Constraining the quasar population with the broad-line width distribution

S. Fine,[1]⋆ S. M. Croom,[1] P. F. Hopkins,[2] L. Hernquist,[2] J. Bland-Hawthorn,[1]
M. Colless,[3] P. B. Hall,[4] L. Miller,[5] A. D. Myers,[6] R. Nichol,[7] K. A. Pimbblet,[8]
N. P. Ross,[9,10] D. P. Schneider,[10] T. Shanks[9] and R. G. Sharp[3]

[1]*School of Physics, University of Sydney, NSW 2006, Australia*
[2]*Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA*
[3]*Anglo-Australian Observatory, PO Box 296, Epping, NSW 1710, Australia*
[4]*Department of Physics and Astronomy, York University, 4700 Keele Street, Toronto, ON M3J 1P3, Canada*
[5]*Department of Physics, Oxford University, Keble Road, Oxford OX1 3RH*
[6]*Department of Astronomy, University of Illinois, Urbana, IL, USA*
[7]*Institute of Cosmology and Gravitation, Mercantile House, University of Portsmouth, Portsmouth PO1 2EG*
[8]*Department of Physics, University of Queensland, Brisbane, Qld 4072, Australia*
[9]*Department of Physics, Durham University, South Road, Durham DH1 3LE*
[10]*Department of Astronomy and Astrophysics, The Pennsylvania State University, 525 Davey Laboratory, University Park, PA 16802, USA*

ABSTRACT

In this work, we test the assertion that the scatter in the mass of black holes which drive quasars should be luminosity dependent with less scatter in more luminous objects. To this end, we measure the width of the Mg II $\lambda2799$ line in quasar spectra from the Sloan Digital Sky Survey (SDSS), 2df QSO Redshift survey (2QZ) and 2dF SDSS LRG And QSO (2SLAQ) surveys and, by invoking an unnormalized virial mass estimator, relate the scatter in linewidth to the scatter of mass in the underlying black hole population. We find conclusive evidence for a trend such that there is less scatter in linewidth, and hence black hole mass, in more luminous objects.

However, the most luminous objects in our sample show such a low degree of scatter in linewidth that, when combined with measures for the intrinsic scatter in the radius–luminosity relation for the broad-line region (BLR) in active galaxies, an inconsistency arises in the virial technique for estimating black hole masses. This analysis implies that, at least for the most luminous quasars, either there is little-to-no intrinsic scatter in the radius–luminosity relation or the Mg II broad emission-line region is not totally dominated by virial velocities.

Finally, we exploit the measured scatter in linewidths to constrain models for the velocity field of the BLR. We show that the lack of scatter in broad-line widths for luminous quasars is inconsistent with a pure planar/disc-like geometry for the BLR. In the case of a BLR with purely polar flows, the opening angle to luminous quasars must be less than ∼55°. We then explore the effects of adding a random or spherically symmetric component to the velocities of gas clouds in the BLR. Assuming an opening angle to quasars of 45°, a planar field can be made consistent with our results if ∼ 40–50 per cent of the velocities are randomly distributed.

Key words: Galaxies: evolution – quasars: emission lines – quasars: general – cosmology: observations.

1 INTRODUCTION

It is now apparent that the majority of massive galaxies harbour a super-massive black hole (SMBH) at their centre. Dynamical

⋆E-mail: sfine@physics.usyd.edu.au

studies of the sphere of influence of these SMBHs have been successful in determining the mass for some tens of systems. In these cases, the mass of the central SMBH has been observed to correlate strongly with properties of their host spheroid such as luminosity (Magorrian et al. 1998), velocity dispersion (σ; Ferrarese & Merritt 2000; Gebhardt et al. 2000) and concentration/Sersic index (Graham et al. 2001). More recently, Hopkins et al. (2007a) have

demonstrated that these relationships can all be regarded as various projections of a 'black hole fundamental plane' (BHFP), relating black hole mass to the potential well of the galaxy.

The existence of these correlations and, in particular, the similarity of the BHFP to the fundamental plane of elliptical galaxies point to an intimate link between the growth of SMBHs and galaxy evolution. Moreover, the BHFP is consistent with the notion that SMBHs grew in a self-regulated manner, from gas gravitationally confined in the galaxy centre, which was eventually expelled by feedback processes (Hopkins, Richards & Hernquist 2007a). Recent hydrodynamical simulations that incorporate radiative cooling, star formation, black hole growth and feedback from both supernovae and nuclear activity have shown that mergers between gas-rich galaxies of comparable mass provide a mechanism for concentrating gas in galaxy centres through tidal effects (e.g. Barnes & Hernquist 1991, 1996), fuelling the growth of SMBHs (e.g. Di Matteo, Springel & Hernquist 2005), reproducing the correlations (e.g. Robertson et al. 2006) and explaining observed properties of quasars as a phase of evolution prior to the termination of black hole accretion (e.g. Hopkins et al. 2005a,b,c).

Further evidence for a merger-driven origin of quasars is provided by comparing the evolution of the abundance of luminous quasars with that of the cosmic star formation rate (SFR). From an empirical determination of the bolometric quasar luminosity function, Hopkins et al. (2007a) infer that the quasar luminosity density peaks at $z = 2.15 \pm 0.05$ and rapidly declines towards higher redshifts. Some observational estimates suggest a similar behaviour for the evolution of the cosmic SFR (e.g. Fan et al. 2001; Hopkins & Beacom 2006). However, at high redshifts, incompleteness, cosmic variance of the surveys and uncertain corrections owing to dust extinction complicate this analysis. From a measurement of the opacity of the Lyman α forest (Faucher-Giguere et al. 2008a), Faucher-Giguere et al. (2008b) have shown that the optical depth at $z = 4$ is incompatible with a steep decline in the cosmic SFR at $z > 3$, as suggested by, for example, the results of Hopkins & Beacom (2006), but in accordance with theoretical modelling (e.g. Hernquist & Springel 2003; Springel & Hernquist 2003) which predicts that the SFR should peak at $z > 4$. The implied offset between the peak in the cosmic SFR and the quasar luminosity density indicates that the growth of SMBHs is not directly tied to star formation or gas density, but is related to a secondary process. By employing estimates of the rate of halo mergers, Hopkins et al. (2008) have shown that the quasar luminosity function can be reproduced if SMBH growth occurs primarily in major mergers of gas-rich galaxies.

Motivated by these various lines of argument, Hopkins et al. (2006a,b, 2007b,c) have developed a model for galaxy evolution in which mergers, starbursts, quasars, SMBH growth and the formation of ellipticals are connected through an evolutionary sequence. As part of this work, Hopkins et al. (2005a) used simulations of merging galaxies to quantify the phases of evolution associated with quasar activity and showed, in particular, that quasar lifetimes depend not only on the instantaneous luminosity of a quasar, but also on its 'peak' luminosity. Convolving these lifetimes with estimates of the merger rates of galaxies, Hopkins et al. (2005b,c, 2006a, 2007b) were able to reproduce the observed optical and X-ray luminosity functions of quasars.

In this interpretation of the quasar luminosity function, the brighter ($>L^*$) objects are all massive black holes, accreting near the Eddington rate towards the end of their growth phase. Less luminous quasars can be either low-mass systems accreting rapidly and undergoing significant growth, or larger black holes accreting at comparatively lower rates.

This model implies that the range of accretion rates in quasars should be luminosity dependent. That is, the range accretion rates (and by extension SMBH masses) at a given luminosity should be larger for lower luminosity objects, and decrease as the luminosity increases.

In this work, we aim to derive the dispersion of the SMBH mass of quasars as a function of luminosity. We measure the Mg II broad emission-line width in spectra from three large spectroscopic surveys of type I active galactic nuclei (AGN). Assuming there exists a virial relation for calculating SMBH masses from the Mg II linewidth, we relate the measured distribution of emission linewidths to the SMBH mass distribution.

2 VIRIAL SMBH MASS ESTIMATORS

By far, the most accurate, robust and believable method for measuring the mass of SMBHs is to perform dynamical studies of stars, gas or masers in the potential of the SMBH (e.g. Herrnstein et al. 1999; Gebhardt et al. 2003; de Francesco, Capetti & Marconi 2006). Rarely, however, is this viable for type I AGN which for the most part are too distant to resolve the sphere of influence of the SMBH, and have nuclei so bright that it is difficult if not impossible to obtain detailed photometric or spectroscopic information of this environment.

The most direct method for calculating SMBH masses of type I AGN is via reverberation mapping of the broad-line region (BLR) and the virial theorem (Peterson et al. 2004). In this case, the size of the BLR (r_{BLR}) is given by the time delay between continuum and emission line variations, and the velocity dispersion of the BLR (V_{BLR}) is measured as the width of the emission line in the variable spectrum. The mass of the SMBH is then estimated as

$$M_{\mathrm{BH}} = f \frac{r_{\mathrm{BLR}} V_{\mathrm{BLR}}^2}{G}. \tag{1}$$

The factor f defines what is not known about the BLR: its geometry, velocity field and orientation. The value of f is of the order of unity and authors take various approaches assigning it a value. One can find $f = 3/4$ applicable to random orbits (Peterson & Wandel 1999), $f = 1/(4 \sin^2\theta)$ for a disc inclined at an angle θ to the observer (McLure & Dunlop 2001) or even $f = 1$ for simplicity (McLure & Jarvis 2002). Onken et al. (2004) measured $\overline{f} = 1.4$ by comparing the M–σ relation in a group of reverberation-mapped AGN in local spheroids.

This last value is somewhat higher than expected compared to the simple theoretical values. However, the interpretation of this is unclear and the discrepancy may be due to the selection bias (Lauer et al. 2007) and/or cosmological evolution in the M–σ relation (Woo et al. 2006), both of which could bias the measured value of f high.

Due primarily to time constraints, the number of reverberation-mapped systems is in the tens and the luminosity range which these measurements span is not huge (Kaspi 2007). However, one key result which has come out of reverberation mapping is the radius–luminosity relation (Kaspi et al. 2005). While there is significant scatter in this relation, including significant intrinsic scatter of up to 40 per cent, it does offer a simple single epoch method for estimating the radius of the BLR in AGN.

If one takes the instantaneous full width at half-maximum (FWHM) of a broad emission line as a measure of the virial velocity in the BLR, then the r–L relation provides a method for estimating the unscaled SMBH mass with single epoch observations. This

Table 1. Summary of the surveys from which we obtained spectra. Successive surveys have fewer spectra but go deeper, increasing our luminosity range at a given redshift. Note that the magnitude limits quoted for the SDSS QSO survey are those for the primary QSO survey. The high-redshift sample goes deeper, and included are sources observed under different selection criteria and also QSOs identified as part of other surveys.

Survey	No. of objects	Mag. limits	Resolution (km s^{-1})	Dispersion (Å pixel^{-1})	$\overline{S/N}$ (pixel^{-1})
SDSS (DR5)	77,429	$19 > i > 15$	$\sim$165	$\sim$1.5	$\sim$13
2QZ	23,338	$20.85 > b_J > 18.25$	$\sim$465	$\sim$4.3	$\sim$5.5
2SLAQ	8,492	$21.85 > g > 18.00$	$\sim$465	$\sim$4.3	$\sim$5.5

technique has become known as the virial method. To date, virial relations have been calibrated for the Hβ (Kaspi et al. 2000; Vestergaard & Peterson 2006), Mg II (McLure & Jarvis 2002; McLure & Dunlop 2004; Kollmeier et al. 2006; Salviander et al. 2007; McGill et al. 2008) and C IV lines (Vestergaard 2002; Vestergaard & Peterson 2006). In each case, these calibrations result in expressions of the form

$$M_{BH} = A(\lambda L_\lambda)^\alpha \text{FWHM}^2, \qquad (2)$$

where FWHM is the full width at half-maximum of the spectral line in question and λL_λ is the monochromatic luminosity of the continuum near that line. A is a normalization constant and the exponent α gives the luminosity dependence of the r–L relation.

2.1 Problems with virial mass estimators

When using virial SMBH masses, it is prudent to state some caveats. First and most obviously, the virial relations are statistical in nature. Hence, while they may be accurate when averaged over a large number of systems, individual measurements should be viewed with caution.

Secondly, a number of studies have investigated the differing emission regions for high- and low-ionization lines (e.g. Richards et al. 2002; Baskin & Laor 2004; Elvis 2004) which raises questions as to whether all broad lines can be used as virial mass indicators. This is, however, still a subject open to debate. Vestergaard & Peterson (2006) reanalysed Baskin & Laor's data and showed that the apparent discrepancies between Hβ and C IV virial mass estimates are nullified by applying more appropriate selection criteria.

Thirdly, it is important to note that the calibrations of the Mg II and C IV virial relations are not direct. These lines have not been sufficiently studied in reverberation mapping campaigns to calibrate any potential r–L relation for them. Instead, the virial relations for these lines are normalized through comparisons with SMBH masses calculated from the Hβ line.

Finally, there is simply a dearth of solid information constraining the velocity field and geometry of the BLR. Reverberation mapping as a technique has the potential to describe the BLR in detail given sufficient quality data (Welsh & Horne 1991). However, in reality this sort of idealized precision is unlikely, and to date these types of results have not materialized. Theoretical models for BLRs exist (e.g. Emmering, Blandford & Shlosman 1992; Königl & Kartje 1994; Murray & Chiang 1997) but the lack of strong observational constraints makes a proper comparison between these difficult. However, almost all of these models imply a strong virial component to the velocity field of the BLR.

There is good evidence that the virial method, while imprecise, can on average give accurate black hole masses. Comparisons between reverberation masses and bulge velocity dispersion show an M–σ relation analogous to that observed in nearby quiescent galaxies (Onken et al. 2004). Comparisons between virial and re-

verberation masses also show an agreement (McLure & Jarvis 2002; Vestergaard & Peterson 2006). But, to obtain strong results from virial masses large samples are required to beat down random errors in these estimates.

In this work, we assume a virial relation for the Mg II line as given by equation (2). We do not, however, assume a parametrization for this relation since it is not required. We take a large number of quasi-stellar object (QSO) spectra from several spectroscopic surveys and bin them by luminosity and redshift. We invoke the r–L relation and assume objects in a luminosity bin will have the same r_{BLR}, and the scatter in SMBH masses within that bin is simply given by the scatter in broad-line widths. Furthermore, looking at the scatter in log space for sufficiently small luminosity bins, we can ignore the L term in equation (2), and the coefficient and

$$\text{Disp}(\log(M_{BH})) = 2 \times \text{Disp}(\log(\text{FWHM}_{Mg II})), \qquad (3)$$

where Disp() denotes the dispersion in the given variable. Note that we have ignored all extraneous sources of scatter in this equation, some of which will be significant to our calculations. We discuss these in more detail in Sections 6.3 and 6.4.

3 DATA

We take as our sample all of the quasar spectra from the Sloan Digital Sky Survey (SDSS; York et al. 2000) data release five (DR5; Adelman-McCarthy et al. 2007) as compiled by Schneider et al. (2007). To increase our luminosity range, we also take all QSO spectra from the 2df QSO Redshift survey (2QZ; Croom et al. 2004) and the 2dF SDSS LRG And QSO survey (2SLAQ; Richards et al. 2005; Croom et al. in preparation) as well. Table 1 shows a brief summary of the number of objects and magnitude limits in each sample.

We thus have spectra from three different surveys taken with two different instruments. We give here a brief description of the spectra.

3.1 SDSS spectra

Details of the Sloan telescope and spectrograph are given in Gunn et al. (2006) and Stoughton et al. (2002). The spectra have a logarithmic wavelength scale translating to a dispersion of $\sim$1–2 Å pixel^{-1} and a resolution $\lambda/\Delta\lambda \sim 1800$ in the wavelength range 3800–9200 Å. Objects are observed initially for 2700 s and then reobserved in 900 s blocks until the median signal-to-noise ratio (S/N) is greater than $\sim$4 pixel^{-1} resulting in a S/N distribution with a mean at $\sim$13 pixel^{-1}.

The spectra are extracted and reduced with the SPECTRO2D pipeline and automatically classified with SPECTRO1D (Stoughton et al. 2002). However, in creating the Sloan QSO sample used in this paper,

Figure 1. Example of a 2dF spectrum with poor response at the blue end. In the top panel, the original spectrum is shown as reduced by 2dFDR and the dashed line shows the response correction from Lewis et al. (2002) scaled for comparison. In the bottom panel, this correction is applied to the spectrum. We find that, while this does improve the shape of the spectrum, it is inadequate in this case to correct for the sharp drop in received flux at the blue end.

Schneider et al. (2007) visually inspect all of the candidate spectra to determine their classification.

3.2 2dF spectra

Both 2QZ and 2SLAQ spectra were taken with the 2dF instrument on the Anglo-Australian Telescope with the 300B grating (Lewis et al. 2002). Spectra have a dispersion of 4.3 Å $pixel^{-1}$ and a resolution of $\sim$9 Å in the wavelength range 3700–7900 Å. 2QZ observations were between 3300 and 3600 s compared with 14 400 s for 2SLAQ. The increase in exposure time for the fainter 2SLAQ sample results in S/N distributions that are almost indistinguishable. Both peak at $\sim$5.5 $pixel^{-1}$ for positive QSO IDs. 2dF spectra are extracted and manually classified during the observing run with the 2dFDR pipeline (Bailey & Glazebrook 1999) and AUTOZ redshifting code (Croom et al. 2001).

The main difference between reduced Sloan and 2dF spectra is the lack of flux calibration for 2dF sources. An average flux calibration for the 300B grating has been calculated by Lewis et al. (2002) as part of the 2dF Galaxy Redshift Survey, and in our analysis we do apply this correction. However, a quick examination of 2dF spectra shows the inadequacy of this median calibration to correct for the variations in response between differing spectra as illustrated in Fig. 1.

4 ANALYSIS

Our goal is to measure the dispersion in the widths of broad Mg II lines in a large sample of QSOs. For this, we require an automated routine to measure the width of Mg II in a consistent way across our sample and a robust method for analysing the results. In this section, we discuss the line-fitting code used to measure the linewidths, and our analysis of the output, but first a quick note on absolute magnitudes.

4.1 Magnitude calculations

Magnitudes are given in the b_J band in the 2QZ catalogue, and Sloan *ugriz* bands in both the SDSS and 2SLAQ catalogues. Richards

et al. (2005) compared the b_J and g pass bands and found them to be roughly equivalent. Examining 2QZ QSOs with g-band imaging, they found a consistent $\langle g - b_J \rangle = -0.045$. In our analysis, we use g-band magnitudes for all objects taken from the SDSS and 2SLAQ data sets, and for the 2QZ we apply this correction to their b_J magnitudes.

Throughout this paper, we use a flat $(\Omega_m, \Omega_\Lambda) = (0.3, 0.7)$, $H_0 = 70 \, km \, s^{-1} \, Mpc^{-1}$ cosmology when calculating absolute magnitudes. We use the *K*-corrections laid out by Cristiani & Vio (1990) and correct for galactic extinction as advised in the relevant catalogue paper.

4.2 Line measurements

Our spectral analysis routine implements two separate fitting processes which are linked via iteration. First, a combined iron and continuum model is fit to the region of the spectrum not affected by the Mg II line, then a Gaussian profile is fit to the line itself. We iterate the procedure to more accurately define the region of the spectrum affected by Mg II emission, and so improve our iron and continuum fits. This significantly improves our spectral line calculations.

4.2.1 Continuum and iron correction

Continuum emission from QSOs is well described by a power law in the optical-UV region. However, 2dF spectra are not flux calibrated. Variations in response modulate the shape of 2dF spectra most often characterized by a drop off in flux at the blue end. Since we cannot accurately model these spectra with a power law, we describe the continuum with a quadratic fit. This can then simultaneously correct for response effects in 2dF spectra and, since we only fit this to a limited region of the spectrum, approximate a power-law shape for Sloan objects. Fig. 2 shows our iron and continuum fits to both a 2dF spectrum which has a drop off in flux at the blue end, and a well-calibrated SDSS spectrum.

To remove iron emission from our spectra, we use the template for QSO iron emission derived from the narrow-line Seyfert I object I Zw 1 by Vestergaard & Wilkes (2001). This template is derived from an intrinsically narrow-line object which made it possible to isolate the species responsible for each component of the emission. Unfortunately, because the template is derived from a real object, it is least well defined in the region around strong emission lines where de-blending differing species becomes difficult.

This limitation is of particular importance directly beneath the Mg II line where the template shows no iron emission. This is not due to a quantified lack of emission, but to the difficulty of de-blending the iron and Mg II emission in this region. It has now become common practise to add flux to the template in this region (e.g. Kurk et al. 2007; Salviander et al. 2007) to make it more consistent with theoretical models (Sigut & Pradhan 2003). We follow the method of Kurk et al. (2007) and add a constant level of flux under the Mg II line at 20 per cent the intensity of the average in the 2930–2970 Å (rest-frame) region. We find that the level of flux included under the Mg II line does affect our results. Given the range of likely levels of iron emission in this region, however, the effect is not significant and does not impact our conclusions. This will be discussed further in Section 5.2.

Since the iron template is derived from a narrow-line object, it must be smoothed to properly describe the iron emission in a broader-line QSO. We follow the same procedure as followed by

Figure 2. These figures illustrate the iron and continuum fitting process on two very different spectra. In each case, the top panels show the Mg II line region in the original spectrum, heavy dashed and solid lines show the continuum and iron + continuum model fit to it, respectively. The bottom panel then shows the residual when this is subtracted from the spectrum. (a) The fit to the spectrum of J095421.6−000152 as observed as part of the 2QZ. This is the same spectrum as given in Fig. 1 and shows a strong turnover in received flux at the blue end. (b) The spectrum of J000042.89+005539.5 as observed with the Sloan telescope with a classic quasar power-law continuum. For both cases, the quadratic continuum we use produces an excellent fit to the data. In all plots, the *y*-axis gives the normalized flux.

previous authors (e.g. Boroson & Green 1992) and make a selection of iron templates smoothed by Gaussians of the width 500, 750, 1000, 1500, 2500, 5000, 7500 and 10 000 km s^{-1}. We fit all of these templates to this spectrum and accept the best fit in terms of the χ^2.

We fit for continuum and iron emission simultaneously with the SVDFIT routine (Press et al. 1989). The fit is performed on a region bounded at the blue end at 2450 Å by the [Ne IV] λ2424 line, and at the red end at 3085 Å which is the limit of the iron template (all wavelengths are rest frame). We also mask the Mg II line out of our fitting.

On the first attempt, we use a fiducial $\pm$50 Å mask for the line but on subsequent iterations we use a Gaussian fit to the line to define its boundaries.

It is important to keep in mind that once we have subtracted the iron and continuum contributions from the data, we introduce a strong covariance into our spectrum and individual pixel values can no longer be thought of as independent in subsequent error calculations.

4.2.2 Gaussian fitting, iteration and linewidth calculations

We subtract the iron and continuum contribution to the spectrum and fit a single Gaussian to the Mg II line with the MRQMIN routine (Press et al. 1989). We then perform the iron and continuum fit to the original spectrum again masking out data within $\pm$1.5 FWHM of the centre of the fitted Gaussian. This process is repeated until successive Gaussian fits have FWHMs consistent to within 1/2 their error.

In this analysis, some thought must be given to what statistic to use when defining the width of a spectral line. The width of single and/or multiple Gaussian/Lorentzian fits to the line is often used to describe the linewidth (e.g. Wang, Lu & Zhou 1998; McLure & Dunlop 2004; Shen et al. 2008). We have already performed Gaussian fits to our line and we do find that despite the clear non-Gaussian nature of the Mg II line, these give a reliable measure for the width when compared with the other methods outlined here. However, we find non-parametric estimators for the linewidth more appealing for this analysis, primarily because they

make no assumptions as to the line profile, but also because their errors can readily be calculated including the contribution of covariance introduced into the spectrum in our iron and continuum correction.

A common statistic used for describing linewidth is the FWHM (Vestergaard 2002; Vestergaard & Peterson 2006). In high S/N spectra, this gives a good determination of the linewidth, and the error for the FWHM can readily be calculated including the contribution from covariance in the spectrum. Unfortunately, the measured FWHM is quite susceptible to noise. How to define the maximum flux density of a line in an unbiased way is unclear and multiple crossings of the half-maximum value in noisy spectra demand some sort of averaging which can also bias the measurement. Further, we find that the susceptibility of the FWHM to the values of a relatively small number of pixels in a spectrum makes it an unstable and often inaccurate measure of the line's true width.

Another measure of the linewidth that is becoming increasingly widespread is the dispersion of the line σ (Fromerth & Melia 2000; Vestergaard & Peterson 2006; Wilhite et al. 2007). However, we find that the excessive weighting this measure gives to the values of pixels in the wings of the lines makes it an unreliable estimator of linewidth in low S/N spectra.

Finally, inter-percentile values (IPV) can be measured (Whittle 1985). While at first glance the process of measuring an IPV width is similar to measuring the FWHM, the dependence of IPV widths on the cumulative flux distribution rather than on the flux density at a given point makes the IPV measurements considerably more robust. Like the dispersion, IPV widths are somewhat affected by noise in the wings of lines; in particular, this can affect the total flux of a line and how one defines the zero-point of the cumulative flux distribution. However, when calculating the dispersion the weight given to a single pixel is proportional to the square of the displacement of that pixel from the line centre. This power-of-two dependence makes σ very susceptible to noise in the wings of a line which is not a problem for IPV widths.

Our calculation for the IPV width is performed only on the part of the spectrum within $\pm$1.5 FWHM of the Gaussian fit to the Mg II line. In this region, we find the cumulative flux distribution

and search this for the first crossing points at 1/4 and 3/4 the total flux. We then interpolate between the adjacent pixels on either side of the crossing to obtain sub-pixel accuracy in our IPV estimate.

Errors on the IPV widths are calculated from the spectral variance array including the contribution of covariance introduced by the iron and continuum subtraction. We find the mean error of these measurements to be ± 0.05 dex.

Finally, we subtract the resolution of the spectrograph in quadrature from our measured linewidth under the assumption of a Gaussian profile for both the emission line and instrumental resolution. Since the Mg II line is not Gaussian in shape, this is only an approximate correction, and to test its validity we evaluate the effect of smoothing Lorentzian profiles by a Gaussian. In the worst case scenario of a Lorentzian line with the smallest IPV width, we measure in our sample (~ 1200 km s^{-1}; see Fig. 6), smoothed by a Gaussian with FWHM 675 km s^{-1} (equivalent to the 9 Å resolution of the 2dF spectrograph at $\lambda = 4000$ Å), the correction above results in an overestimate of the linewidth by <5 per cent. For broader spectral lines, and in cases where the spectral resolution is greater, this effect is diminished and hence we make no further attempt to correct for instrumental smoothing.

4.2.3 IPV widths versus FWHM

In taking the IPV width as our measure of linewidth, we assume there is a linear scaling relation between the differing linewidth parametrizations. Specifically, we assume that equation (3) holds equally true for IPV widths. To test this assumption, we also measure the FWHM of the Mg II lines in our sample to compare with the IPV widths.

These are measured by defining the maximum flux density of the spectral line as the average of the highest pixel value and the two pixels adjacent to it. Then, the lower and upper bounds of the FWHM are found by averaging all crossings of the half-maximum value on the red and blue sides of the lines centre. Fig. 3 shows a comparison between the IPV width and FWHM measurements for our data set. In the figure, we have multiplied the IPV widths by 1.75, the ratio between IPV and FWHM for a Gaussian. There


Figure 3. Comparison of 50 per cent IPV widths and FWHMs for our sample. The dashed line gives the 1:1 relation, and the solid line shows the *y* on *x* least-squares fit. IPV widths have been multiplied by 1.75 for the comparison which is the expected ratio for a Gaussian line.

is considerably more scatter in the FWHM direction due to the susceptibility of this statistic to noise in the spectrum. However, a *y* on *x* least-squares fit to these data gives

$$\log(\text{FWHM}) = 0.98 \log(IPV) + 0.1, \qquad (4)$$

a linear relation between the IPV width and FWHM. The offset is due to the profile of the Mg II line which tends to have more flux in the wings of the line compared to a pure Gaussian.

4.2.4 Rejection criteria

With an automated routine, it is impossible to accurately fit every spectrum and there is always the chance of catastrophic failure. We make a series of checks and cuts to the results to try and reduce the number of spurious fits in our final sample.

Failure to produce a satisfactory fit is normally due to a property of the spectrum being analysed. Absorption features, low S/N, extreme curvature in the underlying spectrum, bad pixels and residual telluric features can all cause failures, and commonly a failure is due to a combination of these.

To test the effect of S/N on the accuracy of our fitting routine, we simulate large numbers of noisy spectra by adding Gaussian noise to high-quality Sloan and 2dF spectra. We then measured the width of the Mg II line in these degraded spectra and compared our results with the measurement from the high S/N spectrum.

Down to a S/N ~ 1.5 pixel^{-1}, we find no systematic deviation between the average linewidth measured from noisy spectra and the true linewidth. Furthermore, we find that above this S/N the errors on our measurements correctly describe the scatter we observe in the measured linewidths. We therefore apply a spectral S/N cut of S/N > 1.5 pixel^{-1} to our sample.

It is worth noting that while our error calculations assume Gaussianity in linear space, we find that the distribution of IPV widths measured for a given spectrum at low S/N is better described by a lognormal. This may not be surprising since the widths are randomly distributed and limited to be >0, however, that our errors transfer from linear to log space is encouraging. We transform the linear error σ_{lin} to the log error σ_{log} via its ratio to the measured linewidth IPV. If we propagate the error through the log, then

$$\sigma_{\text{log}}^2 = \left[\frac{\partial}{\partial IPV} (\log_{10}(IPV)) \right]^2 \sigma_{\text{lin}}^2 \qquad (5)$$

$$\sigma_{\text{log}} = \frac{\sigma_{\text{lin}}}{IPV \ln(10)}. \qquad (6)$$

We make three further constraints on our data as to whether we will accept a particular fit. We apply a redshift cut to avoid contamination by the many telluric features towards the red end of our spectra. We limit the number of times we iterate our fitting procedure as described in Section 4.2.2, and finally we try to eliminate broad absorption line (BAL) objects which could contaminate our data.

Telluric features pervade unreduced spectra redwards of 7100 Å. While both SPECTRO2D and 2dFDR try to remove sky features, the reduced spectra often contain residuals in this region. To avoid the worst of this, we make a simple redshift cut at $z = 1.5$ which ensures that the Mg II line and much of our continuum fit avoid this region.

For a typical spectrum, we iterate the fitting procedure outlined above two to three times to obtain convergence. We limit the maximum number of iterations allowed to 20 and reject any fits which have not converged by this time as unreliable.

Table 2. Table outlining our final data set. For each of the samples, we give the total number of objects for which we have attempted to analyse the Mg II line as well as the apparent (uncorrected for Galactic extinction) and absolute (extinction corrected) magnitude range of these objects. We also give the number of these fits which have been rejected from our analysis by each of our criteria. These are low S/N spectra, possible BAL objects and objects for which the fitting code did not converge on a result within 20 iterations. Note that the redshift cut at $z = 1.5$ to avoid sky contamination has already been factored into the number of spectra in Column 2. We also give the total number of fits rejected, since an object can be rejected for a number of reasons this does not equal the sum of the previous three columns.

Sample	No. fit	App. mag. range	Abs. mag. range	S/N	BAL	Itt.>20	All
					No. of fits rejected		
2SLAQ	2684	$25.09 > g > 18.12$	$-19.48 > M_g > -26.38$	44	53	142	193
2QZ	7209	$20.80 > g > 18.20$	$-21.07 > M_g > -26.74$	161	75	356	425
SDSS	23725	$23.33 > g > 15.32$	$-20.98 > M_g > -29.27$	83	141	653	786
All	33618	$25.09 > g > 15.32$	$-19.48 > M_g > -29.27$	288	269	1151	1404

BAL objects are a contaminant in our data and we try to identify and reject these during our fitting. In the 2QZ and 2SLAQ catalogues, many of the most severe BAL objects are flagged as such and are not included in our analysis. Those which are not flagged in a catalogue are identified on the basis that they will have pixels which are significantly deviant from the rest of the spectrum. After every Gaussian fit, we reject any pixels which deviate by $>3\sigma$ from the fit and we do not include these pixels in subsequent iterations. Once convergence is obtained in our fitting process, we reject any object which has consecutive rejected pixels spanning $500\,\mathrm{km\,s^{-1}}$.

Note that this process of rejecting pixels does not only affect BAL objects and any pixels with outlying values will be excised. If this process removes more than one-fifth of the pixels within ± 1.5 FWHM of the Mg II line, we also reject the fit from our data.

Details of how many objects are rejected due to these various cuts are given in Table 2. In total, ~ 4 per cent of the objects are rejected from our sample. More than 80 per cent of these are objects rejected due to the fitting routine not converging on a solution which occurs primarily on low S/N spectra. This, combined with the imposed S/N cut, could create a bias at low luminosity in our data, in particular if there is a strong correlation between IPV width and luminosity. However, as we will see there is only slight evidence for a correlation between luminosity and IPV width. In addition to this, the small number of objects rejected mean that any bias will be negligible. On the other hand, the BAL rejection is designed to expel true outliers from our data and should not bias the results in any way.

Fig. 5 shows the results from our fitting plotted on a linewidth–magnitude diagram, and Fig. 4 shows how these objects are distributed on the redshift–luminosity plane. Each point represents an object with an accepted linewidth measurement. The fitting results for these objects are available from the 2SLAQ website (www.2slaq.info).

Added to the plot are lines of constant SMBH mass (dotted) and Eddington ratio (dashed). These are calculated using the McLure & Dunlop (2004) calibration for the Mg II virial relation, assuming their B-band bolometric correction and taking $g \sim B - 0.11$ to calculate the bolometric luminosity of the sources. We then use a bolometric correction to the continuum luminosity at $3000\,\mathrm{\AA}$ of 5.2 (Richards et al. 2006) to calculate the continuum luminosity at this wavelength from the bolometric luminosity. The McLure & Dunlop calibration was based on linewidths measured in a different fashion to this work. They fit two Gaussians to the Mg II line and take the FWHM of the broader component as their linewidth. We correct the lines in Fig. 5 for our use of IPV widths assuming a Gaussian profile for the Mg II lines, but McLure & Dunlop's use of two Gaussians in their fitting will likely bias the plotted lines high by a factor of 1.5 to 2. These lines are plotted more as a guide as to how lines of


Figure 4. Distribution of our final sample of objects in the redshift–luminosity plane. Magnitude limits for the various surveys can be made out as lines where the density of objects increases. The very few objects with g-band magnitudes below the 21.85 cut for 2SLAQ were selected as potential high-redshift objects in the i band.

constant mass and Eddington ratio would lie on the diagram rather than being exact in normalization.

We will discuss the implications of this diagram further in Section 6, but already if one compares contours on the left- and right-hand side of this plot it appears to suggest that there is less scatter in broad-line widths for more luminous quasars. Furthermore, the average linewidth shows very little variation across the luminosity range sampled, a result consistent with previous analyses (Corbett et al. 2003; Shen et al. 2008).

4.3 Dispersion analysis

To investigate how the dispersion in M_{BH} depends on the luminosity of QSOs in our sample, we bin our sample by luminosity and calculate the dispersion in linewidths in each bin. To remove possible redshift evolution in this relation, we also bin by redshift. As with measuring linewidths, some consideration must be given to choosing the most robust estimate for the dispersion in each bin.

The overall distribution of QSO linewidths is shown in Fig. 6. We find it is roughly lognormal although with extended wings and a slight asymmetry. This further motivates our use of logarithms in equation (3) to derive the dispersion in M_{BH}.

The obvious way to measure dispersion is with the rms in log space. This approach is, however, very susceptible to outliers. Objects with either very large or very small results for their linewidth

Figure 5. This plot gives the results of our line-fitting analysis. We plot the absolute *g*-band magnitude of the source versus the measured IPV velocity width of the Mg II line. Overplotted are lines of constant SMBH mass (dotted) and Eddington ratio (L/L_{Edd}; dashed) as a guide to where these objects fall in mass-accretion space. Masses are labelled in units of $M_\odot$. We have added contours of even density to the plot to highlight the shape of the distribution. The contours are evenly spaced in terms of log(density).

Figure 6. Distribution of Mg II linewidths for our sample (heavy histogram) with a Gaussian fit for comparison (fine line). We find this to be roughly lognormal, although with slightly more objects in the wings in particular towards larger linewidths.

are more likely to be due to poor spectral fitting and tend to have large errors on these measurements. We find that the IPV method is less biased by these outliers, although for the most part we find very little difference between the two statistics.

We must be aware that the distribution of measured linewidths is not a direct representation of the intrinsic distribution of linewidths in our sample. Instead, the measured distribution represents the intrinsic distribution convolved with the distribution of errors on these estimates. In the perfect case, where both are normally distributed in log space the intrinsic rms, rms_i, is related to the measured rms, rms_m, and the average error in linewidth $\overline{\sigma^2}$ by

$$rms_i = \sqrt{rms_m^2 - \overline{\sigma^2}}. \qquad (7)$$

We are not using the rms in this analysis, instead we take the 68.3 per cent IPV width to provide an equivalent statistic. Further, we do not take $\overline{\sigma^2}$ to describe the average error in our data since we find this can be skewed by a small number of objects with very large errors on their linewidths. To avoid these outliers, we take the square of the median error in a bin to estimate $\overline{\sigma^2}$.

5 RESULTS

We calculate the dispersion in linewidths for each $L - z$ bin and correct by the median error on the linewidths in that bin. To ensure reliable dispersions, we only consider $L - z$ bins that contain more than 40 objects. Fig. 7 displays this corrected dispersion versus luminosity for each of the 2QZ (*a*), 2SLAQ (*b*) and SDSS (*c*) samples separately, and then for the combined data set (*d*). Note that while we bin our data by redshift, we do not see evidence for redshift dependence in our sample (see Section 5.1) and so we plot the data from all redshift bins in Fig. 7.

Error bars in the plots are propagated from the errors on the linewidth measurements. For the most part, we have a large number of objects in each $L - z$ bin, and thus are dominated not by these random errors but by systematics in the analysis.

Figs 7(a) and (b) have considerable scatter and exhibit little-to-no trend in the plots, although the shallow trend seen in the 2QZ data is significant [a Spearman rank test gives $r_s = 0.56$ with $P(r_s) = 0.01$]. Fig. 7(c) shows a strong relation with little scatter due to the larger numbers in the SDSS sample (Table 2) and its higher quality spectra.

Fig. 7(d) shows the results when we combine all of our data. This is heavily dominated by SDSS objects down to absolute magnitudes $M_g \sim -23$ where the 2QZ and then 2SLAQ samples become important.

Figure 7. Plots of the dispersion in log (*IPV*) in absolute magnitude bins for the 2SLAQ (a), 2QZ (b) and SDSS (c) samples, as well as the combined sample (d). Since we find no redshift dependence in our data, we plot the results from all redshift bins in one figure, hence more than one point per magnitude interval. The open symbols show the dispersion in measured linewidth before correction (rms_m; see equation 7) for errors on our measurement and filled symbols with error bars give the corrected dispersion (rms_i).

Shen et al. (2008) perform a similar analysis on the SDSS DR5 quasar sample. They find a consistent level of variance in their data but report no dependence on luminosity, although the inspection of their fig. 4 does suggest a trend albeit slight. It may be that their following the McLure & Dunlop (2004) prescription for measuring the width of the Mg II line and/or the use of large luminosity bins in this figure could explain their not finding as strong a relation as apparent in Fig. 7.

5.1 Redshift dependence

We find in Fig. 7 that points from the same luminosity bin but differing redshift bins lie on top of each other in the plot suggesting the dispersion in linewidths has little-to-no redshift dependence. This is further illustrated in Fig. 8. In this plot, we show the dispersion in IPV linewidth plotted against redshift for $L - z$ bins. Note that these are not the same $L - z$ bins as used in the rest of this paper, we have doubled the number of redshift bins and halved the number of luminosity bins to make the plot clearer.

In the plot, we connect points which are in the same luminosity bin and it is evident the dispersion in IPV linewidths in these bins has at most a weak redshift dependence. A Spearman rank test performed only on the three luminosity bins which have full redshift coverage gives $r_s = -0.14$ significant at only the 50 per cent level. This is consistent with the changing distribution of luminosities within each L bin as the luminosity function evolves with redshift.

5.2 The effects of the iron template

As discussed above, the Vestergaard & Wilkes (2001) iron template used in our analysis includes no iron emission directly under the


Figure 8. Here, we plot the dispersion in log (*IPV*) versus redshift for our sample. Luminosity bins are 1 mag in width and are centred on (top to bottom) $M_g = -21, -22, -23, -24, -25, -26$ and -27. Bins with the same luminosity cuts are given the same symbol and connected in the plot, and it is evident that the dispersion in IPV widths in these bins is only very weakly dependent on redshift.

Mg II line and we correct for this by adding flux at a level suggested by theoretical models for QSO iron emission (Sigut & Pradhan 2003). To test what effect this has on our results, we refit all of our data with an iron template with no iron emission under Mg II and twice as much iron emission as used in the main analysis. Finally, we also try an iron template where we interpolate between the iron emission peaks at ~2750 and 2840 Å (see e.g. Fig. 13; Sigut & Pradhan 2003); this template includes more than five times the flux in the Mg II region as in our primary template.

There is a systematic trend between our results and the amount of iron emission added under the Mg II line. The more emission, the more scatter we find in our linewidths. However the effect is small. We find an offset of ∼0.015 between the dispersion in log(*IPV*) for objects fit with no iron under the Mg II line and twice as much iron emission as assumed in this analysis; the data presented in Fig. 7 lie in between these two. In the case of the template where we interpolate between the iron peaks on either side of Mg II, we find an offset of ∼0.08. While this approach produces a significant offset, this model has no physical basis and was only implemented to test the extremes of the distribution.

Furthermore, regardless of which iron template we use in our fitting, the trend in our dispersion analysis is the same. We always find that more luminous objects show less scatter in linewidths than fainter objects.

5.3 Completeness and homogeneity

While the 2SLAQ, 2QZ and primary SDSS QSO surveys all have high spectroscopic completeness overall (∼90 per cent), towards the faint flux limit of each survey this drops off. Furthermore, the amalgam of many separate surveys in the SDSS quasar catalogue and our adding to this the 2SLAQ and 2QZ spectra as well makes for a very inhomogeneous sample of objects in terms of selection criteria. However, these factors are not of great importance in terms of biasing our results or the trend shown in Fig. 7.

Incompleteness at faint magnitudes should not have a major effect since the inspection of Fig. 5 shows there is little variation in the average linewidth with luminosity. Hence, the dispersion in IPV widths in a luminosity bin is not affected strongly by the size of a luminosity bin, or the distribution of objects within a given bin.

Inhomogeneity in selection criteria is also not responsible for the observed trend in Fig. 7. This is evident since we find exactly the same trend in each sample separately. Furthermore, we find the same trend if we only take the primary SDSS quasar sample, albeit with significantly reduced luminosity coverage.

6 DISCUSSION

In Fig. 5, the results of our line fitting are plotted along with contours of constant SMBH mass and Eddington ratio. While there is uncertainty in the normalization of these lines, the distribution of points in the plot relative to them is quite suggestive. The linewidth distribution appears to be constrained on at least two sides parallel to these lines.

First, the bottom of the distribution follows the line at $L/L_{\mathrm{Edd}} = 1$. Potentially, this indicates that the Eddington rate does represent an upper limit for allowable accretion in QSOs, constraining our results to lie above this line.

Secondly, brighter than $M_g \sim -24.5$, the top of the distribution of linewidths follows the $M_{\mathrm{BH}} = 10^{10}\,\mathrm{M}_{\odot}$ contour on the plot. This is likely due to the drop off in the SMBH mass function at high mass. The space density of SMBHs falls off dramatically for masses above $10^9\,\mathrm{M}_{\odot}$, and this creates an upper limit to the distribution of linewidths plotted in Fig. 5.

Finally, and more speculatively, fainter than $M_g \sim -24.5$, the top of the linewidth distribution appears to be constrained along a line of constant Eddington ratio around $L/L_{\mathrm{Edd}} \sim 0.01$. The implication of this drop off of objects towards low accretion rates is that there is a preferred level of accretion on to SMBHs for QSOs. This indicates that the accretion distribution for SMBHs as a whole


Figure 9. This shows how the skew in the IPV linewidth distribution is dependent on the magnitude of QSOs. The skew is defined by equation (8), and the error bars are calculated from the errors on the IPV width measurements propagated through this equation.

is bimodal, with quiescent and active SMBHs occupying distinct areas of mass-accretion space.

As a test as to how the Eddington limit affects the accretion efficiency distribution, we measure the skew of the linewidth distribution in magnitude bins. We take the dimensionless skew defined by the third moment of the IPV width distribution

$$\mathrm{Skew} = \frac{1}{N} \sum_{i=1}^{n} \left(\frac{\mathrm{IPV}_i - \overline{\mathrm{IPV}}}{\sigma_{\mathrm{IPV}}} \right)^3. \tag{8}$$

Fig. 9 plots the skew in the IPV width distributions against the magnitude of the bins.

This figure shows that in more luminous QSOs, the linewidth distribution consistently has a positive skew. This skew towards larger linewidths and hence lower accretion efficiencies may be the result of the accretion efficiency distribution being truncated at the fast end by the Eddington rate.

There is some evidence that this skew reverses for fainter objects. At these luminosities, the mean of the accretion efficiency distribution is significantly lower than for brighter objects, hence a cut at the Eddington rate will not have as strong an effect. The reversal of the trend is likely due to the underlying mass function of active SMBHs. With many more small SMBHs, there would be a skew towards higher accretion efficiencies and hence smaller broad-line widths at a constant luminosity.

Taken at face value, the distribution of points in Fig. 5 fits well with simple expectations about the QSO SMBH population, suggesting that the virial mass estimates work relatively well. However, see Section 6.3.1 for a discussion on some concerns with the virial mass technique which these data highlight. None the less, this gives us confidence that we can make a meaningful comparison between our data and theoretical models for activity in the SMBH population.

6.1 Comparisons with models

In Fig. 10, we compare our results to the model predictions of Hopkins et al. (2005b). The models are dependent on the distributions of galaxy merger rates at a given epoch. This function is not well constrained by current observations and two possible realizations are applied. Parametrically, the merger rate function is described by a double power law and the differing models assume differing slopes for the low-mass end. The first assumes a steep drop

Figure 10. Comparison between our data and the models of Hopkins et al. (2005b). The model lines give the predicted dispersion in log (M_{BH}) as a function of luminosity at the mean redshift of the bin (top right of each panel). The two different curves give the predictions from differing merger rate functions, either with a sharp cut-off at lower masses (solid) or with a flatter cut-off (dashed). Our measured dispersion in log (IPV) is multiplied by two for the comparison (equation 3).

off of merging systems towards lower mass and gives only a narrow mass range for merging systems at a given redshift. Secondly, a broad range of systems is assumed to be merging at any one time with a flat slope at the low-mass end of the merger function. We label these models as sharp or flat as describes their low-mass slope; the sharp and flat models are shown as the solid and dashed lines, respectively, in Fig. 10.

We find that our data points match the Hopkins et al. (2005b) models extremely well in general. However, the models exhibit some evolution with redshift which we do not observe in our data.

These models are, to an extent, bound to the QSO luminosity function which shows strong redshift evolution. Over the redshift range, we sample the break in the luminosity function, M^*, changes by more than a magnitude, and the Hopkins et al. models follow this to an extent. We find no such trend in the relation shown in Figs 7 and 10. To illustrate this, Fig. 11 shows exactly the same data as Fig. 7(d) except the dispersion is plotted against $M^* - M$, i.e. the object's luminosity relative to the break in the luminosity function. For this plot, we take the quadratic parametrization of M^* from Croom et al. (2004). Comparisons between Figs 7(d) and 11 clearly show that the observed trend towards lower dispersion in broad-line width for brighter objects is dependent on the object's luminosity, and not its position on the luminosity function.

The Hopkins et al. models evolve with redshift in a similar sense to the QSO luminosity function albeit less strongly, hence we observe a slight discrepancy. However, it was unlikely that these simple models would be exact in their determination of the dispersion in M_{BH}. Of more importance is the trend, which is predicted by the models and well echoed in the data. Our data show that more luminous QSOs show less scatter in their Mg II linewidths than those of


Figure 11. Same data as Fig. 7(d), except the dispersion is plotted versus $M^* - M$ rather than M. Comparing the scatter evident in this plot with Fig. 7(d) clearly demonstrates the observed trend is dependent on the brightness of a quasar, and not its position on the luminosity function. In this plot, we use the quadratic parametrization for M^* given in Croom et al. (2004).

lower luminosity, implying they have less scatter in their black hole masses.

6.2 Is this trend a selection effect?

Recently, Babić et al. (2007) showed that a broad intrinsic distribution of Eddington ratios which truncates at the Eddington rate, convolved with a double power-law SMBH mass function for QSOs, naturally leads to a selection effect such that samples with a fainter flux limit will find a broader range of Eddington ratios.

Figure 12. These plots illustrate the effect described by Babić et al. (2007). The dispersion in SMBH mass at a fixed luminosity is plotted assuming Eddington ratio distributions with means $\langle\lambda\rangle = 0.3$ and 0.7, and with a variety of widths (σ_λ).

To test whether this effect could be responsible for the trend observed in Fig. 7, we recreate this situation and compare it with our results. In our analysis, we assume that the distribution of Eddington rates ($L/L_{Edd} = \lambda$) is lognormal with a given mean, $\langle\lambda\rangle$, and dispersion, σ_λ, and is truncated at $\lambda = 1$. To begin with, we assume the SMBH mass function can be described as a double power law, and constrain this by convolving with the Eddington distribution and fitting to the observed luminosity function (we use the Hopkins et al. 2007a *B*-band luminosity function at redshift 1), producing a (almost) unique solution for the mass function. We can then derive the dispersion in SMBH mass at a given luminosity.

We find that the effect discussed by Babić et al. (2007) is reproduced in the situation where $\langle\lambda\rangle$ and σ_λ conspire such that the $\lambda = 1$ cut is not too many σ_λs away from $\langle\lambda\rangle$; i.e. in situations where the cut has a significant effect on the Eddington distribution. Hence, in situations where $\langle\lambda\rangle$ is small σ_λ must be large to produce the effect. If σ_λ is small, then $\langle\lambda\rangle$ must be comparable to 1 for a pronounced effect. Fig. 12 gives an example of this. We plot the dispersion in SMBH mass in these models as a function of luminosity in the case where $\langle\lambda\rangle = 0.3$ and 0.7, for a variety of values of σ_λ.

While both of these plots show a decrease in the dispersion in SMBH mass with luminosity, neither give a good representation of our results. In the case where $\langle\lambda\rangle = 0.3$ σ_λ must be so high to get a pronounced gradient that the normalization of the relation, in particular at the bright end, is too high when compared with our results. In the case where $\langle\lambda\rangle = 0.7$, σ_λ can be smaller and the normalization is closer to that observed. However, the model dispersions flatten a magnitude brighter than M^*, an effect we do not observe.

In all cases, the dispersion asymptotes to a constant value above M^* since at these luminosities one is sampling both from a luminosity function and by extension SMBH mass function, which follow a single power law. The Babić et al. effect does not occur without sampling objects from a mass function which is steepening in the log sense.

This is, however, an artefact of our assigning a double power law for SMBH mass function. If we assume a Schechter function instead the high-mass end drops off exponentially and the Babić et al. effect is apparent at the higher luminosities sampled. On the other hand, with this model the fit to the observed QSO luminosity function (which does follow a double power law) is significantly degraded.

Finally, since these models are tied directly to the luminosity function (i.e. are relative to M^*), they evolve strongly with redshift. The observed evolution in the QSO luminosity function must be due to either a similarly evolving active SMBH mass function or an equivalently varying accretion efficiency distribution. Either will strongly affect the Babić et al. effect.

As discussed in Section 5.1, we find no evidence for a redshift dependence in our results. As shown in Fig. 11, the observed correlation between dispersion in IPV widths and luminosity is degraded when plotted versus luminosity relative to M^*.

In light of the above discussion, it is clear that while the Babić et al. (2007) effect may well bias our results, it does not accurately describe both the results presented in Fig. 7 and the observed luminosity function simultaneously. The lack of observed evolution in our results may suggest that this effect is not significant; however, without a better understanding of the underlying SMBH mass function or accretion efficiency distribution, it is difficult to gauge the magnitude of this effect and we cannot rule out a significant bias in our data.

6.3 Disp(log(IPV)) versus Disp(log(M_{BH})) and virial masses revisited

While we observe a strikingly good agreement between our data and the Hopkins et al. (2005b), we caution against reading too much into the normalization of our data with respect to the dispersion in SMBH mass. In fact, a closer examination of this normalization raises a question as to just how accurately the dispersion in log (IPV) traces that in log (M_{BH}).

In using equation (3), we are assuming there are no other sources of scatter in the QSO SMBH population which will affect these results. This is almost certainly not the case and we must consider where other sources of scatter might occur, their magnitude and their effect on our results. Shen et al. (2008) recently performed simulations of the QSO SMBH population to investigate the biases associated with virial mass estimations, and we follow a similar path to their work.

They describe the QSO and SMBH populations with four variables. The SMBH mass, the bolometric luminosity (and by extension the Eddington ratio), the monochromatic luminosity as used in the virial relation (equation 2) and the FWHM of the spectral line being analysed (in our case the IPV width of Mg II). We add to this the radius and velocity dispersion of the BLR: r_{BLR} and v_{BLR}. There are possible sources of scatter to each of these variables and here we discuss each of them, their likely magnitude and their effect on our measurements.

(i) **Scatter in luminosity**

The primary source of scatter in luminosity at a fixed SMBH mass comes from the distribution of accretion efficiencies for QSOs. This is precisely the scatter we are trying to measure in this work and is not a contaminant to our calculations.

The only other sources of scatter to the luminosity are photometric errors, host galaxy light and extinction. Each of these would add extra scatter into the luminosities we measure independently from any virial mass equations. Hence, this scatter would not propagate through to the linewidths we measure in the way scatter in accretion rate would. This has the unique effect of making the SMBH mass distribution we measure at a given luminosity narrower relative to the intrinsic mass distribution (see Shen et al. 2008).

However, these sources of scatter can only play a minor role in our measurements and calculations. Photometric errors are ~0.1 mag in the b_J band used by the 2QZ (~0.04 dex in L; ~ 0.02 dex in M_{BH}) and smaller for the SDSS observations. Host galaxy emission and extinction will only affect the fainter objects in our sample, and even at these luminosities will not be a major factor. Hence, these sources of scatter will not be significant to our calculations.

Shen et al. also include a scatter due to variations in continuum shape when converting bolometric luminosities to monochromatic luminosities. They quote a value for this of ~0.1 dex translating to ~0.05 dex scatter in M_{BH} and so again this is not of great importance to our analysis.

(ii) **Scatter in r_{BLR}**

The r–L relation shows that the radius of the BLR around QSOs is driven by the luminosity of the source. Hence, any scatter in the intrinsic luminosity of the source (i.e. not due to photometric errors/host galaxy emission/extinction) will also scatter r_{BLR}. Furthermore, the r–L relation has been observed to exhibit intrinsic scatter of its own at a level of 40 per cent (Kaspi et al. 2005) translating to at least 0.15 dex, and this intrinsic scatter will further broaden the r_{BLR} distribution. This extra scatter will propagate through to the IPV width distribution we measure, and will bias our final measurement for the scatter in M_{BH} high by the same factor. 0.15 dex of scatter in M_{BH} is significant to our discussion and we expand on the implications of this in Section 6.3.1.

(iii) **Scatter in v_{BLR}**

The distribution of velocity dispersions will have all of the above scatters folded into it apart from the scatter in luminosity due to photometric errors/host galaxy emission/extinction. If we take v_{BLR} to be the virial velocity of the BLR, then there will be no additional sources of scatter which affect this quantity since the virial equation (equation 1) is exact.

However, other (non-virial) factors may affect the velocity of the BLR. Shen et al. include an additional scatter to the linewidth distribution due to non-virial velocities in the BLR. This will bias the scatter we measure in IPV widths (and hence M_{BH}) high, although in this work we do not attempt to investigate this further.

(iv) **Scatter in IPV widths**

Beyond all of the scatters which affect v_{BLR}, the IPV width distribution we measure can be affected by further sources of scatter. First, errors on our measurements will artificially broaden the measured IPV width distribution, and we account for this in our dispersion analysis (see Section 4.3 and equation 7).

Secondly, the linewidth we measure may not accurately describe the velocity of the BLR. Selective absorption (e.g. the model proposed for the C IV BLR by Richards et al. 2002) and/or orientation-dependent linewidths would further broaden the measured IPV width distribution. There is little evidence for absorption in the Mg II line for non-BAL QSOs, and we do not attempt to model its effects here. The potential effects of orientation on our measurements are discussed in Section 6.4.

Of these sources of scatter, only the uncorrelated scatter in luminosity can bias the measured dispersion in IPV widths low. All of the other sources of scatter propagate through to the IPV distribution and will bias the results high.

Little is known about the magnitude of many of the sources of scatter discussed above, but most can be ruled out as negligible when compared to the observed dispersion in IPV widths. The only source of scatter which is known to be comparable to that measured for the IPV widths is the intrinsic scatter in the r–L relation.

6.3.1 Implications of intrinsic scatter in the r–L relation

Equation (3) assumes that there is no variation in r_{BLR} within a luminosity bin. This is not the case, and the distribution of linewidths in a given luminosity–redshift bin is not directly analogous to the intrinsic distribution of SMBH masses. Instead, the distribution of linewidths is a convolution of the black hole mass distribution and the distribution of BLR radii in a given bin (as well as our error distribution; equation 7).

With this in mind, a better estimate of the true dispersion in M_{BH} is the calculated dispersion in log (IPV) ($\times 2$; equation 3) with both the dispersion in luminosities in a single bin ($\times \alpha$; equation 2) and the intrinsic scatter in the r–L relation subtracted from it in quadrature.

The scatter in log (L) for 0.5 mag bins will be ~0.1 dex in the faint bins, and decrease to brighter luminosities as the luminosity function steepens. This translates to <0.05 dex of scatter in M_{BH} and hence is relatively unimportant to this discussion.

On the other hand, the intrinsic scatter in the r–L relation is not negligible with respect to our results. Since the radius of the BLR is dependent on the luminosity of the quasar (and not the other way around), any intrinsic scatter in this relation will also tend to bias our results high. Due to the relatively small number of objects with good reverberation data, the scatter in the r–L relation is not definitively constrained, and not defined at all for the brighter luminosities in our sample. However, Kaspi et al. (2005) do measure this for the Hβ BLR with 35 of the best mapped AGN and quote a value of ~40 per cent for the intrinsic scatter in r_{BLR}. This corresponds to at least ~0.15 dex in the log (0.17 if the error is propagated as in equation 6), and is greater than the scatter we find in SMBH mass in the brightest magnitude bins.

This leaves a dilemma: either there is significantly less intrinsic scatter in the r–L relation than that quoted by Kaspi et al. (at least at high L) or there is an intrinsic problem with our approach to finding the scatter in M_{BH}.

The 40 per cent scatter quoted by Kaspi et al. (2005) is not directly applicable to our analysis for two reasons: (1) the r–L relation studied in that work is for the Hβ emission region not Mg II and (2) the r–L relation is only defined for relatively faint Seyfert 1 type objects, while the lowest scatter we find is for the most luminous quasars in our sample. In effect, our use of the Kaspi et al. value for the scatter in this relation equates to an extrapolation over ~2 orders of magnitude in luminosity.

In their analysis, Kaspi et al. suggest possible sources of scatter in the r–L relation such as intrinsic reddening, contributions by the host galaxy or effects of variability due to non-contemporaneous observations. Of these, galaxy contamination and potentially reddening will play a smaller role at higher luminosities. Bentz et al.

(2006) used *HST* imaging to correct a subset of the objects in Kaspi et al.'s sample for host galaxy emission and found that the intrinsic scatter in the *r*–*L* relation could be as low as ∼30 per cent. At higher luminosities, where host contamination would be lessened, it is possible that the scatter is even less.

If there were no scatter in the *r*–*L* relation, our data imply there is only ∼0.14 dex scatter in M_{BH} in luminous QSOs. This may be plausible considering that at the extreme luminosities we are considering here, we would expect most quasars to be radiating at or around their Eddington luminosity.

An alternative interpretation may be that the dispersion in log (*IPV*) does not properly represent the dispersion in log(M_{BH}). That is, QSO broad lines do not show enough variation to account for the expected variation in M_{BH} when accounting for the intrinsic scatter in the *r*–*L* relation.

This raises a question as to the validity of the analysis performed here, and the virial method for estimating M_{BH}.

There is evidence for a virialized BLR. Time lags between continuum and emission line variations have been observed to be shorter in the wings as opposed to the core of lines (Kollatschny & Dietrich 1996; Onken & Peterson 2002; Kollatschny 2003). The same authors also find that in objects which have had more than one line mapped, their time delays appear to follow a virial ($\tau \propto V_{BLR}^{-2}$) relation. In addition, studies comparing virial masses to bulge velocity dispersion have shown a correlation analogous to that found in local quiescent galaxies (Onken et al. 2004; Woo et al. 2006) indicating that the virial method works as a tracer of M_{BH}.

An explanation could be that the BLR of QSOs is not *wholly* virialized. There is a significant component of the BLR which shows very little object-to-object velocity variation. If this were the case, virial motion would still be the primary cause of the variation between broad-line widths in QSOs, and observations such as the *M*–σ relation would be reproduced. However, globally we would not see a range of linewidths comparable to the range of black hole masses.

Alternatively, and more controversially, it may be that the BLR is not virialized at all and the virial relations only prove accurate due to their luminosity dependence. The SMBH mass of a QSO and its luminosity must be strongly covariant and the ability of virial estimators to determine M_{BH} may be due to this simple relation. The linewidth term may be redundant. Indeed, we (and other authors, e.g. Corbett et al. 2003; Shen et al. 2008) find the average linewidth changes very little over the luminosity range of our sample. To be consistent with the virial mass estimators, this requires that both SMBH mass and Eddington ratio for QSOs vary as luminosity to the power of ∼1/2. A conspiracy which must be viewed with some caution.

The reason for the lack of observed scatter in broad-line widths is uncertain, but it is apparent that this observation has consequences for black hole mass estimation in AGN. While this may imply that virial masses are not necessarily unbiased estimators of M_{BH}, there is evidence from other studies that they can be used as an indicator of the mass of a black hole.

6.4 Constraining BLR geometry

The low dispersion in broad-line width found for luminous QSOs is also a strong constraint on the velocity field of the BLR.

In all but the case of a spherically symmetric BLR, the velocity dispersion we measure from a spectrum will depend on the viewing angle to the QSO. Given a model for the BLR, we can calculate the expected dispersion in measured linewidths due to variations in observing angle. If the expected dispersion is greater than we find


Figure 13. A sketch showing the assumed geometry of AGN. The central SMBH and BLR are surrounded by an obscuring torus which constrains the allowed observing angle (θ) to be less than some opening angle (θ_m).

in our data, we can rule out that model for the BLR. Since we take no account of the myriad of other sources of scatter in the linewidth distribution, this equates to a very strong constraint on a given BLR model.

6.4.1 Pure planar/polar BLRs

The first models we consider are planar velocity fields in which all the velocities are confined to the $z = 0$ plane (cylindrical coordinates), and polar fields in which all the velocities are in the $\hat{z}$ direction. In these cases, we expect our measured IPV widths to be modified by

Planar: $\log(IPV) \propto \log(\sin(\theta))$

Polar: $\log(IPV) \propto \log(\cos(\theta))$,

where θ is our observing angle to the QSO (see Fig. 13). Assuming an opening angle (θ_m), which we take to be constrained by an obscuring torus (Fig. 13), we can calculate what the expected dispersion in measured linewidths would be due solely to orientation effects for planar and polar velocity fields.

Fig. 14 shows how much scatter in linewidths we expect in each case. On both plots, we indicate with a dashed line 0.07 dex dispersion which is the lowest we measure in our sample (Fig. 7).

In the planar case, the expected dispersion is ∼0.2 dex for all opening angles. We find somewhat less dispersion in IPV widths than this in all but two of our *L*–*z* bins. Hence, the velocity field of the BLR cannot solely be planar but must include some other component.

For the polar velocity field, the expected dispersion is a strong increasing function of θ_m. For modest opening angles, the resulting dispersion is low and it is not until $\theta_m \sim 55°$ that this becomes comparable to the dispersions we measure in our data. Hence, if the BLR is characterized by polar flows, the opening angle to luminous quasars must be less than this value.

6.4.2 Planar/polar BLRs with a random/spherically symmetric component

It is not likely that in any model for the BLR all the velocities are confined to a single plane/direction. Instead, there will always be some random component to the velocity field, and we include this

Figure 14. These plots show how the dispersion in measured linewidth is dependent on the opening angle of the source in the special cases where all of the BLR velocities are in either the planar or polar directions. On each plot, the dashed line shows the dispersion = 0.07 dex level, the lowest dispersion in linewidth we measure in our data.

in our models with the following parametrization:

Planar: $\log(IPV) \propto \log(a\sin(\theta) + (1-a))$

Polar: $\log(IPV) \propto \log(a\cos(\theta) + (1-a))$.

In these models, the trigonometric term represents a geometrically constrained component to the BLR velocity, and the $(1-a)$ term represents a random/spherically symmetric component. Hence, a model with $a = 1$ represents the pure planar/polar cases outlined above, and if $a = 0$ the BLR velocity field is spherically symmetric. In these cases, given values for a and θ_m, we can again calculate the expected dispersion in $\log(IPV)$ to compare with our results. These calculations are illustrated in Fig. 15.

In this figure, the grey-scale indicates the expected scatter at a point in the $a - \theta_m$ plane. The contours are at increments of 0.035 dex, and so the second contour (heavy line) represents the 0.07 dex scatter we observe for the most luminous QSOs in our sample. Hence, for these luminous objects, the region of parameter space above the solid line is ruled out in our analysis.

In the planar case, for which the constraints are stronger, taking a believable value for the opening angle ($\sim 45°$), our data show that roughly half of the contribution to the velocity field of the BLR must come from a symmetric component.

6.4.3 A hybrid BLR

As a final model, we consider a BLR with both planar and polar components to the velocity field as well as a spherically symmetric component. We model this with

$$\log(IPV) \propto \log(a\sin(\theta) + b\cos(\theta) + (1-a-b)).$$

So, for a given (a, b, θ_m), we can calculate the expected dispersion in measured linewidth. Fig. 16 shows the expected dispersion in the $a - b$ plane for differing opening angles. In each plot, the origin ($a = b = 0$) represents a spherically symmetric BLR and the line $a + b = 1$ represents a BLR with no symmetric component. As one moves from left to right in each plot, the BLR becomes more planar, and from bottom to top the velocity field is more polar.

As expected, for small opening angles our data are not strong constraints on the velocity field of the BLR, and only the most disc-like models are rejected. For larger opening angles there is a larger region of the parameter space ruled out by our data, predominantly the very planar or very polar models.

7 CONCLUSIONS

We have measured the 50 per cent IPV width of the Mg II line in QSO spectra from the SDSS, 2QZ and 2SLAQ surveys and find a strong correlation between the dispersion in IPV widths and the optical luminosity of QSOs. If we assume there exists a virial relation of the form of equation (2), this implies that there is an equivalent reduction in the dispersion in M_{BH}. At face value, this is in excellent agreement with models for the QSO population proposed by Hopkins et al. (2005b).

However, the remarkably low scatter we find in our IPV width measurements, in particular for the more luminous objects, has implications as to the validity of virial mass estimators. We find less scatter in IPV widths of luminous QSOs than is intrinsic to the r–L relation. While the r–L relation is not defined for these very luminous objects, these results are at odds with the practise of virial black hole estimation. Possible explanations for this observation include very low scatter in the r–L relation for bright QSOs or a BLR which is not fully virialized. In either case, it is clear we are yet to gain a full understanding of the process of virial black hole mass estimation, and one must be cautious when performing or interpreting these estimates.

Finally, we show that the observed scatter in IPV widths can be used to constrain models for the velocity field of the BLR. We show how variations in observing angle to sources affect the scatter in measured IPV widths for a series of simplified geometries and derive constraints on these models from our data. Perhaps of most interest is the rejection of a pure planar BLR regardless of the assumed opening angle to the source.

A table containing the Mg II IPV widths for the quasars studied in this work is available through the 2SLAQ web site (http://www.2slaq.info/qso_ipv.dat).

Figure 15. Here, we demonstrate how adding a spherically symmetric component to the BLR velocity field affects the expected dispersion in broad-line widths. Again, we assume that all of the other velocities are either in the planar or in the polar directions. In our parametrizations, $a \sim 0$ indicates a BLR dominated by a spherically symmetric velocity field, while $a \sim 1$ indicates no symmetric component equivalent to the models shown in Fig. 14. The grey-scale and contours indicate the expected dispersion. Contour levels are in increments of 0.035 dex, and the second contour (heavy) shows the 0.07 dex scatter we measure for the most luminous objects in our sample. The parameter space above this is inconsistent with our analysis.
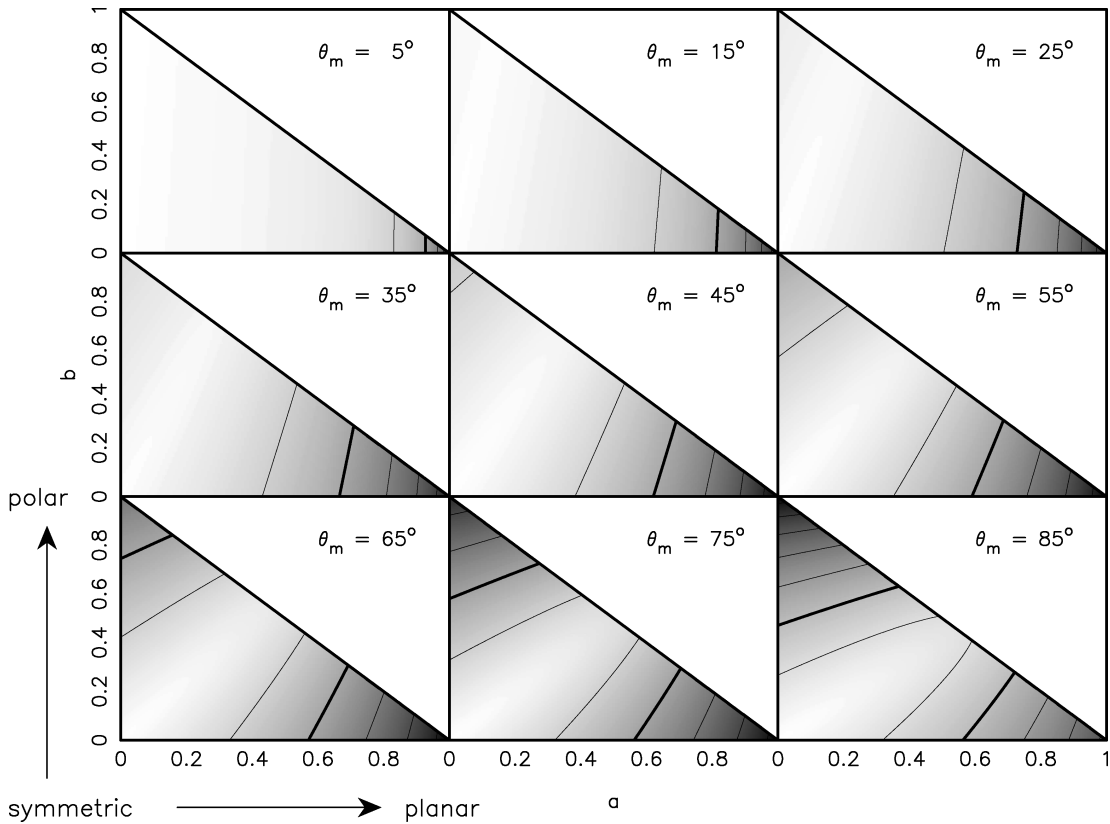
Figure 16. These plots show the expected dispersion in measured linewidths for our composite BLR model which include planar, polar and symmetric components. Each plot shows the expected dispersion (grey-scale + contours as in Fig. 15) for a given AGN opening angle (θ_m). In our parametrization, as a increases the BLR is more planar, and as b increases more polar. (0,0) indicates a totally symmetric velocity field while all models on the line $a + b = 1$ have no symmetric component.

REFERENCES

Adelman-McCarthy J. K. et al., 2007, ApJS, 172, 634
Babić A., Miller L., Jarvis M. J., Turner T. J., Alexander D. M., Croom S. M., 2007, A&A, 474, 755
Bailey J., Glazebrook G., 1999, 2dF User Manual. Anglo-Australian Observatory
Barnes J. E., Hernquist L., 1991, ApJ, 370, L65
Barnes J. E., Hernquist L., 1996, ApJ, 471, 115
Baskin A., Laor A., 2004, MNRAS, 350, L31
Bentz M. C., Peterson B. M., Pogge R. W., Vestergaard M., Onken C. A., 2006, ApJ, 644, 133
Boroson T. A., Green R. F., 1992, ApJS, 80, 109
Corbett E. A. et al., 2003, MNRAS, 343, 705
Cristiani S., Vio R., 1990, A&A, 227, 385
Croom S. M., Smith R. J., Boyle B. J., Shanks T., Loaring N. S., Miller L., Lewis I. J., 2001, MNRAS, 322, L29
Croom S. M., Smith R. J., Boyle B. J., Shanks T., Miller L., Outram P. J., Loaring N. S., 2004, MNRAS, 349, 1397
Di Matteo T., Springel V., Hernquist L., 2005, Nat, 433, 604
de Francesco G., Capetti A., Marconi A., 2006, A&A, 460, 439
Elvis M., 2004, in Richards G. T., Hall P. B., eds, ASP Conf. Ser. Vol. 311, AGN Physics with the Sloan Digital Sky Survey. Astron. Soc. Pac., San Francisco, p. 109
Emmering R. T., Blandford R. D., Shlosman I., 1992, ApJ, 385, 460
Fan X. et al., 2001, AJ, 121, 54
Faucher-Giguere C.-A., Prochaska J. X., Lidz A., Hernquist L., Zaldarriaga M., 2008a, ApJ, 681, 831
Faucher-Giguere C. A., Lidz A., Hernquist L., Zaldarriaga M., 2008b, ApJ, 682, L9
Ferrarese L., Merritt D., 2000, ApJ, 539, L9
Fromerth M. J., Melia F., 2000, ApJ, 533, 172
Gebhardt K. et al., 2000, ApJ, 539, L13
Gebhardt K. et al., 2003, ApJ, 583, 92
Graham A. W., Erwin P., Caon N., Trujillo I., 2001, ApJ, 563, L11
Gunn J. E. et al., 2006, AJ, 131, 2332
Hernquist L., Springel V., 2003, MNRAS, 341, 1253
Herrnstein J. R. et al., 1999, Nat, 400, 539
Hopkins A. M., Beacom J. F., 2006, ApJ, 651, 142
Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Martini P., Robertson B., Springel V., 2005a, ApJ, 630, 705
Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2005b, ApJ, 630, 716
Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2005c, ApJ, 632, 81
Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2006a, ApJS, 163, 1
Hopkins P. F., Hernquist L., Cox T. J., Robertson B., Springel V., 2006b, ApJS, 163, 50
Hopkins P. F., Richards G. T., Hernquist L., 2007a, ApJ, 654, 731
Hopkins P. F., Hernquist L., Cox T. J., Robertson B., Krause E., 2007b, ApJ, 669, 45
Hopkins P. F., Hernquist L., Cox T. J., Robertson B., Krause E., 2007c, ApJ, 669, 67
Hopkins P. F., Hernquist L., Cox T. J., Keres D., 2008, ApJS, 175, 356
Kaspi S., 2007, in Ho L. C., Wang J.-M., eds, ASP Conf. Ser. Vol. 373, The Central Engine of the AGN. Astron. Soc. Pac., San Francisco, p. 13
Kaspi S., Smith P. S., Netzer H., Maoz D., Jannuzi B. T., Giveon U., 2000, ApJ, 533, 631
Kaspi S., Maoz D., Netzer H., Peterson B. M., Vestergaard M., Jannuzi B. T., 2005, ApJ, 629, 61
Kollatschny W., 2003, A&A, 407, 461
Kollatschny W., Dietrich M., 1996, A&A, 314, 43
Kollmeier J. A. et al., 2006, ApJ, 648, 128
Königl A., Kartje J. F., 1994, ApJ, 434, 446
Kurk J. D. et al., 2007, ApJ, 669, 52
Lauer T. R., Tremaine S., Richstone D., Faber S. M., 2007, ApJ, 670, 249
Lewis I. J. et al., 2002, MNRAS, 333, 279
Magorrian J. et al., 1998, AJ, 115, 2285
McGill K. L., Woo J.-H., Treu T., Malkan M. A., 2008, ApJ, 673, 703
McLure R. J., Dunlop J. S., 2001, MNRAS, 327, 199
McLure R. J., Dunlop J. S., 2004, MNRAS, 352, 1390
McLure R. J., Jarvis M. J., 2002, MNRAS, 337, 109
Murray N., Chiang J., 1997, ApJ, 474, 91
Onken C. A., Peterson B. M., 2002, ApJ, 572, 746
Onken C. A., Ferrarese L., Merritt D., Peterson B. M., Pogge R. W., Vestergaard M., Wandel A., 2004, ApJ, 615, 645
Peterson B. M., Wandel A., 1999, ApJ, 521, L95
Peterson B. M. et al., 2004, ApJ, 613, 682
Press W. H., Flannery B. P., Teukolsky S. A., Vetterling W. T., 1989, Numerical Recipes. Cambridge Univ. Press, Cambridge
Richards G. T., Vanden Berk D. E., Reichard T. A., Hall P. B., Schneider D. P., SubbaRao M., Thakar A. R., York D. G., 2002, AJ, 124, 1
Richards G. T. et al., 2005, MNRAS, 360, 839
Richards G. T. et al., 2006, ApJS, 166, 470
Robertson B., Hernquist L., Cox T. J., Di Matteo T., Hopkins P. F., Martini P., Springel V., 2006, ApJ, 641, 90
Salviander S., Shields G. A., Gebhardt K., Bonning E. W., 2007, ApJ, 662, 131
Schneider D. P. et al., 2007, AJ, 134, 102
Shen Y., Greene J. E., Strauss M., Richards G. T., Schneider D. P., 2008, ApJ, 680, 169
Sigut T. A. A., Pradhan A. K., 2003, ApJS, 145, 15
Springel V., Hernquist L., 2003, MNRAS, 339, 312
Stoughton C. et al., 2002, AJ, 123, 485
Vestergaard M., 2002, ApJ, 571, 733
Vestergaard M., Peterson B. M., 2006, ApJ, 641, 689
Vestergaard M., Wilkes B. J., 2001, ApJS, 134, 1
Wang T. G., Lu Y. J., Zhou Y. Y., 1998, ApJ, 493, 1
Welsh W. F., Horne K., 1991, ApJ, 379, 586
Whittle M., 1985, MNRAS, 213, 1
Wilhite B. C., Brunner R. J., Schneider D. P., Vanden Berk D. E., 2007, ApJ, 669, 791
Woo J. H., Treu T., Malkan M. A., Blandford R. D., 2006, ApJ, 645, 900
York D. G. et al., 2000, AJ, 120, 1579

This paper has been typeset from a TₑX/LₐTₑX file prepared by the author.