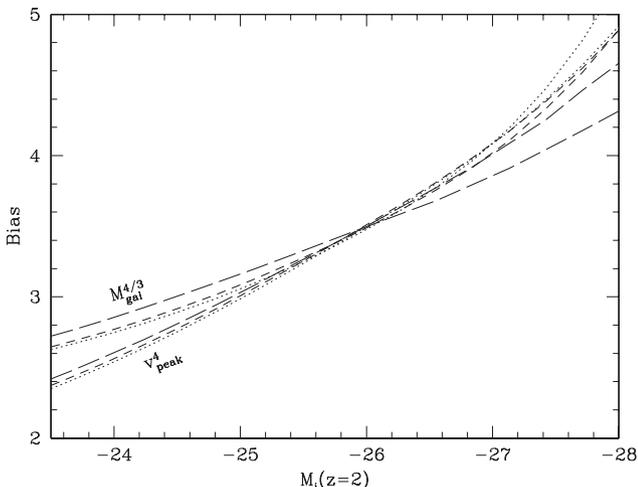


Figure 15. The predicted run of bias with magnitude for quasar samples 2 mag wide in $M_i(z=2)$ centred on the x -ordinate. There are two sets of curves, each with scatter in L_{obs} of $\sigma = 25$ (dotted), 50 (dashed) and 100 per cent (long dashed). The steeper curves correspond to the model with $L_{\text{bol}} \sim v_{\text{peak}}^4$, while the flatter curves correspond to $L_{\text{bol}} \sim M_{\text{gal}}^{4/3}$. In each case the active fraction has been adjusted to match the observed LF, as in Fig. 14. The dependence on the amount of scatter assumed is small because of the 2 mag wide bin that is taken for each sample to increase statistics.

els, but at intermediate z quite a large degree of scatter is allowed (White et al. 2008; Shankar et al. 2010). The scatter can arise from a number of sources including luminosity-dependent lifetime, but also stochasticity in the relations between halo mass and galaxy mass, galaxy mass and central potential well depth, potential well depth and black hole mass, and black hole mass and optical luminosity. For reasonable values of halo-observable slope, luminosity bin width and stochasticity, one obtains quite flat $b(L)$ – whether or not there is a sharp cut-off in the halo mass distribution. Thus while it is definitely plausible that quasar lifetime is luminosity dependent, it is not strictly required by the current clustering data at these redshifts.

The run of bias with halo mass becomes quite shallow at the low-mass end. If quasar luminosity is additionally affected by the inefficiency of galaxy formation in lower mass haloes, we expect to see a luminosity-dependent quasar bias primarily at higher luminosities (and redshifts). Unfortunately this is where quasars become increasingly rare, which argues that cross-correlation may be the best means of obtaining a strong clustering signal (e.g. Shen et al. 2009). We explore this issue briefly in Fig. 16, where we show the number of quasar pairs we expect to see in a complete, 10^4 deg^2 survey with separation $< 20 h^{-1} \text{ Mpc}$ in bins of redshift and magnitude. In the black areas, with more than 10^3 pairs, a solid detection of clustering from the auto-correlation should be possible. In the other regions (or for smaller or less complete surveys) cross-correlation with other quasars or a different tracer is likely required.

By contrast the bias at the low-luminosity end is an effective discriminator between models where quasar luminosity depends on halo mass or galaxy mass. Unfortunately BOSS is unable to probe this range of quasar luminosities. Future surveys which probe fainter magnitudes may be able to settle this question. It may also be possible to cross-correlate faint, photometrically selected quasars

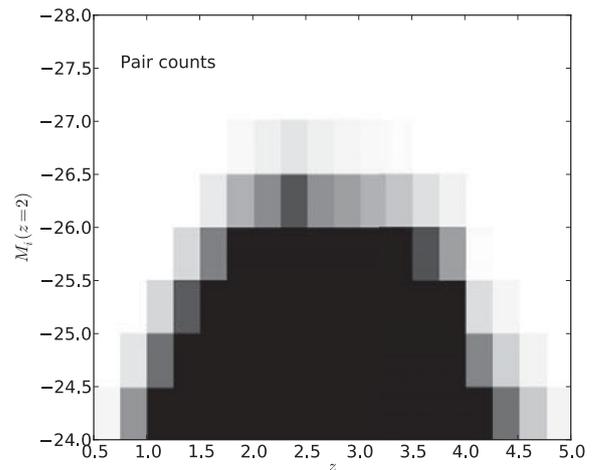


Figure 16. The number of quasar pairs with separation $< 20 h^{-1} \text{ Mpc}$ in the redshift and magnitude bins shown, as predicted by the LF of Croom et al. (2004) as modified by Croton (2009). We assume that all quasars in the redshift and magnitude bins are observed over 10^4 deg^2 of sky. Reducing the sky area reduces the pair counts linearly, while finding only a fraction of the quasars reduces the counts quadratically. The black regions indicate more than 10^3 pairs, the grey regions indicate between 10^2 and 10^3 pairs and the white regions indicate fewer than 10^2 pairs. A solid detection of clustering from the auto-correlation should be possible in the black regions, while in the white regions it is likely necessary to cross-correlate with more numerous samples or to make wider bins in redshift or magnitude.

with the brighter spectroscopic quasars from the BOSS sample (using e.g. the methods in Myers, White & Ball 2009).

4.3 Redshift evolution

In Fig. 17 we compare the clustering of our quasar sample to that of other well-studied objects at $z \sim 2$. The quasar clustering amplitude is similar to that of submillimetre galaxies, suggesting that they live in similar mass haloes, and stronger than the typical star-forming population. This is consistent with an evolutionary picture in which a merger triggers a massive starburst that creates a submillimetre galaxy and which is then quenched by the formation of a bright quasar (e.g. Alexander & Hickox 2012). The descendants of our quasar hosts will have comparable clustering amplitudes to the quasars themselves, indicating that they will likely evolve into massive, luminous early-type galaxies at low redshift.

To better understand the possible fate of quasar host haloes we employ another high-resolution N -body simulation which allows us to track haloes and subhaloes down to $z = 0$ (for details, see White, Cohn & Smit 2010). We select all haloes at $z = 2.4$ which are central subhaloes of haloes which lie within an octave (i.e. factor of 2) in mass centred on $2 \times 10^{12} h^{-1} M_{\odot}$ ($>10^4$ particles). Essentially all of these haloes are the most massive in their local environment. One-quarter of the hosts fall into a large halo (becoming a satellite) and then lose more than 99.9 per cent of their mass (falling below the resolution limit of the simulation) without merging with the central galaxy or another satellite subhalo. To the extent that this process is well resolved, the stars in any galaxies hosted by these subhaloes would likely contribute to an intra-halo light component, while the black holes would form a freely floating component or be associated with highly stripped satellites. Of the remaining 75 per cent of the subhaloes, three-quarters are the most massive progen-

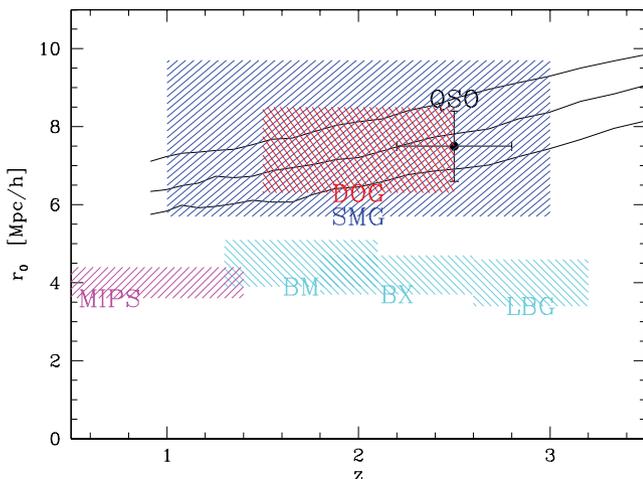


Figure 17. Comparison of our measured correlation length (square with error bars) to that of other objects at $z \sim 2$. In each case the horizontal error bar or width of the shaded region shows the redshift range over which the measurement is performed, while the vertical error bar or height of the shaded region shows the $\pm 1\sigma$ region. The other data sets are mid-IR-selected star-forming galaxies (MIPS; Gilli et al. 2007), BM, BX and LBG star-forming galaxies (Adelberger et al. 2005), submillimetre galaxies (SMG; Hickox et al. 2012) and dust-obscured galaxies (DOG; Brodwin et al. 2008). The lines show the clustering of haloes with peak circular velocity 250, 300 and 350 km s^{-1} as measured from the simulations described in Section 4.1. See Krumpke, Miyaji & Coil (2010), Figs 9–11, for a similar comparison at lower z .

itor in all their subsequent mergers and remain central galaxies to $z = 0$. The remaining quarter become satellites which survive to $z = 0$ inside more massive haloes. As a whole the population inhabits $z \simeq 0$ haloes over a broad range of masses ($2 \times 10^{12} - 10^{15} h^{-1} M_{\odot}$) peaking at $(1-2) \times 10^{13} h^{-1} M_{\odot}$. Haloes of this group scale host galaxies of a few L_* today, a population dominated by elliptical galaxies. This wide diversity of outcomes is reminiscent of the varied fates of $z \sim 2$ star-forming galaxies (e.g. Conroy et al. 2008). As constraints on the stellar masses of galaxies within haloes become tighter, comparison of the stellar masses of quasar hosts with that of their $z \sim 0$ descendants will put constraints on the star formation history of these objects.

The evolution of clustering with time can place strong constraints on how episodic quasar activity can be. As emphasized by Croom et al. (2005), if the typical host of quasars does not evolve significantly with redshift, then quasars cannot be repeated bursts of the same black hole because such black holes would live in haloes which grow in mass as the Universe evolves. However, the quantitative strength of this statement, and the allowed fraction of objects which could burst more than once, is difficult to assess. One issue is the size of the observational errors, the other is that higher redshift samples typically probe more massive black holes than lower redshift samples (as emphasized e.g. by Hopkins et al. 2007a).

For example, the highest redshift clustering measurement comes from Shen et al. (2007), at $z \simeq 4$. Their result is consistent with host haloes of $(2-10) \times 10^{12} h^{-1} M_{\odot}$. Such hosts would grow in mass by a factor of approximately 4 between $z = 4$ and 2.4. Similarly, the progenitors of our $2 \times 10^{12} h^{-1} M_{\odot}$, $z = 2.4$ haloes are approximately 4 times less massive at $z = 4$. This suggests on average an order-of-magnitude mismatch in halo masses, but with a large error. Assuming no evolution in Eddington ratio between $z = 4$ and 2.4 the quasars in the Shen et al. (2007) sample are a factor of about 5 brighter than the BOSS quasars. If the power-law index of the $M_{\text{BH}}-M_{\text{h}}$ relation is close to unity, there is little tension from clustering⁷ in assuming quasars episodically burst.

The time span between $z = 4$ and 2.4 in our adopted cosmology is $1 h^{-1} \text{ Gyr}$. The Universe is another $1 h^{-1} \text{ Gyr}$ older at $z \simeq 1.5$. Taking r_0 at $z = 1.57$ from Ross et al. (2009) and converting it to a host halo mass we obtain $\sim 5 \times 10^{12} h^{-1} M_{\odot}$. The descendants of our $2 \times 10^{12} h^{-1} M_{\odot}$, $z = 2.4$ haloes are approximately twice as massive at $z = 1.5$, again leading to little tension in a model with episodic outbursts.

We can see the issues most clearly in Fig. 12. If haloes moved with the same large-scale velocities as the dark matter and never merged, the decrease of their bias would approximately cancel the increased clustering of the matter (Fry 1996; see fig. 1 of White et al. 2007). For a highly biased sample such passive evolution corresponds to almost constant clustering strength, for slightly less-biased objects the clustering strength grows slowly with time. This result is shown as the dashed line in Fig. 12, corresponding to objects with $b(z = 0) = 1.8$. Merging of haloes and including a finite range of halo masses alter the details of this evolution, but keep the sense unaltered. By contrast, haloes of a fixed mass predict a clustering strength which drops slowly with time – shown as the dotted line in Fig. 12 for haloes of $2 \times 10^{12} h^{-1} M_{\odot}$. Again, a more realistic scenario with a finite range of halo masses has the same sense of the evolution. The current data are not in strong conflict with either scenario. If a random fraction of quasars repeats at later

⁷ The dramatic decrease in quasar numbers to higher redshift does impose constraints.

epochs while new quasars always appear in haloes of a fixed mass, the measured clustering resembles a sample with bias $f_{\text{repeat}}b_{\text{repeat}} + (1 - f_{\text{repeat}})b_{\text{new}}$. The measured values of r_0 are consistent with being roughly constant over $1 < z < 3$. For this reason it is difficult to put a strong upper limit on the fraction of quasars that turn on more than once or the maximum number of times a given quasar can burst.

5 CONCLUSIONS

We have performed real- and redshift-space clustering measurements of a uniform subsample of quasars observed by the BOSS. These quasars lie in the redshift range $2.2 < z < 2.8$ where there has previously been a gap in coverage due to the fact that the quasar and stellar loci cross and the difficulty in targeting quasars in large numbers to faint fluxes. We detect clustering at high significance in both real space and redshift space for the entire sample and subsamples split by redshift and luminosity (see Table 3). We do not detect a luminosity or redshift dependence of the clustering strength, although our sensitivity to this dependence is weak due to the limited dynamic range in both variables in our sample.

The two-point correlation functions are consistent with power laws over the range of scales measured, with an underlying real-space clustering of the form $\xi(r) = (r_0/r)^2$. The correlation length, r_0 , does not appear to evolve strongly over the redshift range $z \simeq 3$ to 1. This result is consistent with passive evolution of a highly biased population, although this interpretation is by no means unique. Our results are consistent with quasars living in haloes of typical mass $10^{12} h^{-1} M_{\odot}$ at $z \simeq 2.4$, in line with expectations from earlier surveys. The measured bias and space density of quasars can be used to infer their duty cycle (Cole & Kaiser 1989). For our best-fitting models the duty cycle peaks at 1 per cent, implying an activity time of $\sim 10^7$ yr. This time is comparable to the activity times inferred for quasars at lower redshift, although the Hubble time at $z = 2.4$ is shorter than at lower z and hence the active fraction is larger.

The typical host halo mass is similar to the inferred hosts of submillimetre galaxies and is more massive than the inferred hosts of typical star-forming galaxies at the same redshift. This interpretation is in turn consistent with an evolutionary picture in which a massive starburst creates a submillimetre galaxy and is quenched by the formation of a bright quasar. While the typical descendant of the haloes that host BOSS quasars is likely to host a luminous, elliptical galaxy at the present time, we find a wide diversity in descendants in N -body simulations.

Using abundance matching to infer the properties of quasar host galaxies we find evidence for evolution in the $M_{\text{BH}}-M_{\text{gal}}$ relation in the sense that black holes must be ≈ 5 times more massive at fixed galaxy mass at $z = 2.4$ than at $z \simeq 0$. We find that the predictions for how quasar activity and clustering (bias) depend on luminosity differ depending on whether we take as our fundamental relationship a black hole–halo correlation or a black hole–galaxy correlation. This is because the efficiency of galaxy formation is strongly (halo) mass dependent for the haloes of interest at these redshifts, leading to strong curvature in the $M_{\text{gal}}-M_{\text{h}}$ relation. In either scenario a modest scatter between halo or galaxy mass and observed quasar luminosity (arising, for example, from a combination of scatters in the black hole–bulge, bulge–galaxy and galaxy–halo correlations and the Eddington ratio) leads to a shallow dependence of clustering on quasar luminosity, as observed.

Future surveys of quasars which probe different regions of the luminosity–redshift plane will inform models of quasar formation and evolution. In this regard BOSS continues to measure quasar

redshifts, and we expect that the number of quasars in the luminosity and redshift range discussed here will be more than doubled by the end of the survey. It may be possible to incorporate the additional, BONUS quasars through cross-correlation or to cross-correlate spectroscopic and photometric quasar samples to better allow us to break the sample by luminosity, spectral or radio properties. In addition BOSS is measuring redshifts for a large sample of $z > 3$ quasars, and analysis of those data will be crucial in understanding the early phases of quasar growth.

ACKNOWLEDGMENTS

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation and the US 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.

The analysis made use of the computing resources of the National Energy Research Scientific Computing Center, the Shared Research Computing Services Pilot of the University of California and the Laboratory Research Computing project at Lawrence Berkeley National Laboratory.

MW was supported by the NSF and NASA. ADM is a research fellow of the Alexander von Humboldt Foundation of Germany. JM-E is supported by Spanish grants AYA2009-09745 and PR2011-0431. JB was partially supported by a NASA Hubble Fellowship (HST-HF-51285.01).

REFERENCES

- Abazajian K. et al., 2004, *AJ*, 128, 502
- Adelberger K. L., Steidel C. C., Pettini M., Shapley A. E., Reddy N. A., Erb D. K., 2005, *ApJ*, 619, 697
- Adelman-McCarthy J. et al., 2007, *ApJS*, 172, 634
- Aihara H. et al., 2011, *ApJS*, 193, 29
- Alexander D. M., Hickox R. C., 2012, *New Astron. Rev.*, preprint (arXiv:1112.1949)
- Anderson L. et al., 2012, *MNRAS*, submitted
- Behroozi P. S., Conroy C., Wechsler R. H., 2010, *ApJ*, 717, 379
- Blanton M. R., Lin H., Lupton R. H., Maley F. M., Young N., Zehavi I., Loveday J., 2003, *AJ*, 125, 2276
- Bonoli S., Marulli F., Springel V., White S. D. M., Branchini E., Moscardini L., 2009, *MNRAS*, 396, 423
- Booth C. M., Schaye J., 2010, *MNRAS*, 405, L1
- Bovy J., Hogg D. W., Roweis S. T., 2009, *Ann. Applied Statistics*, 5, 1657
- Bovy J. et al., 2011, *ApJ*, 729, 141
- Bovy J. et al., 2012, *ApJ*, 749, 41
- Brodwin M. et al., 2008, *ApJ*, 687, L65
- Bromley J. M., Somerville R. S., Fabian A. C., 2004, *MNRAS*, 350, 456

- Budavári T. et al., 2001, *AJ*, 122, 1163
 Cao X., 2010, *ApJ*, 725, 388
 Carlberg R. G., 1990, *ApJ*, 350, 505
 Cattaneo A., Haehnelt M. G., Rees M. J., 1999, *MNRAS*, 308, 77
 Ciotti L., Ostriker J. P., 1997, *ApJ*, 487, 105
 Ciotti L., Ostriker J. P., 2001, *ApJ*, 551, 131
 Cohn J. D., White M., 2005, *Astropart. Phys.*, 24, 316
 Cole S., Kaiser N., 1989, *MNRAS*, 237, 1127
 Conroy C., Shapley A., Tinker J. L., Santos M. R., Lemson G., 2008, *ApJ*, 679, 1192
 Croom S. M. et al., 2004, *MNRAS*, 349, 1397
 Croom S. M. et al., 2005, *MNRAS*, 356, 415
 Croom S. M. et al., 2009, *MNRAS*, 392, 19
 Croton D. J., 2009, *MNRAS*, 394, 1109
 Croton D. J. et al., 2006, *MNRAS*, 365, 11
 da Angela J. et al., 2008, *MNRAS*, 383, 565
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 De Graf C., Di Matteo T., Khandai N., Croft R., Lopez J., Springel V., 2011, preprint (arXiv:1107.1254)
 Decarli R., Falomo R., Treves A., Labita M., Kotilainen J. K., Scarpa R., 2010, *MNRAS*, 402, 2453
 Di Matteo T., Khandai N., DeGraf C., Feng Y., Croft R. A. C., Lopez J., Springel V., 2012, *ApJ*, 745, L29
 Doi M. et al., 2010, *AJ*, 139, 1628
 Efron B., Gong G., 1983, *American Statistician*, 37, 36
 Efstathiou G., Rees M. J., 1988, *MNRAS*, 230, 5
 Eisenstein D. J. et al., 2011, *AJ*, 142, 72
 Fakhouri O., Ma C.-P., 2009, *MNRAS*, 394, 1825
 Fan X., 1999, *AJ*, 117, 2528
 Fanidakis N. et al., 2012, *MNRAS*, 419, 2797
 Ferrarese L., 2002, *ApJ*, 578, 90
 Fine S. et al., 2006, *MNRAS*, 373, 613
 Fry J. N., 1996, *ApJ*, 461, L65
 Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, *AJ*, 111, 1748
 Gilli R. et al., 2007, *A&A*, 475, 83
 Gorski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
 Graham A., 2012, *ApJ*, 746, 113
 Granato G. L., de Zotti G., Silva L., Bressan A., Luigi D., 2004, *ApJ*, 600, 580
 Gunn J. E. et al., 1998, *AJ*, 116, 3040
 Gunn J. E. et al., 2006, *AJ*, 131, 2332
 Haiman Z., Hui L., 2001, *ApJ*, 547, 27
 Haiman Z., Loeb A., 1998, *ApJ*, 503, 505
 Haiman Z., Ciotti L., Ostriker J. P., 2004, *ApJ*, 606, 763
 Haring N., Rix H., 2004, *ApJ*, 604, L89
 Hartwick F. D. A., Schade D., 1990, *ARA&A*, 28, 437
 Hawkins E. et al., 2003, *MNRAS*, 346, 78
 Hennawi J. F. et al., 2006a, *AJ*, 131, 1
 Hennawi J. F. et al., 2006b, *ApJ*, 651, 61
 Hewett P. C., Wild V., 2010, *MNRAS*, 405, 2302
 Hickox R. et al., 2011, *ApJ*, 731, 117
 Hickox R. et al., 2012, *MNRAS*, 421, 284
 Hopkins P. F., Somerville R. S., Hernquist L., Cox T. J., Robertson B., Li Y., 2006, *ApJ*, 652, 864
 Hopkins P., Lidz A., Hernquist L., Coil A. L., Myers A. D., Cos T. J., Spergel D. N., 2007a, *ApJ*, 662, 110
 Hopkins P., Hernquist L., Cox T. J., Robertson B., Krause E., 2007b, *ApJ*, 669, 45
 Hopkins P. F., Hernquist L., Cox T. J., Keres D., 2008, *ApJS*, 175, 356
 Hopkins P. et al., 2010, *ApJ*, 724, 915
 Kaiser N., 1987, *MNRAS*, 227, 1
 Kauffmann G., Haehnelt M., 2000, *MNRAS*, 311, 576
 Kauffmann G., Haehnelt M., 2002, *MNRAS*, 332, 529
 Kormendy J., Richstone D., 1995, *ARA&A*, 33, 581
 Krumpke M., Miyaji T., Coil A. L., 2010, *ApJ*, 713, 558
 Krumpke M., Miyaji T., Coil A. L., Aceves H., 2012, *ApJ*, 746, 1
 Lacey C., Cole S., 1993, *MNRAS*, 262, 627
 Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
 Li C., Kauffmann G., Jing Y. P., White S. D. M., Borner G., Cheng F. Z., 2006, *MNRAS*, 368, 21
 Lidz A., Hopkins P. F., Cox T. J., Hernquist L., Robertson B., 2006, *ApJ*, 641, 41
 Lupton R., Gunn J. E., Ivezić Z., Knapp G. R., Kent S., 2001, in Hamden F. R., Jr, Primini F. A., Payne H. E., eds, *ASP Conf. Ser. Vol. 238, Astronomical Data Analysis Software and Systems X*. Astron. Soc. Pac., San Francisco, p. 269
 Lynden Bell D., 1969, *Nat*, 223, 690
 Malbon R. K., Baugh C. M., Frenk C. S., Lacey C. G., 2007, *MNRAS*, 382, 1394
 Martini P., Weinberg D. H., 2001, *ApJ*, 547, 12
 Marulli F., Crociani D., Volonteri M., Branchini E., Moscardini L., 2006, *MNRAS*, 368, 1269
 Matarrese S., Coles P., Lucchin F., Moscardini L., 1997, *MNRAS*, 286, 115
 Merloni A. et al., 2010, *ApJ*, 708, 137
 Monaco P., Fontanot F., Taffoni G., 2007, *MNRAS*, 375, 1189
 Moster B. P. et al., 2010, *ApJ*, 710, 903
 Myers A. D. et al., 2006, *ApJ*, 638, 622
 Myers A. D., Brunner R. J., Nichol R. C., Richards G. T., Schneider D. P., Bahcall N. A., 2007a, *ApJ*, 658, 85
 Myers A. D., Brunner R. J., Nichol R. C., Richards G. T., Schneider D. P., Bahcall N. A., 2007b, *ApJ*, 658, 99
 Myers A. D., White M., Ball N., 2009, *MNRAS*, 399, 2279
 Nandra K. et al., 2007, *ApJ*, 660, L11
 Natarajan P., 2012, in *Proc. XVth Congress of Philosophy & Foundations of Science*, to be published by the American Institute of Physics, preprint (arXiv:1105.4902)
 Padmanabhan N. et al., 2008, *ApJ*, 674, 1217
 Padmanabhan N., White M., Norberg P., Porciani C., 2009, *MNRAS*, 397, 1862
 Peng C. Y., Impey C. D., Rix H.-W., Falco E. E., Keeton C., Kochanek C. S., Lehar J., McLeod B. A., 2006, *New Astron. Rev.*, 50, 689
 Percival W. J., Scott D., Peacock J. A., Dunlop J. S., 2003, *MNRAS*, 338, L31
 Pier J. R., Munn J. A., Hindsley R. B., Hennessy G. S., Kent S. M., Lupton R. H., Ivezić Z., 2003, *AJ*, 125, 1559
 Porciani C., Norberg P., 2006, *MNRAS*, 371, 1824
 Porciani C., Magliocchetti M., Norberg P., 2004, *MNRAS*, 355, 1010
 Reid B. A. et al., 2012, *MNRAS*, preprint (arXiv:1203.6641)
 Richards G. T. et al., 2001a, *AJ*, 121, 2308
 Richards G. T. et al., 2001b, *AJ*, 122, 1151
 Richards G. T. et al., 2002a, *AJ*, 123, 2945
 Richards G. T. et al., 2002b, *AJ*, 124, 1
 Richards G. T. et al., 2006, *AJ*, 131, 2766
 Richards G. T. et al., 2009, *ApJS*, 180, 67
 Richards G. T. et al., 2011, *AJ*, 141, 167
 Ross N. P. et al., 2007, *MNRAS*, 381, 573
 Ross N. P. et al., 2009, *ApJ*, 697, 1634
 Ross N. P. et al., 2012, *ApJS*, 199, 3
 Salpeter E. E., 1964, *ApJ*, 140, 796
 Schneider D. P. et al., 2010, *AJ*, 139, 2360
 Shankar F., 2009, *New Astron. Rev.*, 53, 57
 Shankar F., Weinberg D. H., Shen Y., 2010, *MNRAS*, 406, 1959
 Shanks T., Croom S. M., Fine S., Ross N. P., Sawangwit U., 2011, *MNRAS*, 416, 650
 Shen Y., 2009, *ApJ*, 704, 89
 Shen Y. et al., 2007, *AJ*, 133, 2222
 Shen Y. et al., 2009, *ApJ*, 697, 1656
 Shlosman I., Frank J., Begelman M. C., 1989, *Nat*, 338, 45
 Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
 Silk J., Rees M. J., 1998, *A&A*, 331, L1
 Silverman J. D. et al., 2008, *ApJ*, 679, 118
 Smith J. A. et al., 2002, *AJ*, 123, 2121

- Smith R. E. et al., 2003, MNRAS, 341, 1311
 Springel V., Di Matteo T., Hernquist L., 2005, MNRAS, 361, 776
 Swanson M. E. C., Tegmark M., Hamilton A. J. S., Hill J. C., 2008, MNRAS, 387, 1391
 Tegmark M. et al., 2004, ApJ, 606, 702
 Tremaine S. et al., 2002, MNRAS, 574, 740
 Vanden Berk D. E. et al., 2001, AJ, 122, 549
 Volonteri M., Stark D. P., 2011, MNRAS, 417, 2085
 Volonteri M., Haardt F., Madau P., 2003, ApJ, 582, 559
 Weedman D., 1986, Quasar Astronomy. Cambridge Univ. Press, Cambridge
 Weinstein M. A. et al., 2004, ApJS, 155, 243
 Wetzel A., Cohn J. D., White M., 2009, MNRAS, 394, 2182
 White M., 2002, ApJS, 579, 16
 White M., Zheng Z., Brown M. J. I., Dey A., Jannuzi B. T., 2007, ApJ, 655, L69
 White M., Martini P., Cohn J. D., 2008, MNRAS, 390, 1179
 White M., Cohn J. D., Smit R., 2010, MNRAS, 408, 1818
 White M. et al., 2011, ApJ, 728, 126
 Wyithe J. S. B., Loeb A., 2002, ApJ, 581, 886
 Wyithe J. S. B., Loeb A., 2003, ApJ, 595, 614
 York D. G. et al., 2000, AJ, 120, 1579

APPENDIX A: REDSHIFT ERRORS

At the signal-to-noise ratio and redshift at which BOSS is working, it is difficult to obtain precise redshifts for quasars. Emission lines, such as Mg II which are good redshift indicators and which can be used at lower redshift, have redshifted into a relatively noisy part of the spectrum or off of the device altogether. 1-h integrations on a 2.5-m telescope make it difficult to measure redshifts for quasars with weak lines.

The BOSS pipeline measures quasar redshifts by fitting their spectra to a set of principal component analysis (PCA) templates (Aihara et al. 2011) plus a cubic polynomial to allow for changes in continuum slope. The reduced χ^2 versus redshift is mapped in steps of $\Delta \log_{10}(\lambda) = 10^{-4}$ from $z = 0.0033$ to 7 and the template fit with the best reduced χ^2 is selected as the redshift. In addition, redshifts are computed by fitting any lines in two groups (forbidden and allowed). A comparison of redshifts determined from different lines (Hennawi et al. 2006b; Shen et al. 2007), and visual inspections, suggests that the automatic redshifts are good to $\Delta z/(1+z) \simeq 0.003$. At $z = 2.5$ this corresponds to an error in the line-of-sight distance of roughly $10 h^{-1}$ Mpc (comoving), which is significant compared to the correlation length of quasar clustering. We are attempting to improve our quasar redshift determination, but for now we simply account for the residual line-of-sight smearing induced by redshift errors in our fitting.

In the limit $b \gg 1$ the redshift-space halo correlation function is approximately isotropic and a power law. If the redshift errors on quasars are uncorrelated and Gaussian distributed with fixed amplitude σ_z the observed correlation function is

$$\xi_{\text{obs}}(s) = \frac{1}{2} \int \frac{dz}{s} \int \frac{dZ}{\sqrt{2\pi}\sigma} \xi\left(\sqrt{s^2 - z^2}, Z\right) e^{-(z-Z)^2/2\sigma^2}, \quad (\text{A1})$$

where $\sigma = \sqrt{2}(c\sigma_z)/H(z)$. The integral, ξ_{obs} , divided by the input power law is what we refer to in the main text as $F(s)$.

Fig. A1 compares this model, with a power-law correlation function of slope 1.8, to haloes from N -body simulations with applied Gaussian line-of-sight velocity errors. As expected, the agreement is excellent. In this test the redshift errors were all drawn from a Gaussian of the same dispersion. In the observational sample we might reasonably expect the errors to depend on the properties of each quasar. In this situation we should interpret σ as a pair-weighted,

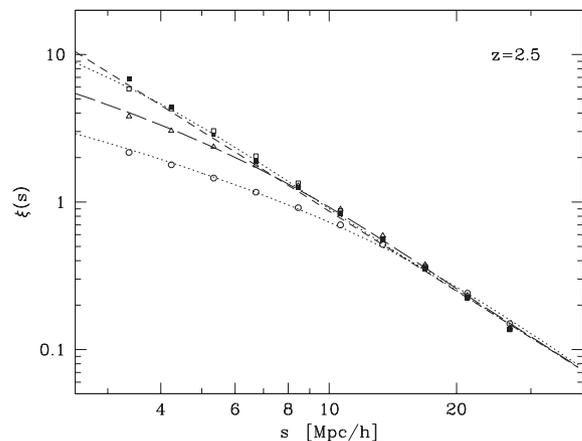


Figure A1. The effects of redshift errors on the redshift-space (monopole) correlation function. The dashed line shows a power-law correlation function with slope 1.8, and the other lines are for equation (A1) with per-object redshift errors corresponding to 2.5, 5 and $10 h^{-1}$ Mpc. The points are for mock quasars as in Section 4.2 with Gaussian errors added to the line-of-sight velocities.

‘effective’, redshift error. The above tests, plus the clustering measurements themselves, are consistent with a per quasar redshift error corresponding to $10 h^{-1}$ Mpc, and we use that as our fiducial value throughout.

APPENDIX B: QUASAR MODEL

Models of the quasar phenomenon come in several basic flavours, but are in fact all quite similar. The majority of models assume that quasar activity occurs due to the major merger of two gas-rich galaxies, since this scenario provides the rapid and violent event needed to funnel fuel to the centre of the galaxy (e.g. via the bars-within-bars instability; Shlosman, Frank & Begelman 1989) and feed the central engine while at the same time giving a connection between black hole fuelling and the growth of a spheroidal stellar component. If black hole growth is feedback limited, it is only a rapidly growing potential well that can host a rapidly accreting hole. A notable exception is the models of Ciotti & Ostriker (1997, 2001) which postulate that the fuel is funnelled to the centre by thermal instabilities, which provide the rapid growth of the spheroidal component necessary for black hole growth. As we shall see, both sets of models can predict very similar halo occupancy.

Some models implement the physics directly in numerical simulations which attempt to track the hydrodynamics of the gas, with subgrid models for quasar and star formation and the associated feedback (Sijacki et al. 2007; Hopkins et al. 2008; De Graf et al. 2011). Other models work at the level of dark matter haloes, but follow the same physics – in simplified form – semi-analytically (Cattaneo et al. 1999; Kauffmann & Haehnelt 2000, 2002; Volonteri, Haardt & Madau 2003; Bromley, Somerville & Fabian 2004; Granato et al. 2004; Croton et al. 2006; Malbon et al. 2007; Monaco, Fontanot & Taffoni 2007; Bonoli et al. 2009; Fanidakis et al. 2012). Even more idealized are models which are built upon dark matter haloes but use scaling relations or convenient functional forms to relate the quasar properties to those of their host haloes (Efsthathiou & Rees 1988; Carlberg 1990; Wyithe & Loeb 2002, 2003; Haiman, Ciotti & Ostriker 2004; Lidz et al. 2006; Marulli et al. 2006; Croton 2009; Shen 2009; Booth & Schaye 2010). A final level of abstraction is to simply provide a stochastic recipe for populating dark matter

haloes with quasars which is tuned to reproduce the observations as best as possible while not attempting to follow the underlying physics (Porciani et al. 2004; White et al. 2008; Padmanabhan et al. 2009; Shankar et al. 2010; Volonteri & Stark 2011; Krumpe et al. 2012; Kirkpatrick et al., in preparation).

The modelling is simplified by several facts. Quasars are rare, their activity times are short and the fraction of binary quasars is small (Hennawi et al. 2006a; Myers et al. 2007b). The hydrodynamic and semi-analytic models tend to reproduce the observed, $z = 0$ relation between black hole mass and halo, which is used as input to the scaling models. The models agree on the level of scatter in the relation (roughly a factor of 2) and in broad brush on the evolution in the amplitude and slope of this relation with redshift. Some models invoke quasar feedback-limited accretion explicitly, while others achieve the same scaling relations without such a limit – e.g. by coupling the mechanisms by which bulges and black holes grow. At these masses and redshifts the satellite galaxy fraction is tiny, so the halo to stellar mass relation is set by abundance matching and any model which reproduces the galaxy stellar mass function will reproduce this relation. At high redshift bright quasars radiate near Eddington, so the models predict similar halo occupancies.

The probability that a halo will undergo a major merger in a short redshift interval is only weakly dependent on the mass of the halo (Lacey & Cole 1993; Percival et al. 2003; Cohn & White 2005; Fakhouri & Ma 2009; Wetzel, Cohn & White 2009; Hopkins et al. 2010), i.e. the mass function of such haloes is almost proportional to the mass function of the parent population. Similarly, the clustering properties of recently merged haloes are similar to a random sample of the population with the same mass distribution (Percival et al. 2003; Wetzel et al. 2009). Thus in any interval, Δz , the fraction of haloes of mass M_h which undergo a quasar event is almost independent of M_h and z and can be regarded as a random selection. This makes it difficult to infer that quasars arise from mergers simply from their large-scale clustering, but also implies that for the purposes of modelling the one- and two-point functions of the quasar population it is sufficient to specify the halo occupation of the parent population (more complex models may be needed if correlations between quasars and properties of e.g. galaxies were required).

To determine whether a quasar candidate makes it into any sample, it is necessary to relate the observed luminosity to the peak luminosity which is determined by M_{BH} . This is done either by specifying a light curve (e.g. a power law, a power law with a varying slope or an exponential) or directly $P(L|M_h)$ using the fact that

the triggering rate is understood. For constant triggering rate a light curve $L \sim t^{-\gamma}$ implies $P(L) \sim L^{-(\gamma+1)/\gamma}$, so a wide range of γ maps to a narrow range in $P(L)$ index. For the high-luminosity thresholds we only see objects very near their peak brightness. Taking into account the factor of ~ 2 scatter in M_{BH} at fixed bulge mass, we expect $P(L|M_h)$ to be roughly lognormal with a similar width (it can be slightly broader due to variation in the Eddington ratio at peak). As we probe lower luminosities we are more likely to see older, more massive black holes leading to a low- L tail in $P(L|M_h)$ and more lower Eddington-ratio objects (Lidz et al. 2006; Shen 2009; Cao 2010). Unless we cut off the probability that a halo hosts a quasar at low and high halo mass, we will overproduce low- and high- L sources (Lidz et al. 2006; Croton 2009; Shen 2009). In the physical models these limits occur due to lack of fuel in low-mass haloes and the inability of gas to cool in high-mass haloes (e.g. Cattaneo et al. 1999). As discussed in Section 4.2, models which match the M_h – M_* relation for galaxies tend to roughly match the required quasar suppression.

We are interested in the probability that a given halo hosts a quasar in our sample, e.g. $f_{\text{on}}(M_h) \propto P(> L_{\text{min}}|M_h)$. From the arguments above we expect this relation to be an approximately lognormal function which is asymmetric towards high M_h at low L . It is quite difficult to put constraints on the detailed form of this function using LF and clustering measurements.

We argued above that the steepness of the halo mass function and the high bias of quasars implies that the quasar satellite fraction is small. The LF thus provides a constraint on the duty cycle. (This extra degree of freedom reduces the ability of large-scale structure measurements to constrain the halo occupancy, compared to modelling galaxies.) The steeply falling mass function also implies that the number of quasars hosted in very massive haloes is small regardless of the occupancy statistics of such haloes. On scales larger than the virial radius of the typical quasar host halo (i.e. 200–300 h^{-1} kpc) the two-point function is dominated by pairs of quasars in different haloes, and thus primarily measures the quasar-weighted halo bias which allows us to infer the mean, quasar-weighted halo mass. This result remains true for cross-correlation studies too, provided that we work on scales larger than the virial radius of the quasar hosts. On smaller scales the amplitude and slope of the correlation function allow us to measure a combination of the satellite fraction of quasars and the low-mass cut-off.

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