

Quasar clustering and duty cycle measurements at $0 \leq z \leq 4$ with the Gaia-unWISE Catalog

MARIONA GINER MASCARELL ^{1,2}, ANNA-CHRISTINA EILERS ^{2,3} AND KATE STOREY-FISHER ⁴

¹*Department of Physics, University of Richmond, Richmond, Virginia, VA 23173, USA*

²*Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

³*MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

⁴*Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, 452 Lomita Mall, Stanford, CA 94305, USA*

ABSTRACT

We measure the two-point correlation function of a uniformly selected, all-sky sample of ~ 1.3 million quasars with magnitudes $G \leq 20.5$ from the *Gaia*–unWISE Quasar Catalog (Quaia) over the redshift range $0 \leq z \leq 4$ to trace the evolution of the quasar clustering strength across cosmic time. We find a steady increase in the correlation length r_0 with redshift, i.e. $r_0 = 6.8 \pm 0.2 h^{-1} \text{Mpc}$ at $0 \leq z < 1$, $r_0 = 8.0 \pm 0.2 h^{-1} \text{Mpc}$ at $1 \leq z < 2$, $r_0 = 10.8 \pm 0.2 h^{-1} \text{Mpc}$ at $2 \leq z < 3$, and $r_0 = 13.9 \pm 1.2 h^{-1} \text{Mpc}$ at $3 \leq z < 4$, and slopes consistent with $\gamma \approx 2$. Our measurements suggest a slightly weaker clustering signal at $z > 3$ than previous studies, and thus we find a smooth, monotonic rise in clustering strength. Using a bias-halo mass relation and a step-function for the halo occupation distribution, we infer characteristic minimum halo masses of quasar hosts of $\log_{10}(M_{\text{min}}/M_{\odot}) \approx 12.8$ across all redshifts. Combining these with the observed quasar number densities yields duty cycles that rise from $f_{\text{duty}} \approx 2\%$ to $\approx 7\%$ with increasing redshift, corresponding to integrated quasar lifetimes of $t_{\text{QSO}} \sim 10^8$ years. These results suggest that both the characteristic halo mass of active quasars and their typical lifetimes have remained remarkably stable over more than $12 \sim \text{Gyr}$ of cosmic time, implying a self-regulated growth process largely independent of epoch.

1. INTRODUCTION

Nearly all galaxies in the local universe are known to host central supermassive black holes (SMBHs). Well-established empirical correlations between SMBH mass and their host properties (e.g., bulge mass and velocity dispersion) point towards coupled growth between the SMBHs and their hosts (J. Kormendy & L. C. Ho 2013; N. J. McConnell & C.-P. Ma 2013). The classical Soltan argument (A. Soltan 1982) links the time galaxies shine as luminous quasars and during which the bulk of the black hole growth is believed to occur to the local SMBH mass density (see also Q. Yu & S. Tremaine 2002), suggesting that a substantial fraction of SMBH assembly occurs during radiatively efficient, UV-luminous quasar episodes.

Quasars are expected to trace the most massive peaks of the cosmic density field, serving as powerful probes between baryonic mass and the underlying dark matter density field (e.g. M. White et al. 2008; P. F. Hopkins et al. 2006). The quasars’ clustering amplitude and evolution encode how SMBHs populate dark matter halos and how accretion onto the SMBHs is triggered across

cosmic time. In this work, we will study the clustering of quasars focusing on the redshift interval $0 \leq z \leq 4$, which spans the rise and decline of cosmic star formation and quasar activity (P. Madau & M. Dickinson 2014), providing a key window into both the assembly of massive galaxies and the growth of their central black holes.

By measuring the clustering strengths of quasars via the auto-correlation function we will determine their correlation length scale r_0 , and the quasars’ minimum host dark matter halo mass M_{min} . Our cosmological model ΛCDM provides a framework to link the host halo mass to their number density, and thus a comparison between the number densities of observed quasars to the number densities of their host dark matter halos inferred from clustering studies will also allow us to constrain the quasars’ duty cycle, which is the fraction of cosmic time that SMBHs grow as luminous quasars (P. Martini & D. H. Weinberg 2001; C. Conroy & M. White 2013). In particular, we can estimate the duty cycle as

$$f_{\text{duty}} \simeq \frac{t_{\text{QSO}}}{t_{\text{H}}(z)} \simeq \frac{n_{\text{QSO}}}{n_{\text{DM}}(> M_{\text{min}})} \quad (1)$$

where t_{QSO} denotes the typical quasar lifetime, $t_{\text{H}}(z)$ is the Hubble time at redshift z , n_{QSO} is the observed co-moving number density of quasars, and $n_{\text{DM}}(> M_{\text{min}})$

is the comoving number density of dark matter haloes above the characteristic minimum host mass inferred from the clustering analysis.

Previous clustering analyses have established this framework for quasars across a wide range of redshifts and luminosities. At moderate redshift ($0.8 \leq z \leq 2.1$), the Sloan Digital Sky Survey (SDSS) Data Release 5 (DR5; Y. Shen et al. 2007) provided one of the first large, uniform quasar clustering measurements, based on a sample of $\sim 30,000$ quasars. Later SDSS releases extended these analyses with larger samples and improved redshift coverage (e.g., M. White et al. 2012; P. Laurent et al. 2017), but the derived clustering amplitudes remain broadly consistent with DR5 results. These studies find correlation lengths of $r_0 = 5.45^{+0.35}_{-0.45} h^{-1} \text{Mpc}$, corresponding to characteristic halo masses of order $10^{12} h^{-1} M_\odot$ and duty cycles of a few percent. At higher redshifts ($2.2 \leq z \leq 2.8$), S. Eftekharzadeh et al. (2015) measured $r_0 = (8.12 \pm 0.22) h^{-1} \text{Mpc}$, consistent with halo masses of $\sim 2 \times 10^{12} h^{-1} M_\odot$ and similarly low duty cycles, with little dependence on luminosity. However, at higher redshifts of $3 \lesssim z \lesssim 4$, clustering analyses of SDSS quasars have revealed a pronounced increase in bias with redshift (Y. Shen et al. 2007; M. White et al. 2012; S. Eftekharzadeh et al. 2015; J. Arita et al. 2023). For instance, Y. Shen et al. (2007) found $r_0 = 16.9 \pm 1.7 h^{-1} \text{Mpc}$ at $2.9 \leq z \leq 3.5$ and $r_0 = 24.3 \pm 2.4 h^{-1} \text{Mpc}$ at $z \geq 3.5$, corresponding to halo masses of $\sim 2\text{--}6 \times 10^{12} h^{-1} M_\odot$ and quasar lifetimes of several hundred Myr, implying duty cycles approaching unity.

Beyond $z \gtrsim 6$, quasar auto-correlation measurements become increasingly challenging due to the low number density of luminous quasars. Using a faint ($M_{1450} \gtrsim -25$), but more numerous population of quasars a first attempt of the auto-correlation measurement at $z \sim 6$ has recently been reported by the Subaru High- z Exploration of Low-luminosity Quasars (SHELLQs) Collaboration (J. Arita et al. 2023), indicating host dark matter halo masses similar or slightly higher than those found at lower redshifts, but uncertainties are large. However, while auto-correlation measurements of quasars at these high redshifts remain challenging, complementary quasar-galaxy cross-correlation measurement is still viable. Recent measurements of quasar-galaxy cross-correlations leveraging data from the James Webb Space Telescope (JWST) reveal similar minimum host masses of $\log_{10}(M_{\text{halo,min}}/M_\odot) = 12.43^{+0.13}_{-0.15}$ to those in the lower redshift universe, and extremely short UV-bright duty cycles $f_{\text{duty}} \lesssim 1\%$, suggesting only brief luminous black hole growth phases (A.-C. Eilers et al. 2024; E. Pizzati et al. 2024). At even earlier cosmic times, the

picture becomes less clear; for example, a recent study of two quasar fields at $z \approx 7.3$ suggests a decrease in host halo mass down to $\log_{10}(M_{\text{halo,min}}/M_\odot) \sim 11.6 \pm 0.6$ and duty cycles of $f_{\text{duty}} \sim 0.1\%$ (J.-T. Schindler et al. 2025), potentially indicating a non-monotonic redshift evolution. Despite these advances, current measurements are highly heterogeneous in their sample selection, making direct comparisons across redshift challenging. Thus, a homogeneous clustering analysis is essential to understand how quasar activity and host halo properties evolve across cosmic history to the present day.

In this work, we leverage the *Gaia*-unWISE Quasar Catalog (“Quaia”; K. Storey-Fisher et al. 2024), an all-sky, homogeneously observed sample of ~ 1.3 million quasars. Its combination of full-sky coverage, relatively uniform depth, and a tractable selection function provide an unprecedented data set for measuring the quasar clustering signal. We measure the projected two-point correlation function $w_p(r_p)$ in four redshift bins between $0 \leq z \leq 4$, and derive the real-space correlation length r_0 and slope γ . We then use a halo-bias framework (J. L. Tinker et al. 2008, 2010) to infer the corresponding minimum host dark matter halo mass M_{min} and derive the quasars’ duty cycle f_{duty} as a function of redshift.

This paper is structured as follows. In Section 2, we describe the Quaia catalog and the construction of our quasar sample, along with the selection functions and catalogs of randomly distributed quasars used to correct for observational biases. Section 3 presents our methodology for computing the projected two-point correlation function, and measuring the correlation length-scales. We also contrast our clustering results with previous studies in the literature. In Section 4, we infer the characteristic host dark-matter halo masses of Quaia quasars and estimate their duty cycles by combining our clustering measurements with quasar number densities. Finally, Section 5 summarizes our main findings and discusses their implications for quasar evolution and large-scale structure at high redshift. Throughout, we adopt a Planck Collaboration (2020) cosmology: $\Omega_m = 0.315$, $\Omega_\Lambda = 0.685$, $\Omega_b = 0.049$, $\Omega_c = 0.265$, $\Omega_r \simeq 9 \times 10^{-5}$, $\Omega_k \simeq 0$.

2. DATA

2.1. Quaia: Gaia’s quasar sample

We use the all-sky quasar catalog Quaia (K. Storey-Fisher et al. 2024), constructed from a parent sample of *Gaia* DR3 quasar candidates (Gaia Collaboration 2023) cross-matched with *unWISE* infrared photometry (D. Lang 2016; A. M. Meisner et al. 2019). The Quaia quasar sample is described in detail in K. Storey-Fisher




Figure 1. Two-dimensional histogram of the Quaia quasar sample as a function of redshift and absolute magnitude M_{1450} .

et al. (2024), but we will summarize the properties here briefly.

The parent *Gaia* DR3 quasar candidate sample contains 6,649,162 sources with low-resolution BP/RP spectra (330–1050 nm; $R \sim 30$ –100) and pipeline redshift estimates (“QSOC”), but was designed for completeness over purity (estimated purity $\sim 52\%$; *Gaia* Collaboration 2023). K. Storey-Fisher et al. (2024) decontaminated this *Gaia* parent set by applying a proper-motion cut and color cuts in *Gaia* matched with unWISE colors. This procedure reduced the known contaminants by $\sim 4\times$ while removing only $\sim 1.2\%$ of SDSS-confirmed quasars (R. Ahumada et al. 2020). The final Quaia sample used in this analysis is limited to magnitudes $G < 20.5$ and contains 1,295,502 quasars. We convert *Gaia*’s G -band magnitude to an absolute magnitude at rest-frame 1450 Å, M_{1450} , by simply applying a K-correction assuming a power-law continuum with spectral index $\alpha_\nu = -0.61$ (E. Lusso et al. 2015), and show the sample in Fig. 1.

Additionally, the redshift estimates are improved in Quaia by training a k -nearest-neighbors model on SDSS spectroscopic labels (R. Ahumada et al. 2020) using *Gaia*+unWISE colors and the “QSOC” redshift estimates from *Gaia* DR3 as features. For the $G < 20.5$ sample the catastrophic outlier fraction is reduced to 9% (16%) at $|\Delta z/(1+z)| > 0.2$ (0.1), which is a $\sim 3\times$ ($\sim 2\times$) improvement over the raw “QSOC” values. The typical estimated redshift precision for the quasar sample with $G < 20.5$ is $\Delta z/(1+z) = 0.008$ (K. Storey-Fisher et al. 2024).

Purity/completeness. Quaia is constructed to maximize homogeneity and controllable selection rather than to optimize a single global “purity” number. Empirically, the decontamination step reduces known contaminants by $\sim 4\times$ while retaining $> 98\%$ of SDSS quasars; the residual spatial completeness is modeled with an all-sky selection-function map built via Gaussian-process regression on dust, source-density, and survey-scan templates, with residuals consistent with homogeneous noise away from the Galactic plane (K. Storey-Fisher et al. 2024).

Volume and comparison to other surveys. Despite a lower number density than other optical+IR surveys, Quaia spans a larger effective comoving volume than SDSS DR16Q and the eBOSS quasar sample when matched in magnitude: for $G < 20.5$, $V_{\text{eff}} \approx 6.5 h^{-3} \text{Gpc}^3$, nearly $2\times$ DR16Q and $> 2\times$ eBOSS (B. W. Lyke et al. 2020; A. D. Myers et al. 2015); over the common span $0.8 \leq z \leq 2.2$, the “spanning volume” reaches $\sim 1.44 \times 10^2 h^{-3} \text{Gpc}^3$ (K. Storey-Fisher et al. 2024).

For our main analysis, we split the $G < 20.5$ quasars into four redshift bins:

$$\begin{aligned} 0.0 \leq z < 1.0 : N &= 327,195, \\ 1.0 \leq z < 2.0 : N &= 648,095, \\ 2.0 \leq z < 3.0 : N &= 283,958, \\ 3.0 \leq z < 4.0 : N &= 35,466. \end{aligned}$$

2.2. Selection functions and random catalogs

To estimate the clustering of quasars, we require a catalog of randomly distributed quasars that accurately reproduces both the angular and redshift selection functions of our data set. Measuring the quasars’ auto-correlation function therefore depends on a well-modeled selection function, which is used to create the catalogs of randomly distributed quasars assuming no clustering signal is present. Thus, these random catalogs provide the baseline against which we measure the excess clustering of the real quasar sample. To construct the selection functions for our quasar sample, we adopt the data-driven framework developed in K. Storey-Fisher et al. (2024), which learns how survey systematics modulate the observed quasar number density on the sky. The regression is performed on HEALPIX pixels (K. M. Górski et al. 2005) (NSIDE=64) using a compact set of physically motivated templates. These include the Corrected Schlegel, Finkbeiner, and Davis (CSFD) dust reddening map (Y.-K. Chiang 2023), which removes extragalactic cosmic infrared background (CIB) contamination from the classic SFD map (D. J. Schlegel et al. 1998); stellar density maps constructed from *Gaia* source counts with

$18.5 < G < 20$; unWISE source-density maps (D. Lang 2016); the *Gaia* scanning law and crowding proxy M_{10} , defined as the median G magnitude of sources with ≤ 10 transits from GaiaUnlimited (A. Everall et al. 2023), which encodes the scanning law and crowding effects; the unWISE scanning law, given by the mean number of single-exposure W1 images contributing to each coadded image (A. M. Meisner et al. 2019); and localized versions of the stellar and unWISE density maps restricted to wide regions around the Magellanic Clouds to capture their distinct contamination patterns.

Following K. Storey-Fisher et al. (2024), we regress the (log) quasar counts per pixel on the mean-subtracted templates using a Gaussian-process model with a squared-exponential kernel, which captures the nonlinear response of the number density to the systematics. The prediction is then converted into a relative completeness by normalizing to the mean predicted density in a set of “clean” pixels, which are defined as those with low extinction, low stellar and unWISE counts, high M_{10} , and high unWISE coadds. This normalization ensures values ≤ 1 and avoids sensitivity to a single extremal pixel.

We use selection-function maps and random catalogs at NSIDE=64 for the $G < 20.5$ sample. The random catalogs are generated by Poisson sampling a uniform sky and down-weighting by the selection function, and are provided at approximately $100\times$ the density of the data.

For redshift-split analyses, we adopt the corresponding selection functions and assign redshifts to objects in the random catalog by sampling from the empirical $p(z)$ of each bin. This procedure overall captures the Galactic extinction for quasars close to the plane very well. However, we note that for our analysis in the highest redshift bin, i.e. $3 \leq z \leq 4$, we apply a further cut based on Galactic latitude, excluding quasars close to the Galactic plane at $|b| < 30^\circ$, where modeling the dust distribution and its associated extinction in the optical bands remains uncertain, in order to mitigate contamination as we otherwise observe nonphysical clustering effects on large scales.

3. QUASAR CORRELATION FUNCTION

3.1. Auto-correlation function and scale length measurements

In this section, we present measurements of the projected two-point correlation function of quasars in distinct redshift intervals between $0 \leq z < 4$. We apply the Landy–Szalay estimator (S. D. Landy & A. S. Szalay 1993) as shown in Eqn. 2, which is subsequently integrated along the line of sight to yield the projected

correlation function. We then fit these projected correlation functions with a power-law, to study the clustering properties of the quasar population.

For each redshift bin we generate a corresponding selection function to construct a catalog of quasars randomly distributed on the sky. The Landy–Szalay estimator is defined as

$$\xi(r) = \frac{\langle DD \rangle - 2\langle DR \rangle + \langle RR \rangle}{\langle RR \rangle}, \quad (2)$$

where $\langle DD \rangle$, $\langle DR \rangle$, and $\langle RR \rangle$ denote the normalized counts of data–data, data–random, and random–random pairs, respectively. These pairs are counted as a function of distance, $r = \sqrt{r_p^2 + \pi^2}$, where r_p and π denote the separations perpendicular and parallel to the line of sight, respectively.

We compute the projected two-point correlation function, $w_p(r_p)$, by integrating the redshift-space correlation function, $\xi(r_p, \pi)$, along the line of sight:

$$w_p(r_p) = 2 \int_0^{\pi_{\max}} \xi(r_p, \pi) d\pi. \quad (3)$$

This projection suppresses the impact of redshift-space distortions, yielding an estimator that more directly reflects the underlying real-space clustering of quasars. In our analysis, we adopt $\pi_{\max} = 100 h^{-1}$ Mpc for the line-of-sight integration, a choice commonly used in quasar clustering studies (e.g., Y. Shen et al. 2007; M. White et al. 2012; S. Eftekharzadeh et al. 2015; J. Arita et al. 2023). We compute $w_p(r_p)$ in 14 logarithmically spaced transverse-separation bins spanning $r_p = 1\text{--}200 h^{-1}$ Mpc (corresponding to a bin width of $\Delta \log r_p \simeq 0.16$).

The statistical uncertainties on $w_p(r_p)$ are estimated using delete-one jackknife resampling within each redshift bin. We divide both the quasars and the random catalog into ten equal-width right ascension (RA) stripes covering the sky region of the bin. The projected correlation function is then recomputed ten times, each time omitting one RA stripe. Let $w_p^{(k)}(r_p)$ denote the measurement obtained when stripe k is excluded, and let $\bar{w}_p(r_p)$ be the mean over all jackknife realizations. The jackknife covariance matrix is computed using

$$C_{ij} = \frac{N-1}{N} \sum_{k=1}^N \left[w_p^{(k)}(r_{p,i}) - \bar{w}_p(r_{p,i}) \right] \left[w_p^{(k)}(r_{p,j}) - \bar{w}_p(r_{p,j}) \right], \quad (4)$$

where $N = 10$ is the number of jackknife regions. For visualization and for the tabulated uncertainties, we use




Figure 2. The projected auto-correlation function for the four redshift bins is shown with jackknifing errors for the four redshift bins, with errors estimated using jackknife resampling. The dashed lines and shaded areas show the best power-law fits with 2σ uncertainties. The legend shows the best-fit parameter values of the power-law model fit to the data.

only the diagonal elements of this matrix to define the 1σ error bars,

$$\sigma_{w_p}(r_{p,i}) = \sqrt{C_{ii}}.$$

Our measurements of $w_p(r_p)$, along with the corresponding transverse separations r_p and jackknife uncertainties σ_{w_p} , are shown in Fig. 2 and listed in Table 1 for each redshift bin.

To interpret the measurements of $w_p(r_p)$, we model the two-point correlation function as a power law of the form

$$\xi(r) = \left(\frac{r}{r_0}\right)^{-\gamma}, \quad (5)$$

where r_0 is the correlation length and γ is the slope. This yields the following expression for the projected correlation function, which can be derived by integrating along the line-of-sight:

$$w_p(r_p) = r_p \left(\frac{r_0}{r_p}\right)^\gamma \frac{\Gamma(\frac{1}{2}) \Gamma(\frac{\gamma-1}{2})}{\Gamma(\frac{\gamma}{2})}, \quad (6)$$

where Γ denotes the Gamma function (Y. Shen et al. 2007; M. Davis & P. J. E. Peebles 1983).

We fit this power-law model to the measured $w_p(r_p)$ in each redshift bin to infer the scale length r_0 and slope γ . Parameter estimation uses the ensemble Markov Chain Monte Carlo sampler **emcee** (D. Foreman-Mackey et al. 2013) to explore the posterior $p(r_0, \gamma | \mathbf{w}_p)$. We adopt




Figure 3. Comparison of the projected auto-correlation function of *Quia* sources with SDSS sources analyzed by Y. Shen et al. (2007) for $0.8 \leq z \leq 2.1$, $2.9 \leq z \leq 3.5$, and $z \geq 3.5$ (their *good* fields). Black measurements indicate results by Y. Shen et al. (2007), while red data points show the results from this study.

a Gaussian likelihood with the measured $w_p(r_p)$ values as the data vector and a jackknife-estimated, diagonal covariance (i.e., we use only the jackknife variances $\sigma_{w_p}^2$ for each r_p bin). We impose uniform priors matching of $r_0 \in [1.0, 50.0]$ and $\gamma \in [1.1, 3.0]$. We report the posterior medians as fiducial estimates and the 16th–84th percentiles as the corresponding 1σ credible intervals. The resulting values of r_0 and γ are reported in Table 2, and the best-fit models are shown in Figure 2.

3.2. Comparison to previous work

We compare our measurements to the quasar sample of comparable brightness analyzed by Y. Shen et al. (2007), who studied the large-scale clustering of $\sim 45,000$ optically selected SDSS DR5 quasars spanning $0.8 \leq z \leq 4.0$ and $M_i < -22$ (K-corrected to $z = 2$) over $\sim 4,000 \text{ deg}^2$ of sky. Their sample, drawn from the SDSS color-based quasar selection algorithm, probes bright unobscured quasars with redshifts and luminosities approximately matching those of our *Gaia*–unWISE sample. Thus, both datasets trace a similar population of luminous, optically selected quasars over comparable redshift ranges.

To ensure a consistent comparison, we adopt the same redshift intervals used in the previous study, i.e. $0.8 \leq z < 2.1$, $2.9 \leq z < 3.5$, and $z \geq 3.5$, for which we find $N = 814,858$, $N = 41,745$, and $N = 6,783$ quasars, respectively. Note that we do apply the same cut in Galactic latitude as before, i.e. $|b| > 30^\circ$, for the latter two redshift bins. For this part of the analysis, we recompute the *Quia* selection functions and random catalogs within these bins, ensuring that the projected correlation functions $w_p(r_p)$ are measured under with an identical redshift selection. This allows a direct comparison of the clustering amplitude and slope

between *Quia* and the SDSS spectroscopic results of Y. Shen et al. (2007). Figure 3 shows the projected correlation functions reported by Y. Shen et al. (2007) in black (for their *good* fields, which is a subset of *all* quasar fields analyzed in their work that satisfy a set of quality criteria on the photometry, see their § 2.2 for details) alongside our measurements shown in red. Overall, we find excellent agreement between the two analyses at lower redshift. For the $0.8 \leq z < 2.1$ bin, Y. Shen et al. (2007) reports $r_0 = 6.47 \pm 1.55 h^{-1} \text{ Mpc}$ and $\gamma = 1.58 \pm 0.20$, which are slightly lower than our $r_0 = 7.07 \pm 0.25 h^{-1} \text{ Mpc}$ and $\gamma = 1.89 \pm 0.03$. However, at $2.9 \leq z \leq 3.5$, Y. Shen et al. (2007) finds $r_0 = 17.91 \pm 1.51 h^{-1} \text{ Mpc}$ and $\gamma = 2.37 \pm 0.29$ while we measure slightly lower values of $r_0 = 13.05 \pm 1.9 h^{-1} \text{ Mpc}$ and $\gamma = 2.06 \pm 0.17$. At even higher redshift, at $z \geq 3.5$, this discrepancy increases to $r_0 = 25.22 \pm 2.5 h^{-1} \text{ Mpc}$ and $\gamma = 2.14 \pm 0.24$ reported from the SDSS quasars, while our analysis yields $r_0 = 19.01 \pm 0.6 h^{-1} \text{ Mpc}$ and $\gamma = 2.15 \pm 0.06$. Note, that our results are consistent with measurements by Y. Shen et al. (2007) when *all* fields are taken into account (instead of only their *good* fields) and when they fix the slope to $\gamma = 2.0$, in which case they obtain $r_0 = 14.79 \pm 2.12 h^{-1} \text{ Mpc}$ for $2.9 \leq z \leq 3.5$ and $r_0 = 20.68 \pm 2.52 h^{-1} \text{ Mpc}$ at $z \geq 3.5$.

4. HOST DARK MATTER HALO MASSES AND QUASAR DUTY CYCLE

4.1. Estimating the minimum host dark matter halo mass

In the framework of hierarchical structure formation, quasars are assumed to reside in dark matter haloes, and their clustering properties therefore encode information about the typical halo masses that host them. Within the ΛCDM cosmological model, the strength

Table 1. Projected correlation function $w_p(r_p)$ across redshift bins.

$r_p [h^{-1}\text{Mpc}]$	$w_p(r_p) \pm \sigma_{w_p} [h^{-1}\text{Mpc}]$			
	$0 \leq z < 1$	$1 \leq z < 2$	$2 \leq z < 3$	$3 \leq z < 4$
1.47	79.1 ± 12.8	79.2 ± 12.2	178 ± 13.5	677 ± 517
3.16	37.2 ± 4.39	52.8 ± 4.75	97.5 ± 18.3	326 ± 171
6.81	22.4 ± 2.95	31.9 ± 3.07	67.0 ± 4.81	45.3 ± 24.3
14.7	10.7 ± 2.24	16.3 ± 1.46	29.7 ± 2.28	72.7 ± 23.1
31.6	5.43 ± 0.76	7.07 ± 0.75	11.2 ± 0.76	10.3 ± 10.0
68.1	1.78 ± 0.32	2.18 ± 0.22	4.73 ± 0.54	4.83 ± 3.55
147	0.747 ± 0.168	0.544 ± 0.155	1.38 ± 0.36	3.98 ± 0.85
316	0.313 ± 0.086	0.156 ± 0.074	0.285 ± 0.240	2.49 ± 1.07
681	0.133 ± 0.026	0.121 ± 0.030	0.228 ± 0.119	1.05 ± 0.61

Table 2. Best-fit parameters r_0 and γ in four redshift bins.

Redshift range	γ	$r_0 [h^{-1}\text{Mpc}]$
$0.0 \leq z < 1.0$	1.93 ± 0.03	6.8 ± 0.2
$1.0 \leq z < 2.0$	1.99 ± 0.02	8.0 ± 0.2
$2.0 \leq z < 3.0$	1.95 ± 0.02	10.8 ± 0.2
$3.0 \leq z < 4.0$	2.25 ± 0.02	13.9 ± 1.2

of clustering on large scales (quantified by the correlation length r_0) reflects the typical separation scale over which quasars remain correlated. In other words, r_0 marks the distance at which the real-space two-point correlation function $\xi(r)$ falls to unity, meaning that a pair of quasars separated by r_0 is twice as likely to be found together as if they were randomly distributed. A larger r_0 therefore corresponds to stronger clustering and indicates that quasars inhabit rarer, more strongly biased environments. More strongly clustered quasars correspond to rarer, more massive haloes, while weaker clustering reflects residence in lower-mass haloes. Concretely, the observed r_0 can be mapped onto a minimum host-halo mass $M_{\min}$ by requiring consistency between the quasar number density and the cumulative abundance of dark matter haloes, $n_{\text{DM}}(> M_{\min})$.

We perform this conversion using the publicly available `HALOMOD TracerHaloModel` framework (e.g. [A. P. Hearin et al. 2017](#)), which self-consistently incorporates the halo mass function and halo bias relation to connect the measured clustering amplitude with a characteristic halo mass. In our configuration, `HALOMOD` adopts the halo mass function of [J. L. Tinker et al. \(2008\)](#) together with the large-scale halo bias model of [J. L. Tinker et al. \(2010\)](#), which enter through the two-halo term used to

predict the large-scale real-space correlation function. In practice, we compute the model correlation function $\xi(r|M_{\min})$ for a grid of $M_{\min}$ values, extract the corresponding correlation length $r_0(M_{\min})$ from the condition $\xi(r_0) = 1$, and invert this mapping to obtain the value of $M_{\min}$ that reproduces the observed r_0 in each redshift bin. Tracer occupation is set to unity above a sharp mass threshold (a step-function, or “constant-HOD”, model), leaving the minimum halo mass $M_{\min}$ as the single free parameter that links halo mass to the observed quasar clustering.

For each effective redshift $z_{\text{eff}} = \{0.5, 1.5, 2.5, 3.5\}$, we tabulate the function $r_0(M_{\min})$ on a grid of minimum halo masses spanning $9 \leq \log_{10}(M_{\min}/h^{-1}M_{\odot}) \leq 15$. For every trial $M_{\min}$ we (i) populate haloes above that threshold, (ii) compute the real-space halo auto-correlation function $\xi_h(r)$ over $0.1\text{--}200 h^{-1}\text{Mpc}$, and (iii) record the scale where $\xi_h = 1$. The resulting monotonic $r_0(M_{\min})$ relation is then inverted and evaluated at the observed median r_0 (as well as its 16th and 84th percentiles), yielding the corresponding $M_{\min}$ values with asymmetric 1σ uncertainties.

The results are shown in Figure 4, illustrating this mapping between the measured clustering amplitude and the inferred minimum host-halo mass. Although the observed correlation lengths r_0 vary substantially with redshift, the corresponding $M_{\min}$ values lie within a relatively narrow range of a few $\times 10^{12} h^{-1}M_{\odot}$. This visual trend reflects the redshift evolution of halo bias: at higher redshift, halos of a given mass are rarer and therefore more strongly clustered. As a result, a smaller r_0 at lower redshift can correspond to roughly the same halo mass as a larger r_0 at high- z . The near convergence of the curves highlights that luminous quasars occupy




Figure 4. Minimum dark matter halo mass vs r_0 comparison. Horizontal dashed lines represent measured values through auto-correlation fits; vertical lines are the predicted mass for each measurement.

halos of roughly constant characteristic mass across cosmic time.

The cumulative number density of halos more massive than the threshold is

$$n_{\text{DM}}(> M_{\text{min}}) = \int_{M_{\text{min}}}^{\infty} \frac{dn}{dM} dM,$$

where dn/dM is the mass function (J. L. Tinker et al. 2010) evaluated at z_{eff} . Integrating at the lower, median, and upper M_{min} values propagates the clustering uncertainty into asymmetric errors on n_{DM} .

4.2. Quasar Number Density and Duty Cycle

As quasars temporarily and randomly subsample their host dark matter haloes, we can estimate their duty cycle as established in Eqn. 1. The first expression emphasizes the interpretation of f_{duty} as the fraction of a Hubble time during which a halo is observed in a luminous quasar phase, while the second expression provides a practical estimator based on the fraction of halos more massive than M_{min} that are currently hosting a luminous quasar.

We estimate the quasars' number density n_{QSO} as follows. The HEALPIX $N_{\text{side}} = 64$ map encodes the relative completeness of every sky pixel. If c_p and Ω_p denote the weight and solid angle in steradians of pixel p , respectively, the survey's effective sky fraction is

$$f_{\text{sky,eff}} = \frac{1}{4\pi} \sum_p c_p \Omega_p. \quad (7)$$

For each redshift slice, the comoving shell volume V_c that would be enclosed on the full sky is computed via $V_{\text{shell}} = V_c(z_{\text{max}}) - V_c(z_{\text{min}})$. Multiplication by the factor in Eqn. (7) yields the survey's effective volume, i.e.

$$V_{\text{eff}} = f_{\text{sky,eff}} V_{\text{shell}}. \quad (8)$$

Table 3. Minimum host-halo mass M_{min} and quasar duty cycle f_{duty} in each redshift bin.

Redshift range	$\log_{10}(M_{\text{min}}/M_{\odot})$	f_{duty}
$0.0 \leq z < 1.0$	$12.86^{+0.04}_{-0.05}$	$0.018^{+0.002}_{-0.002}$
$1.0 \leq z < 2.0$	$12.74^{+0.02}_{-0.02}$	$0.022^{+0.001}_{-0.001}$
$2.0 \leq z < 3.0$	$12.79^{+0.03}_{-0.03}$	$0.056^{+0.008}_{-0.007}$
$3.0 \leq z < 4.0$	$12.82^{+0.10}_{-0.12}$	$0.071^{+0.077}_{-0.034}$

Then, given the number of quasars in the given redshift shell as N_{Q} , the mean quasar space density follows as

$$n_{\text{QSO}} = \frac{N_{\text{Q}}}{V_{\text{eff}}}. \quad (9)$$

We obtain duty-cycle measurements ranging from $f_{\text{duty}} \simeq 2\%$ to 7% , increasing with redshift. These values correspond to an integrated quasar lifetime of $t_{\text{QSO}} \sim 10^8$ years across all redshift bins. Note that these uncertainties represent only the statistical component obtained by forward-propagating the errors on r_0 and M_{min} into the halo number densities and Poisson error for N_{QSO} . In practice, the duty cycle is subject to larger systematic uncertainties arising from assumptions about the halo-occupation distribution, the scatter between halo mass and quasar luminosity, and the choice of HOD parametrization (e.g. F. Shankar et al. 2010; F. Shankar 2013; E. Pizzati et al. 2024).

Our measurements for each redshift bin are listed in Table 3 and shown in Figure 5.

5. SUMMARY AND DISCUSSION

Having presented our measurements of the projected correlation functions, dark matter halo masses, and duty cycles, we now turn to their interpretation. In the following, we discuss our results and compare them with previous studies. In this work, we have measured the projected correlation function $w_p(r_p)$ for ~ 1.3 million quasars from the *Gaia*-QUAIA catalog across four redshift bins spanning $0 \leq z \leq 4$. Using these measurements, we inferred the characteristic host-halo masses and corresponding duty cycles, enabling a direct comparison with earlier clustering analyses.

Figure 5 summarizes our results and shows the redshift evolution of the quasars' clustering lengthscale and duty cycle.

Clustering amplitude. We observe that the correlation length r_0 increases steadily with redshift, from $r_0 \approx 6.8 h^{-1} \text{Mpc}$ at $0 \leq z < 1$ to $r_0 \approx 13.9 h^{-1} \text{Mpc}$ at $3 \leq z < 4$. At $z > 3.5$ we obtain $r_0 \approx 19 h^{-1} \text{Mpc}$, which is lower than previous work at similar redshifts (Y. Shen et al. 2007, 2010), as shown in Fig. 3, and thus




Figure 5. Redshift evolution of the quasars’ auto-correlation scale length (left) and the quasars’ duty cycle (right). Our measurements are shown in red, while results from previous studies are shown in different colors (Y. Shen et al. 2007; N. P. Ross et al. 2009; M. White et al. 2012; S. Eftekharzadeh et al. 2015; P. Laurent et al. 2017; J. Arita et al. 2023; A.-C. Eilers et al. 2024; E. Pizzati et al. 2024; J.-T. Schindler et al. 2025). Dashed lines in the left panel show the clustering strength for a constant dark matter halo mass, making the same assumption as in our study, i.e. a step-function HOD with the J. L. Tinker et al. (2010) relation modeled with `halomod`. The dashed lines in the right panel indicate the duty cycle corresponding to a constant integrated quasar lifetime t_{QSO} .

suggests less biased host halos. It is unclear where the discrepancy arises given that the quasars’ redshift range and luminosities are similar. Differences might be due to challenges in modeling the selection function, which is more tractable with Quiaia compared to SDSS. This is supported by the fact that our results and those of Y. Shen et al. (2007) are comparable when *all* quasar fields are taken into account in their analysis instead of only their *good* fields, suggesting that selection effects might play a role. On the other hand the Quiaia sample has less accurate redshift estimates due to the low spectral resolution of the BP/RP spectra and likely contains some contaminants, as discussed in § 2, which could also lead to biases. At lower redshifts, our measurements are slightly higher than those of S. Eftekharzadeh et al. (2015) and M. White et al. (2012), but consistent with N. P. Ross et al. (2009).

The smooth, monotonic rise we observe with redshift appears to continue up to $z \approx 6.25$, spanning over 13 Gyr of cosmic history. We find no evidence for any

rapid change in clustering strength, despite the current lack of measurements at $4 \lesssim z \lesssim 6$. At $z > 7$, the measurement by J.-T. Schindler et al. (2025) suggests a sudden drop in scale length and thus a potentially non-monotonic evolution. However, this study was reporting the cross-correlation of only two faint quasars with their surrounding galaxies and thus the results are likely highly affected by cosmic variance.

Host halo masses. When translated to halo bias using the J. L. Tinker et al. (2010) relation, our measurements imply a nearly constant characteristic halo mass of $\log_{10} M_{\text{min}} \simeq 12.7\text{--}12.9 h^{-1} M_{\odot}$ across $0 \leq z \leq 4$. These values fall squarely within the ballpark established by previous studies, and show remarkably little evolution with redshift. This near constancy confirms results from previous quasar clustering analyses (e.g., S. M. Croom et al. 2005; Y. Shen et al. 2007, 2010; J. Arita et al. 2023; A.-C. Eilers et al. 2024), and supports the idea that quasar activity preferentially occurs in halos of similar mass, independent of epoch. In the

halo-occupation framework, this suggests the existence of a characteristic halo mass, around where quasar triggering appears most efficient.

Note, however, that translating $r_0 \rightarrow M_{\min}$ is model dependent. For instance, despite the higher clustering amplitude measured by Y. Shen et al. (2007), their inferred halo masses are comparable to ours, reflecting their use of a different HOD, i.e. the R. K. Sheth et al. (2001) bias model and A. Jenkins et al. (2001) mass function. These studies predict higher bias values at fixed halo mass than the J. L. Tinker et al. (2010) calibration adopted here, leading to systematically lower $M_{\min}$ estimates for a given r_0 . Similarly, at $z \approx 6.25$, A.-C. Eilers et al. (2024) used a step-function HOD to model both the host halos of quasars as well as those of the surrounding galaxies in the quasars’ environments, while E. Pizzati et al. (2024) modeled the same data taking into account the quasars’ luminosities and their luminosity-dependent halo mass relation. However, both approaches find host halo masses that agree within < 0.1 dex.

Duty cycle and lifetime. We find quasar duty cycle measurements of $f_{\text{duty}} \sim 2\text{--}7\%$ increasing mildly with redshift (see right panel of Fig. 5), which suggests

an integrated quasar lifetime of $t_{\text{QSO}} \sim 10^8$ yr across cosmic time. This alignment is non-trivial: it means that the *integrated* luminous time per halo required to match the observed quasar abundance at fixed $M_{\min}$ is roughly constant from $z \sim 0$ to $z \sim 4$. Physically, a lifetime of $t_{\text{QSO}} \sim 10^8$ yr is comparable to a few “Salpeter” timescales, or e -folding timescales of the black hole for radiatively efficient, Eddington-limited growth, as expected in our standard picture of black hole growth (e.g., Q. Yu & S. Tremaine 2002). Read together, these points favor a self-regulated black hole growth picture in which luminous accretion can proceed in one long or multiple shorter accretion episodes (A.-C. Eilers et al. 2017; I. S. Khrykin et al. 2019; K. A. Morey et al. 2021) whose total integrated duration is approximately $\sim 10^8$ yr per halo across cosmic time.

ACKNOWLEDGMENTS

The authors would like to thank Yue Shen, David W. Hogg, Joe Hennawi, and Mahlet Shiferaw for helpful discussions.

M.G.M. acknowledges the support of the MIT Summer Research Program (MSRP).

REFERENCES

- Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, The Astrophysical Journal Supplement Series, 249, 3, doi: [10.3847/1538-4365/ab929e](https://doi.org/10.3847/1538-4365/ab929e)
- Arita, J., Kashikawa, N., Matsuoka, Y., et al. 2023, ApJ, 954, 210, doi: [10.3847/1538-4357/ace43a](https://doi.org/10.3847/1538-4357/ace43a)
- Chiang, Y.-K. 2023, ApJ, 958, 118, doi: [10.3847/1538-4357/acf4a1](https://doi.org/10.3847/1538-4357/acf4a1)
- Conroy, C., & White, M. 2013, ApJ, 762, 70, doi: [10.1088/0004-637X/762/2/70](https://doi.org/10.1088/0004-637X/762/2/70)
- Croom, S. M., Boyle, B. J., Shanks, T., et al. 2005, MNRAS, 356, 415, doi: [10.1111/j.1365-2966.2004.08499.x](https://doi.org/10.1111/j.1365-2966.2004.08499.x)
- Davis, M., & Peebles, P. J. E. 1983, ApJ, 267, 465, doi: [10.1086/160884](https://doi.org/10.1086/160884)
- Eftekharzadeh, S., et al. 2015, MNRAS, 453, 2779, doi: [10.1093/mnras/stv1800](https://doi.org/10.1093/mnras/stv1800)
- Eilers, A.-C., Davies, F. B., Hennawi, J. F., et al. 2017, ApJ, 840, 24, doi: [10.3847/1538-4357/aa6c60](https://doi.org/10.3847/1538-4357/aa6c60)
- Eilers, A.-C., Mackenzie, R., Pizzati, E., et al. 2024, ApJ, 974, 275, doi: [10.3847/1538-4357/ad778b](https://doi.org/10.3847/1538-4357/ad778b)
- Everall, A., et al. 2023, GaiaUnlimited: Mapping Gaia’s scanning and completeness using open-source tools,, <https://gaiaunlimited.space/>
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306, doi: [10.1086/670067](https://doi.org/10.1086/670067)
- Gaia Collaboration. 2023, A&A, 674, A29, doi: [10.1051/0004-6361/202243161](https://doi.org/10.1051/0004-6361/202243161)
- Górski, K. M., et al. 2005, ApJ, 622, 759, doi: [10.1086/427976](https://doi.org/10.1086/427976)
- Hearin, A. P., Campbell, D., Tollerud, E., et al. 2017, AJ, 154, 190, doi: [10.3847/1538-3881/aa859f](https://doi.org/10.3847/1538-3881/aa859f)
- Hopkins, P. F., Hernquist, L., Cox, T. J., et al. 2006, ApJS, 163, 1, doi: [10.1086/499298](https://doi.org/10.1086/499298)
- Jenkins, A., Frenk, C. S., White, S. D. M., et al. 2001, Monthly Notices of the Royal Astronomical Society, 321, 372, doi: [10.1046/j.1365-8711.2001.04029.x](https://doi.org/10.1046/j.1365-8711.2001.04029.x)
- Khrykin, I. S., Hennawi, J. F., & Worseck, G. 2019, MNRAS, 484, 3897, doi: [10.1093/mnras/stz135](https://doi.org/10.1093/mnras/stz135)
- Kormendy, J., & Ho, L. C. 2013, ARA&A, 51, 511, doi: [10.1146/annurev-astro-082708-101811](https://doi.org/10.1146/annurev-astro-082708-101811)
- Landy, S. D., & Szalay, A. S. 1993, The Astrophysical Journal, 412, 64, doi: [10.1086/172900](https://doi.org/10.1086/172900)
- Lang, D. 2016, AJ, 151, 36, doi: [10.3847/0004-6256/151/2/36](https://doi.org/10.3847/0004-6256/151/2/36)
- Laurent, P., Eftekharzadeh, S., Le Goff, J.-M., et al. 2017, JCAP, 07, 017, doi: [10.1088/1475-7516/2017/07/017](https://doi.org/10.1088/1475-7516/2017/07/017)

- Lusso, E., Worseck, G., Hennawi, J. F., et al. 2015, MNRAS, 449, 4204, doi: [10.1093/mnras/stv516](https://doi.org/10.1093/mnras/stv516)
- Lyke, B. W., et al. 2020, ApJS, 250, 8, doi: [10.3847/1538-4365/aba623](https://doi.org/10.3847/1538-4365/aba623)
- Madau, P., & Dickinson, M. 2014, ARA&A, 52, 415, doi: [10.1146/annurev-astro-081811-125615](https://doi.org/10.1146/annurev-astro-081811-125615)
- Martini, P., & Weinberg, D. H. 2001, ApJ, 547, 12, doi: [10.1086/318333](https://doi.org/10.1086/318333)
- McConnell, N. J., & Ma, C.-P. 2013, ApJ, 764, 184, doi: [10.1088/0004-637X/764/2/184](https://doi.org/10.1088/0004-637X/764/2/184)
- Meisner, A. M., Lang, D., Schlafly, E. F., Green, G., & Finkbeiner, D. P. 2019, PASP, 131, 124504, doi: [10.1088/1538-3873/ab5364](https://doi.org/10.1088/1538-3873/ab5364)
- Morey, K. A., Eilers, A.-C., Davies, F. B., Hennawi, J. F., & Simcoe, R. A. 2021, ApJ, 921, 88, doi: [10.3847/1538-4357/ac1c70](https://doi.org/10.3847/1538-4357/ac1c70)
- Myers, A. D., et al. 2015, ApJS, 221, 27, doi: [10.1088/0067-0049/221/2/27](https://doi.org/10.1088/0067-0049/221/2/27)
- Pizzati, E., Hennawi, J. F., Schaye, J., et al. 2024, Monthly Notices of the Royal Astronomical Society, 534, 3155, doi: [10.1093/mnras/stae2307](https://doi.org/10.1093/mnras/stae2307)
- Planck Collaboration. 2020, A&A, 641, A6, doi: [10.1051/0004-6361/201833910](https://doi.org/10.1051/0004-6361/201833910)
- Ross, N. P., Shen, Y., Strauss, M. A., et al. 2009, ApJ, 697, 1634, doi: [10.1088/0004-637X/697/2/1634](https://doi.org/10.1088/0004-637X/697/2/1634)
- Schindler, J.-T., Hennawi, J. F., Davies, F. B., et al. 2025, arXiv e-prints, arXiv:2510.08455, doi: [10.48550/arXiv.2510.08455](https://doi.org/10.48550/arXiv.2510.08455)
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525, doi: [10.1086/305772](https://doi.org/10.1086/305772)
- Shankar, F. 2013, Monthly Notices of the Royal Astronomical Society, 428, 421, doi: [10.1093/mnras/sts032](https://doi.org/10.1093/mnras/sts032)
- Shankar, F., Weinberg, D. H., & Shen, Y. 2010, Monthly Notices of the Royal Astronomical Society, 406, 1959, doi: [10.1111/j.1365-2966.2010.16804.x](https://doi.org/10.1111/j.1365-2966.2010.16804.x)
- Shen, Y., et al. 2007, AJ, 133, 2222, doi: [10.1086/513517](https://doi.org/10.1086/513517)
- Shen, Y., Hennawi, J. F., Shankar, F., et al. 2010, The Astrophysical Journal, 719, 1693, doi: [10.1088/0004-637X/719/2/1693](https://doi.org/10.1088/0004-637X/719/2/1693)
- Sheth, R. K., Mo, H. J., & Tormen, G. 2001, Monthly Notices of the Royal Astronomical Society, 323, 1, doi: [10.1046/j.1365-8711.2001.04006.x](https://doi.org/10.1046/j.1365-8711.2001.04006.x)
- Sołtan, A. 1982, MNRAS, 200, 115, doi: [10.1093/mnras/200.1.115](https://doi.org/10.1093/mnras/200.1.115)
- Storey-Fisher, K., Hogg, D. W., Rix, H.-W., et al. 2024, ApJ, 964, 69, doi: [10.3847/1538-4357/ad1328](https://doi.org/10.3847/1538-4357/ad1328)
- Tinker, J. L., et al. 2008, ApJ, 688, 709, doi: [10.1086/591439](https://doi.org/10.1086/591439)
- Tinker, J. L., et al. 2010, ApJ, 724, 878, doi: [10.1088/0004-637X/724/2/878](https://doi.org/10.1088/0004-637X/724/2/878)
- White, M., Martini, P., & Cohn, J. D. 2008, MNRAS, 390, 1179, doi: [10.1111/j.1365-2966.2008.13834.x](https://doi.org/10.1111/j.1365-2966.2008.13834.x)
- White, M., et al. 2012, MNRAS, 424, 933, doi: [10.1111/j.1365-2966.2012.21251.x](https://doi.org/10.1111/j.1365-2966.2012.21251.x)
- Yu, Q., & Tremaine, S. 2002, MNRAS, 335, 965, doi: [10.1046/j.1365-8711.2002.05687.x](https://doi.org/10.1046/j.1365-8711.2002.05687.x)