



Universiteit
Leiden
The Netherlands

From a biased perspective: quasars, mergers, and planet-forming discs

Pizzati, E.

Citation

Pizzati, E. (2025, October 29). *From a biased perspective: quasars, mergers, and planet-forming discs*. Retrieved from <https://hdl.handle.net/1887/4281578>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4281578>

Note: To cite this publication please use the final published version (if applicable).

2 | REVISITING THE EXTREME CLUSTERING OF $z \approx 4$ QUASARS WITH LARGE VOLUME COSMOLOGICAL SIMULATIONS

Abstract

Observations from wide-field quasar surveys indicate that the quasar auto-correlation length increases dramatically from $z \approx 2.5$ to $z \approx 4$. This large clustering amplitude at $z \approx 4$ has proven hard to interpret theoretically, as it implies that quasars are hosted by the most massive dark matter halos residing in the most extreme environments at that redshift. In this work, we present a model that simultaneously reproduces both the observed quasar auto-correlation and quasar luminosity functions. The spatial distribution of halos and their relative abundance are obtained via a novel method that computes the halo mass and halo cross-correlation functions by combining multiple large-volume dark-matter-only cosmological simulations with different box sizes and resolutions. Armed with these halo properties, our model exploits the conditional luminosity function framework to describe the stochastic relationship between quasar luminosity, L , and halo mass, M . Assuming a simple power-law relation $L \propto M^\gamma$ with log-normal scatter, σ , we are able to reproduce observations at $z \sim 4$ and find that: (a) the quasar luminosity-halo mass relation is highly non-linear ($\gamma \gtrsim 2$), with very little scatter ($\sigma \lesssim 0.3$ dex); (b) luminous quasars ($\log_{10} L/\text{erg s}^{-1} \gtrsim 46.5 - 47$) are hosted by halos with mass $\log_{10} M/M_\odot \gtrsim 13 - 13.5$; and (c) the implied duty cycle for quasar activity approaches unity ($\varepsilon_{\text{DC}} \approx 10 - 60\%$). We also consider observations at $z \approx 2.5$ and find that the quasar luminosity-halo mass relation evolves significantly with cosmic time, implying a rapid change in quasar host halo masses and duty cycles, which in turn suggests concurrent evolution in black hole scaling relations and/or accretion efficiency.

Published in: **EP**, Joseph F Hennawi, Joop Schaye, Matthieu Schaller, *Revisiting the extreme clustering of $z \approx 4$ quasars with large volume cosmological simulations*, Monthly Notices of the Royal Astronomical Society, Volume 528, Issue 3, March 2024, Pages 4466–4489, doi.org/10.1093/mnras/stae329
Reprinted here in its entirety.

2.1 Introduction

Quasars are extreme manifestations of the supermassive black holes (SMBHs) that are thought to reside at the center of almost every massive galaxy (e.g., Salpeter 1964; Zel'dovich & Novikov 1967; Lynden-Bell 1969; Magorrian et al. 1998; Ferrarese & Merritt 2000; Kormendy & Ho 2013). Investigating the characteristics of these luminous objects has been an active area of research for more than half a century (Schmidt 1963). In the last few years, it has become possible to trace their evolution up to redshift $z \approx 7$ (Yang et al. 2020; Bañados et al. 2018; Wang et al. 2021; see also Fan et al. 2023 for a review). Understanding the properties of quasars such as their abundance, luminosity, and spatial distribution, as well as their evolution with redshift, is a key step in order to study the interplay between supermassive black holes, their host galaxies, and the intergalactic medium (IGM) over cosmic time.

In particular, measuring the clustering of quasars is crucial for gaining information on the large-scale environment in which these objects reside. Like their host halos, quasars are biased tracers of the underlying distribution of dark matter (e.g., Kaiser 1984; Bardeen et al. 1986). For this reason, obtaining an estimate for the linear bias factor of quasars (e.g., by measuring the quasar auto-correlation function) makes it possible to infer the characteristic masses of the halos hosting active quasars. In turn, these masses can shed light on the large-scale environment that quasars inhabit, and – by comparing the number density of quasars with that of the hosting halos – on the fraction of time SMBHs are shining as active quasars (known as the quasar duty cycle; see e.g. Martini & Weinberg 2001; Haiman & Hui 2001; Martini 2004).

Thanks to large-sky surveys such as the Sloan Digital Sky Survey (SDSS, York et al. 2000) and the 2dF QSO redshift survey (2QZ, Croom et al. 2004), measurements of large-scale quasar clustering up to $z \approx 4$ have been available for more than a decade. However, a satisfactory theoretical interpretation of these data at all redshifts is still lacking. This is mainly due to the surprising evidence that the bias factor of quasars is a steep function of redshift (Porciani et al. 2004; Croom et al. 2005; Porciani & Norberg 2006; Shen et al. 2007; Ross et al. 2009; Eftekharzadeh et al. 2015; McGreer et al. 2016; Yue et al. 2021; Arita et al. 2023). While in the local universe quasars trace halos in a way that is similar to optically selected galaxies, with a bias factor close to unity (Croom et al. 2005; Ross et al. 2009), at $z \approx 4$ they are the most highly clustered objects known at that epoch, with a bias factor as high as $b \approx 15$ (or, equivalently, a quasar auto-correlation length of $r_0 \approx 24 \text{ cMpc } h^{-1}$; Shen et al. 2007, hereafter S07). Such a large correlation length implies that quasars are rare objects, arising only in the most massive halos and shining for a large fraction of the Hubble time.

Several theoretical studies have tried to reproduce the results of S07 at $z \approx 4$. White et al. (2008) developed a simple model for quasar demographics that builds on a linear relation between quasar luminosity and host halo mass. They showed that to match the bias measured in S07, the scatter in this relation must be very small ($\lesssim 0.3$ dex). This conclusion poses two fundamental problems. Firstly, such a low scatter in the quasar luminosity-halo mass relation would be very surprising. In fact, the conventional wisdom on the coevolution of quasars and host galaxies/halos implies that there are multiple sources of scatter contributing to determining the luminosity of a quasar at fixed halo mass (the scatter in the relations between black hole mass and quasar luminosity, black hole mass and bulge mass and between bulge mass and halo mass). A second concern is that low scatter in the luminosity of quasars seems to be in contrast with measurements of the relative abundance of quasars at different luminosities (the so-called quasar luminosity function, QLF). It has long been established that the bright end of the QLF is well-fitted by a power-law (e.g., Boyle et al. 2000; Richards et al. 2006), which stands in contrast with the exponentially-declining halo mass and galaxy luminosity functions (Press & Schechter 1974; Schechter 1976). The easiest way to connect these different shapes is via significant scatter in the luminosity of quasars at fixed halo/galaxy mass. Indeed, a number of demographic models have been developed to interpret the abundance of bright quasars and link them to their host halos (e.g., Croton 2009; Conroy & White 2013; Fanidakis et al. 2013; Veale et al. 2014; Ren et al. 2020; Ren & Trenti 2021; Zhang et al. 2023b). All of these studies (sometimes only implicitly) explain the relatively large number of very luminous quasars by demanding a broad range of possible quasar luminosities at a given host mass so that the more abundant population of lower-mass halos can also host a significant (or even dominant; e.g., Zhang et al. 2023a) fraction of the very bright quasars. As pointed out by some of these same studies, however, the masses of the quasar hosts implied by this picture are in plain contrast with the high masses necessary to account for the S07 bias measurement.

In summary, the very strong clustering measured by S07 implies a very small scatter in the luminosity of quasars at a given halo mass, and this is in tension with the large scatter required by physical models of the quasar luminosity function. A first attempt at solving the tension was made by Shankar et al. (2010b), using a model that connects quasar luminosities and black hole masses while accounting for the growth of black holes during cosmic time in a self-consistent way. The authors of this study tried to match simultaneously the value of the bias inferred by S07 and several measurements of the QLF at $z = 3 - 6$ (Shankar & Mathur 2007; Shankar 2009). Assuming a non-linear relation between halo mass and quasar luminosity, they find that a low value of the scatter in this relation can reproduce the measurements of the bright end of the QLF. Even when assuming that all massive halos contribute to the clustering of quasars (i.e., a quasar duty cycle for massive

systems equal to unity), however, their prediction for the $z = 4$ quasar clustering is ≈ 2 standard deviations below the value measured by S07. Wyithe & Loeb (2009) also find that the S07 bias measurement cannot be reproduced when assuming that the bias of dark matter halos is solely a function of their mass, and suggest that stronger clustering could be obtained if quasar activity was sparked by recent mergers (the so-called “assembly/merger bias”, see e.g., Furlanetto & Kamionkowski 2006; Wetzel et al. 2009; see also Wechsler & Tinker 2018). However, Bonoli et al. (2010) (see also Cen & Safarzadeh 2015) used the Millennium Simulation (Springel et al. 2005) to study whether recently merged massive halos were clustered more strongly than other halos of the same mass, but found no evidence for that.

Numerous other studies have compared their predictions for the quasar clustering to the S07 measurements, using a variety of different approaches such as empirical models of quasar-galaxy coevolution (Kauffmann & Haehnelt 2002; Hopkins et al. 2007b; Croton 2009; Shankar et al. 2010a; Conroy & White 2013; Aversa et al. 2015; Shankar et al. 2020), semi-analytic models of galaxy formation (Bonoli et al. 2009; Fanidakis et al. 2013; Oogi et al. 2016) and cosmological hydrodynamical simulations (DeGraf et al. 2012; DeGraf & Sijacki 2017). While these studies are generally successful in reproducing the quasar auto-correlation function (or, equivalently, the quasar linear bias) at lower redshift ($z \lesssim 3$), none of these studies have been shown to be compatible with the strong clustering observed by S07.

In conclusion, despite the efforts that have been devoted to interpreting the auto-correlation function of quasars at high redshift, a number of questions remain open: (a) is the S07 measurement compatible with the standard cosmological model in which clustering is dictated by halo mass or is something akin to assembly bias playing an important role? (b) What is the scatter in the quasar luminosity-halo mass relation? Can small (large) scatter be reconciled with the observed quasar luminosity function (auto-correlation function)? (c) What are the physical properties that can be inferred from jointly modeling the QLF and quasar clustering? Can the characteristic mass of host halos and the quasar duty cycle be determined precisely?

One of the reasons why we have not been able to give definitive answers to these questions in more than a decade, is that modeling the clustering of high redshift quasars is difficult. The works of White et al. (2008), Shankar et al. (2010b), and Wyithe & Loeb (2009) clearly show that the results of their theoretical models are strongly dependent on the assumed functional form for the linear bias-halo mass relation. This is because the different analytical predictions for this relation based on linear theory (e.g. Mo & White 1996; Jing 1998; Sheth et al. 2001) diverge significantly at masses that correspond to peaks in the density perturbations that are already very non-linear (Barkana & Loeb 2001). For the case considered here, a bias of $b \approx 15$,

i.e., the value measured by S07 for $z \approx 4$ quasars, corresponds to a value of the peak height $\nu = \delta_c / \sigma(M, z)$ – with $\delta_c \approx 1.69$ and $\sigma^2(M, z)$ being the variance of the smoothed linear density field – equal to $\nu \approx 4 - 6$, depending on the specific linear bias-halo mass relation and cosmology considered. Such values are rather extreme, implying that the systems contributing to the clustering of $z = 4$ quasars live in very rare and massive environments that depart very early from the behavior expected for a linear density field.

Improving the accuracy of the linear bias-halo mass relation via empirical fits to cosmological N-body simulations (e.g., Shankar et al. 2010b; Tinker et al. 2010; Comparat et al. 2017) does not provide a complete solution to the problem. In fact, the key point here is that the use of the large-scale linear bias as a proxy for the clustering of quasars assumes that the measured data are on quasi-linear scales, where the distribution of quasars is related to the underlying matter distribution by a scale-independent factor. This assumption breaks down for the small scales (as low as ≈ 5 cMpc) and for the highly non-linear environments probed by the S07 data. For the same reason, an approach based on the (analytical) halo model framework (Cooray & Sheth 2002) would also be problematic, as the non-linear bias plays a relevant role in the transition region between the one-halo and the two-halo contributions (e.g., Mead & Verde 2021; Nishimichi et al. 2021).

In this paper, we aim to directly reproduce the observed $z = 4$ quasar auto-correlation function (S07) in its entirety by making use of large-volume cosmological simulations. This is a challenging numerical problem: in order to model the auto-correlation function properly, we need to obtain a large statistical sample of halos with masses up to $M \approx 10^{13} - 10^{14} M_\odot$ (which correspond, at $z = 4$, to the peak heights mentioned above, i.e., $\nu \approx 4 - 6$). Given the fact that the mass function declines exponentially at large masses, these halos are extremely rare ($1 - 10 \text{ cGpc}^{-3}$), and therefore a very large simulated volume of more than $\approx 100 \text{ cGpc}^3$ is needed to obtain a sample of at least $\approx 10^2 - 10^3$ massive halos, that can be used to properly model the quasar auto-correlation function even at the highest masses. This is in agreement with the fact that the comoving volume probed by the SDSS observations used by S07 is around $\approx 50 \text{ cGpc}^3$. A volume larger than the observational one is necessary to build a model for the quasar auto-correlation function that has higher signal-to-noise ratio than the data. At the same time, however, we also want to resolve halos down to $M \approx 10^{11} - 10^{12} M_\odot$ in order to explore the different possible distributions of quasars in halos that can give rise to the observed clustering. To probe these very different halo masses, we make use of a new semi-analytical framework (Sec. 2.2.2.2 and Appendix 2.B) that allows us to employ multiple simulated boxes to widen the range of masses that can be properly modeled by our simulations.

We employ the dark-matter-only versions of the FLAMINGO suite of cosmological simulations (Schaye et al. 2023; Kugel et al. 2023) and focus on two specific box sizes: $L = 2.8 \text{ cGpc}$ and $L = 5.6 \text{ cGpc}$. On top of reproduc-

ing the clustering measurements at $z = 4$, we also consider the constraints coming from the relative abundance of quasars at the same redshift. In other words, we aim to match the observed quasar auto-correlation and luminosity functions simultaneously. We make use of the spatial and mass distribution of halos in the simulated volumes to build a simple quasar population model that can be directly compared with observations. In this way, we are able to investigate the predictive power of quasar observables in a Λ CDM framework and obtain physical constraints on the halo mass distribution of quasar hosts and the quasar duty cycle. We also use our model to analyze the clustering and luminosity function data at a lower redshift ($z \approx 2 - 3$), where the tension between theoretical models and data is not as strong (e.g., [Croton 2009](#); [Conroy & White 2013](#); [Aversa et al. 2015](#)). This serves as a benchmark of the validity of our model and allows us to discuss the evolution of the physical properties of quasars with redshift.

The paper is structured as follows. In Sec. 2.2 we discuss the basic assumptions of the model, outline the cosmological simulations employed in our work, and describe how we extract the physical quantities that are necessary to model the quasar correlation function and luminosity function simultaneously. Sec. 2.3 gives a brief overview of the data we compare our model with, and it provides details on the statistical methodology underlying that comparison. Sec. 2.4 presents the main results of our analysis, while Sec. 2.5 contains a discussion of the implications of our findings and their connections to previous work. Conclusions are provided in Sec. 2.6.

2.2 Methods

In this Section, we describe our model for the distribution of quasars in space and luminosity. We start by outlining the basic framework (Sec. 2.2.1); then, we describe the FLAMINGO cosmological simulations and detail how we extract the mass function and the cross-correlation functions of halos (Sec. 2.2.2). Figure 2.1 shows a summary of the various quantities involved in our analysis, together with a reference to the equations where they are defined.

2.2.1 The conditional luminosity function

We adopt an empirical model that is agnostic to the physics underlying the quasar emission/black hole accretion mechanisms. The only assumptions we make are: (a) every halo above some mass $M_{\min}$ hosts a SMBH at its center, emitting at some luminosity L ; (b) the luminosity of a SMBH depends only on the mass of the host halo, M . Therefore, we can employ a conditional luminosity function approach (CLF; see e.g., [Yang et al. 2003](#); [Ren et al.](#)




Figure 2.1: Summary of the various quantities involved in the analysis. We choose a model for the Conditional Luminosity Function (CLF) that depends on a set of free parameters. We then combine this with the halo mass function and the halo cross-correlation functions taken from the FLAMINGO cosmological simulations to obtain the two main observables of interest, the quasar luminosity function and auto-correlation function, together with other key properties such as the quasar-host mass function, the halo occupation distribution (HOD), and the quasar duty cycle.

2020) and write the 2-d distribution in the black hole luminosity-host halo mass plane, $n(L, M)$, as:

$$n(L, M) = \text{CLF}(L|M) n_{\text{HMF}}(M), \quad (2.1)$$

where $n_{\text{HMF}}(M)$ is the halo mass function.

Note that the luminosity of a SMBH, L , can be interpreted as either a bolometric luminosity or a luminosity in a specific band of the spectrum. The framework that we are introducing here is agnostic to this choice and can be formulated to describe the emission coming from any region of the spectrum. However, for clarity and consistency with previous work on the subject (e.g., White et al. 2008; Shankar et al. 2010b; Conroy & White 2013; Zhang et al. 2023b), in this paper we choose to work with bolometric luminosities. Henceforth, L will always refer to the bolometric luminosity, i.e., $L \equiv L_{\text{bol}}$ ¹.

Within this framework, the Quasar Luminosity Function – $n_{\text{QLF}}(L)$ – is simply the marginalization of $n(L, M)$ over halo mass, M :

$$n_{\text{QLF}}(L) = \int_{M_{\text{min}}}^{M_{\text{max}}} \text{CLF}(L|M) n_{\text{HMF}}(M) dM. \quad (2.2)$$

¹However, note that the data considered in this paper always refer to type I, UV-bright quasars (e.g., Padovani et al. 2017). Hence, the model presented in this work describes only this specific population of active SMBHs.

Therefore, assuming that the halo mass function is known, the QLF can be easily determined once a conditional luminosity function has been adopted. The two limits of integrations, $M_{\min}$ and $M_{\max}$, represent the minimum/maximum mass of a halo that can host a SMBH. In principle, we could have SMBHs in any halos, and set this integration range to be as wide as possible. However, given that the simulations employed in our analysis span a wide but finite range of masses (Sec. 2.2.2), we adopt the following limits: $\log_{10} M_{\min}/M_{\odot} = 11.5$, and $\log_{10} M_{\max}/M_{\odot} = 14$ ($\log_{10} M_{\max}/M_{\odot} = 14.5$) at redshift $z = 4$ ($z = 2.5$). These limits enclose a range in masses that is sufficiently broad for our redshifts of interest (Sec. 2.4), so that expanding the range would have a negligible impact on our final results.

We use a model for the CLF in which the distribution in luminosity is log-normal at fixed mass (see also Ren et al. 2020; Ren & Trenti 2021):

$$\text{CLF}(L|M) dL = \frac{f_{\text{on}}}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\log_{10} L - \log_{10} L_c(M))^2}{2\sigma^2}\right) d\log_{10} L. \quad (2.3)$$

We then assume a power-law dependence of the characteristic luminosity, L_c , on mass:

$$L_c(M) = L_{\text{ref}} \left(\frac{M}{M_{\text{ref}}}\right)^{\gamma}, \quad (2.4)$$

or, in terms of logarithmic quantities:

$$\log_{10} L_c(M) = \log_{10} L_{\text{ref}} + \gamma (\log_{10} M - \log_{10} M_{\text{ref}}), \quad (2.5)$$

where M_{ref} is simply a reference mass that is associated with the reference luminosity L_{ref} . We fix $\log_{10} M_{\text{ref}}/M_{\odot} = 12.5$. The free parameters of the model are: σ , L_{ref} , γ , and f_{on} . In the following, we assume that these parameters do not depend on the other variables such as halo mass or quasar luminosity, and let them assume different values for the different redshifts we consider in Sec. 2.4.

The factor f_{on} accounts for the fact that not all black holes may be active as quasars at any given time. Therefore, we are implicitly assuming that the CLF is bimodal: the first mode accounts for all luminous quasars and is log-normally distributed, whereas the second mode (not accounted for in eq. 2.3) describes the behavior of the black holes that are too dim to be probed by any observations and is therefore completely irrelevant to our analysis. This bimodality in the CLF has a well-defined physical meaning: black holes are either active as luminous quasars or they are dormant, with a luminosity that is orders of magnitudes lower than any observational limits. However, it is not clear whether the luminosity distribution of black holes is indeed bimodal, or rather shows a continuum between active sources and inactive/faint ones. Observations of very faint quasars ($\log_{10} L/\text{erg s}^{-1} \approx 42 - 45$) can shed light on this question². We will return to this point in Sec. 2.5.3.

²Such observations become very difficult in the distant universe, as faint quasars are often outshined by their host galaxies.

2.2.1.1 The quasar auto-correlation function

In our framework, the correlation function of quasars is identical to the correlation function of the halos that host them, as quasars are temporally subsampling the underlying halo distribution. However, we have to consider that only quasars above some luminosity threshold L_{thr} are accounted for when measuring the correlation function in a survey. Therefore, we are effectively considering a “biased” halo mass distribution traced by the quasars above this luminosity threshold: we will refer to it as the “Quasar-Host Mass Function” (QHMF). This quantity can be expressed in terms of the halo mass function and another marginalization integral of the CLF:

$$n_{\text{QHMF}}(M|L > L_{\text{thr}}) = n_{\text{HMF}}(M) \int_{L_{\text{thr}}}^{\infty} \text{CLF}(L|M) dL. \quad (2.6)$$

The clustering of quasars can then be determined by computing the correlation function of a sample of halos that are distributed according to $n_{\text{QHMF}}(M|L > L_{\text{thr}})$. Here, we use an approach that allows us to quickly determine the quasar auto-correlation function for different $n_{\text{QHMF}}(M)$ distributions: we create different mass bins, and – selecting halos in these bins – extract the cross-correlation functions for halos with different masses from a cosmological simulation (see Sec. 2.2.2 for more details). Let us call these cross-correlation terms $\xi_h(M_j, M_k; r)$, with $M_{j,k}$ being the centers of the mass bins. We can then compute the quasar auto-correlation function, $\xi(r)$, by simply weighting the cross-correlations terms, $\xi_h(M_j, M_k; r)$, according to the quasar-host mass function, n_{QHMF} :

$$\xi(r) = \sum_{j,k} p_j p_k \xi_h(M_j, M_k; r), \quad (2.7)$$

where the weights $p_{j,k}$ are defined as:

$$p_i = \frac{n_{\text{QHMF}}(M_i|L > L_{\text{thr}}) \Delta M}{\int_0^{M_{\text{max}}} n_{\text{QHMF}}(M|L > L_{\text{thr}}) dM}, \quad (2.8)$$

with ΔM being the width of the mass bins. We present how to derive these equations in Appendix 2.A.

Once $\xi(r)$ is known, other related quantities such as the projected auto-correlation function, $w_p(r_p)$, can be easily obtained by integrating along the parallel direction π . The projected auto-correlation function is relevant since it can be directly compared to observational data (see Sec. 2.3.1). Setting a maximum value for the parallel distance π_{max} , which is chosen in accordance with the one used for observational data, e.g. $\pi_{\text{max}} = 100 \text{ cMpc } h^{-1}$ for the S07 measurements, the projected auto-correlation function reads:

$$w_p(r_p) = \int_{-\pi_{\text{max}}}^{\pi_{\text{max}}} \xi(r_p, \pi) d\pi = 2 \int_{r_p}^{\sqrt{r_p^2 + \pi_{\text{max}}^2}} \frac{r \xi(r)}{\sqrt{r^2 - r_p^2}} dr. \quad (2.9)$$

2.2.1.2 Halo occupation distribution and duty cycle

From the CLF, we can extract other quantities that will be relevant to our analysis. In particular, the integral of the CLF above some threshold luminosity L_{thr} represents the aggregate probability for a halo of mass M to host a quasar with a luminosity above the threshold value. Therefore, it is equivalent to a Halo Occupation Distribution (HOD; see e.g., [Berlind & Weinberg 2002](#)):

$$\text{HOD}(M) = \frac{n_{\text{QHMF}}(M|L > L_{\text{thr}})}{n_{\text{HMF}}(M)} = \int_{L_{\text{thr}}}^{\infty} \text{CLF}(L|M) dL. \quad (2.10)$$

The HOD is also closely related to the idea of a *quasar duty cycle*. In fact, the duty cycle is defined as the fraction of active quasars (i.e., with a luminosity above the threshold) divided by the fraction of halos that are able to host these quasars. In the standard picture (e.g., [Martini & Weinberg 2001](#); [Haiman & Hui 2001](#)) this fraction is well defined, as it is implicitly assumed that there is a minimum halo mass $\tilde{M}_{\text{min}}$ above which all halos can host quasars, and only a fraction of them is active at the present moment. In other words, the QHMF is:

$$n_{\text{QHMF}}(M) = \varepsilon_{\text{DC}} n_{\text{HMF}}(M) \Theta(\log_{10} M - \log_{10} \tilde{M}_{\text{min}}), \quad (2.11)$$

with ε_{DC} being the duty cycle and Θ the Heaviside step function. However, this definition of the duty cycle is not well-posed in our approach. As described above, we do not assume a specific functional form for the QHMF, but rather we infer this quantity from the CLF (eq. 2.6). As a consequence, we do not define a minimum mass $\tilde{M}_{\text{min}}$ for halos to host bright quasars, but rather consider all halos and compute a probability for them to host bright quasars given their mass, M . This implies that, in principle, even halos with a very low mass could have a small but non-zero probability to host bright quasars. As a result, the above definition of the quasar duty cycle would return artificially small values. Therefore, we opt here for an alternative definition (see also [Ren & Trenti 2021](#)): the duty cycle, ε_{DC} , is the weighted average of the HOD above a threshold mass that is given by the median of the QHMF, $n_{\text{QHMF}}(M|L > L_{\text{thr}})$. In other words, if we define the median of the QHMF as the mass M_{med} satisfying the relation:

$$\int_{M_{\text{med}}}^{M_{\text{max}}} n_{\text{QHMF}}(M) dM = 0.5 \int_{M_{\text{min}}}^{M_{\text{max}}} n_{\text{QHMF}}(M) dM, \quad (2.12)$$

then ε_{DC} can be expressed as:

$$\begin{aligned}\varepsilon_{\text{DC}} &= \frac{\int_{M_{\text{med}}}^{M_{\text{max}}} n_{\text{HMF}}(M) \text{HOD}(M) \text{d}M}{\int_{M_{\text{med}}}^{M_{\text{max}}} n_{\text{HMF}}(M) \text{d}M} = \\ &= \frac{\int_{M_{\text{med}}}^{M_{\text{max}}} n_{\text{QHMF}}(M|L > L_{\text{thr}}) \text{d}M}{\int_{M_{\text{med}}}^{M_{\text{max}}} n_{\text{HMF}}(M) \text{d}M}.\end{aligned}\tag{2.13}$$

2.2.2 Dark matter only simulation setup

In the last section, we have shown how we can make use of the CLF formalism to compute the quasar luminosity and auto-correlation functions – together with other relevant quantities such as the QHMF and the quasar duty cycle – using two fundamental ingredients: the mass function of halos and the cross-correlation functions of halos with different masses (see Figure 2.1 for a summary of this workflow). In this section, we provide details on how we obtain these two ingredients using the Dark-Matter-Only (DMO) version of the FLAMINGO suite of cosmological simulations.

FLAMINGO (Schaye et al. 2023; Kugel et al. 2023) is a suite of state-of-the-art, large-volume cosmological simulations run with the N-body gravity and smooth particle hydrodynamics (SPH) solver SWIFT (Schaller et al. 2024). Gravity is solved using the Fast Multiple Method (Greengard & Rokhlin 1987). The cosmology adopted in FLAMINGO is the “3x2pt + all” cosmology from Abbott et al. (2022) ($\Omega_{\text{m}} = 0.306$, $\Omega_{\text{b}} = 0.0486$, $\sigma_8 = 0.807$, $H_0 = 68.1 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $n_{\text{s}} = 0.967$), with a summed neutrino mass of 0.06 eV . Massive neutrinos are included in the simulation via the δf method of Elbers et al. (2021). Initial conditions (ICs) are set using multi-fluid third-order Lagrangian perturbation theory (3LPT). Partially fixed ICs are used to limit the impact of cosmic variance (Angulo & Pontzen 2016) by setting the amplitudes of modes with $(kL)^2 < 1025$ to the mean variance (k is the wavenumber and L the box size).

In this work, we focus on two specific DMO simulations with box sizes $L = 2800 \text{ cMpc}$ and $L = 5600 \text{ cMpc}$, respectively. Both simulations have 5040^3 cold dark matter (CDM) particles and 2800^3 neutrino particles. The CDM particle masses are $M_{\text{dm}} = 6.72 \times 10^9 M_{\odot}$ and $M_{\text{dm}} = 5.38 \times 10^{10} M_{\odot}$ for the $L = 2800 \text{ cMpc}$ and $L = 5600 \text{ cMpc}$ boxes, respectively. We focus on the DMO version of the simulations because no hydrodynamic version is available for the largest box, and because we are only interested in the spatial distribution of halos, that, in the ΛCDM model, is primarily dictated by gravitational interactions of dark-matter particles only.

We identify halos in the simulated snapshots using the 6-d friends-of-friends code VELOCIRAPTOR (Elahi et al. 2019). Once halos have been




Figure 2.2: *Left:* Halo mass function at $z = 4$ from the simulations considered: $L = 2800 \text{ cMpc}$ (teal diamonds) and $L = 5600 \text{ cMpc}$ (red circles). The gray solid line represents the analytical fit to the simulations (see Sec. 2.2.2.1 for more details). The shaded regions highlight which masses in each simulation are considered for the fit. The bottom panel shows the relative difference between the fit and the two simulations (with the horizontal shaded gray band highlighting the 10% limit). *Right:* Auto-correlation function of halos in different mass bins at $z = 4$. We create 8 mass bins ranging from $\log_{10} M / M_{\odot} = 11.5$ to $\log_{10} M / M_{\odot} = 13.5$ and 0.25 dex wide. Lower mass bins correspond to lower values of the correlation functions, and vice-versa. Teal diamonds refer to the $L = 2800 \text{ cMpc}$ simulation, while red circles refer to the $L = 5600 \text{ cMpc}$ one. Points are staggered in the x-direction for visualization purposes. The gray solid lines represent the fits to the auto-correlation functions (from the lowest mass bin on the bottom to the highest mass bin on top), as described in Appendix 2.B. Relative differences between the fits and the simulated correlation functions are shown in the bottom panel. These differences are generally $\lesssim 10\%$, with the exception of the highest mass bin considered (i.e., $\log_{10} M / M_{\odot} = 13.25 - 13.5$), for which the measurements are noisy due to the small number of halos in that mass range. The gray dashed lines in the top panel show extrapolations of the auto-correlation functions based on our fit for even higher mass bins ($\log_{10} M / M_{\odot} = 13.5 - 13.75$ and $\log_{10} M / M_{\odot} = 13.75 - 14.0$) where measurements from the simulations are not available. More details can be found in Sec. 2.2.2.2 and Appendix 2.B.

identified, their masses are computed using a spherical-overdensity definition based on their density profile. We perform this task using the code SOAP³. We define the radius of a halo as the distance from the most bound particle within which the density reaches a value of 200 times the critical density of the universe ($200\rho_c$). We only include central halos in the analysis and exclude the contribution of sub-halos. As discussed in Sec. 2.5.3, we do not expect this to influence our results significantly.

Once we have obtained a catalogue with the positions and masses of halos in the simulation at a given redshift, we can easily compute key statistical properties such as the halo mass function and the (cross-)correlation functions of halos with different masses. However, this approach is not directly suitable for our purposes. In fact, an important limitation of cosmological simulations is that they give reliable results only in a finite range of masses. The lower limit of this mass range is imposed by resolution: halos with fewer than 50 – 100 dark-matter particles are not well resolved, and thus cannot be trusted. The upper limit, on the other hand, is set by the box size of the simulation: if the number of halos with mass greater than some threshold M is small, these halos are too rare to get a reliable estimate of their statistical properties (e.g., their clustering).

For the problem we are facing here, we need to be able to reproduce the relative abundance of halos and their spatial distribution for a vast range of masses. For this reason, employing a single halo catalogue obtained using a simulation with a fixed box size is not the optimal strategy. Instead, we use here an approach consisting of two key steps: we first compute the quantities of interest (i.e., the halo mass function and the halo cross-correlation functions) from multiple simulations with different box sizes (and mass resolutions), and then we combine these different simulations by making use of analytical fitting functions. In this way, we can predict the abundance and spatial distribution of halos for all the masses that are well captured by the different simulations considered.

Table 2.1 summarizes the properties of the simulations we employ. In brief, we use the two different box sizes $L = 2800$ cMpc and $L = 5600$ cMpc to study the properties of low-mass and high-mass halos, respectively. For the 2800 cMpc box, we select halos in the range of masses $\log_{10} M/M_\odot = 11.5 - 13.0$; for $L = 5600$ cMpc, we focus on halos in the range $\log_{10} M/M_\odot = 12.5 - 13.5$. The lower limits are set to select only halos with at least ≈ 50 particles, whereas the upper limits are set to ensure overlap between the two mass ranges and to guarantee that all mass bins (up to at least $z = 4$) are populated with at least 5000 halos. In the following we describe in detail how we combine these simulations to obtain an analytical description of the halo mass function and of the cross-correlation function of halos with different masses.

³<https://github.com/SWIFTSIM/SOAP>

Table 2.1: Summary of the different FLAMINGO cosmological simulations employed in the analysis. The “fitting mass range” refers to the mass range selected for the fits of the halo mass function and the cross-correlation functions (Sec. 2.2.2.1 and 2.2.2.2, respectively). The redshifts considered in the analysis are $z = 4.0$ (high-redshift data; see Figure 2.2), and $z = 2.5$ (low-redshift data; see Figure 2.9).

Box size [cMpc]	Number of CDM particles	CDM particle mass [$\log_{10} M_{\odot}$]	Fitting mass range [$\log_{10} M_{\odot}$]	Snapshots considered (z)
2800	5040^3	9.83	11.5 – 13.0	2.5, 4.0
5600	5040^3	10.77	12.5 – 13.5	2.5, 4.0

2.2.2.1 Fitting the halo mass function

Following [Tinker et al. \(2008\)](#) (see also [Jenkins et al. 2001](#); [White 2001](#); [Warren et al. 2006](#)), we write the halo mass function in terms of the peak height of the density perturbations, $\nu = \delta_c/\sigma(M, z)$, where $\delta_c \approx 1.69$ is the critical linear density for collapse and $\sigma(M, z)$ is the variance of the linear density field smoothed on a scale $R(M)$ (see [Press & Schechter 1974](#); [Sheth & Tormen 1999](#)). According to this formalism, the mass function can be parametrized in terms of a universal function $f(\sigma)$:

$$f(\sigma; A, a, b, c) = A \left(\left(\frac{\sigma}{b} \right)^{-a} + 1 \right) e^{-c/\sigma^2}, \quad (2.14)$$

where $f(\sigma)$ is related to the mass function via the expression

$$\frac{dn}{dM}(M, z) = f(\sigma) \frac{\rho_{m,0}}{M} \frac{d \ln \sigma^{-1}}{dM}, \quad (2.15)$$

with $\rho_{m,0}$ being the mass density at $z = 0$.

We use the python package COLOSSUS ([Diemer 2018](#)) to compute the value of $\sigma(M, z)$ using the same cosmology as the FLAMINGO simulation (Sec. 2.2.2). We then use χ^2 -minimization to find the best-fitting parameters (A, a, b, c) for the analytical form of the halo mass function. We fit the number density of halos in different mass bins using halo catalogues from two different simulations, using two different (but partially overlapping) mass ranges (see Table 2.1). We assign Poissonian counting errors to every mass bin considered. We also experiment with changing these errors, and find that we achieve a better fit to the data by doubling the errors for the $L = 2800$ cMpc simulation, and halving the ones associated with the $L = 5600$ cMpc box. Note that this choice is arbitrary: our goal is not to provide a physically-motivated fit to the data, but simply to find a good analytical description of the halo mass function coming from simulations.

Figure 2.2 (left panel) shows the best-fitting mass function for $z = 4$, together with the data obtained from the simulations. Analogous results for $z = 2.5$ are shown in Appendix 2.C. The optimal parameter values for this mass function are: $A = 5.68 \times 10^{-5}$, $a = 1.65$, $b = 257$, $c = 1.16$. As shown in the lower left panel of Fig. 2.2, the fit provides a description of the simulated data with an accuracy of $\approx 5 - 10\%$ up to $\log_{10} M/M_\odot \lesssim 13.5$. As we will discuss in Sec. 2.5.3, this level of accuracy for the model is enough to provide a satisfactory description of the observed data.

Finally, we note that the reason why we have performed the fitting of the halo mass functions extracted from our simulations and did not consider the best-fitting parameters provided by [Tinker et al. \(2008\)](#) is because we found that, at $z \geq 4$, differences between the [Tinker et al. \(2008\)](#) model and our simulations were as high as 100% (see also [Yung et al. 2023](#)).

2.2.2.2 Obtaining the cross-correlation functions

We want to obtain the cross-correlation functions of halos with masses M_j and M_k , $\xi_h(M_j, M_k; r)$. In order to achieve this, we create a grid in mass and distance by considering 8 uniformly spaced bins in $\log_{10} M$, with a minimum halo mass of $\log_{10} M_{\min}/M_{\odot} = 11.5$ and a maximum of $\log_{10} M_{\max}/M_{\odot} = 13.5$, and 8 (logarithmically-spaced) bins in the radial direction with a minimum radial distance of $\log_{10} r_{\min}/\text{cMpc} = 0.4$ and a maximum of $\log_{10} r_{\max}/\text{cMpc} = 2.25$. We then use the package CORRFUNC (Sinha & Garrison 2020) to compute the number of halo pairs in the simulated catalogues for every combination of masses and distance, together with the number of pairs obtained assuming that these halos are distributed randomly. The values of the cross-correlation terms are obtained using the Landy & Szalay (1993) estimator:

$$\xi_h(M_j, M_k; r) = \xi_{j,k}(r) = \frac{D_j D_k - D_j R_k - D_k R_j + R_j R_k}{R_j R_k}, \quad (2.16)$$

where $D_j D_k$ stands for the number of pairs of halos in the mass bin j with halos in the mass bin k , whereas $R_j D_k$, $D_j R_k$, and $R_j R_k$ refer to the number of pairs when comparing to a random distribution of the same halos.

We end up with 36 different cross-correlation functions – i.e., the number of independent elements for a symmetric 64-element matrix – which can be used to determine the quasar auto-correlation function according to eq. 2.7. However, once again, we must account for the fact that different simulations probe different mass ranges. We thus fit a parametric analytical function to these cross-correlation functions in a way that allows us to combine different simulated boxes.

Furthermore, in this case the fitting procedure has another critical purpose. Despite the large volume of the simulations employed, the number of simulated halos at the very high mass end is limited by the finite size of the box. For this reason, the obtained cross-correlation terms for the very high-mass halo pairs will suffer from significant uncertainties due to the limited sample size in the simulation. Even for the largest box we consider (i.e., $L = 5600 \text{ cMpc}$), at $z = 4$ this effect starts to be significant for $\log_{10} M/M_{\odot} \approx 13.2 - 13.5$. This is an important limitation for our analysis: in the inference routine we will undertake in the next Section, we want to be able to explore the full parameter space and consider models for which this range of masses (or even higher) plays a significant role. For this reason, we fit the cross-correlation terms with two key objectives: reducing the noise associated with the poor statistics at the high mass end of the halo mass function, and providing a means to sensibly extrapolate the behavior of the cross-correlation functions up to $\log_{10} M/M_{\odot} = 14.0$ ($\log_{10} M/M_{\odot} = 14.5$) at $z = 4$ ($z = 2.5$). This extrapolation allows us to recover well-behaved posterior distributions (see Sec. 2.4) that provide a complete description

of the different models described by our parameters. Its validity and the associated caveats are discussed in detail in Sec. 2.5.3

We provide details on the fitting of the cross-correlation terms $\xi_h(M_j, M_k; r)$ in Appendix 2.B. In short, we divide all the cross-correlation terms, $\xi_h(M_j, M_k; r)$, by a reference correlation function, $\xi_{\text{ref}}(r)$, and fit the results with a 3-d polynomial to capture the residual dependencies on the two masses and on radius. In the rest of this Section, we show the results of the fits for the auto-correlation functions in different mass bins at $z = 4$ (Figure 2.2, right panel; the same plot for $z = 2.5$ is shown in Appendix 2.C). In other words, we plot the correlation functions for bins of equal mass, $\xi_h(M_j, M_j; r)$, together with the fits that are meant to reproduce these functions, $\xi_{h,\text{fit}}(M_j, M_j; r)$ (gray lines)⁴. Lower mass bins correspond to lower values of the auto-correlation functions, and vice-versa.

We assign errors to the $\xi_h(M_j, M_j; r)$ points based on the Poissonian statistics of the pair counts; note that in this way we are underestimating the real uncertainties on the data points because we are not including the effects of cosmic variance and of other sources of systematics. For this reason, when assessing the robustness of our fits, it makes little sense to discuss them in terms of statistical errors. We therefore compare the simulated data and the model fits in terms of relative differences between the two (lower right panel of Figure 2.2). These differences are generally at the level of $\lesssim 10\%$ for all bins but the highest one (i.e., $\log_{10} M/M_\odot = 13.25 - 13.5$), which is easily recognizable because it has the largest Poissonian uncertainties. As already mentioned before, at very large masses correlation measurements from simulations become noisy (and thus unreliable) due to the small number of halos in the snapshots. Even in this extreme case, however, the fit provides a satisfactory description of the shape and normalization of the correlation function in the simulations, with a relative difference that is still smaller than the uncertainties on the S07 observed data (which are at the level of $50 - 100\%$; see Sec. 2.3.1).

Using dashed gray lines, we also plot in Fig. 2.2 the auto-correlations functions for the two bins $\log_{10} M/M_\odot = 13.5 - 13.75$ and $\log_{10} M/M_\odot = 13.75 - 14.0$, as obtained by extrapolating our fitting functions to masses higher than the ones probed by the simulations. We see that the trend of the auto-correlation functions with halo mass is well preserved by these extrapolations; further discussion on this can be found in Sec. 2.5.3 and Appendix 2.C.

Finally, we note that relative differences between our fits and the values of correlation functions extracted from simulations tend to be larger at very large scales ($r \gtrsim 100$ cMpc). This is also due to the fact that simulation-based values become less reliable in this regime. There are two reasons for that:

⁴The global fits to all the cross-correlation terms $\xi_h(M_j, M_k; r)$ at both redshifts are shown in Appendix 2.B.

first, the finite size of the box reduces the number of very large-scale pairs that are available. Secondly, at $r \gtrsim 100$ cMpc the behavior of correlation functions becomes non-trivial due to the presence of the baryon acoustic oscillations (BAO) peak. This is especially difficult to model given the very coarse radial bins we have chosen. Due to these limitations of our model, we simply exclude scales larger than $r \gtrsim 100$ cMpc from our analysis.

2.3 Data-model comparison

In the previous Section, we have described how to obtain the two observables of interest (i.e., the QLF and the quasar auto-correlation function) starting from a CLF and a simulation-based analytical description of the halo mass function and of the halo cross-correlation functions. We now provide more details on the actual comparison between our model and observational data.

2.3.1 Overview of observational data

We start by giving a brief description of the data that we compare the model with. Our main goal is to explain the very strong quasar clustering measured by S07 at $z \approx 4$. Thus, we make use of the S07 data for the projected auto-correlation function (w_p/r_p). Note that the authors assume that the data points are independent (because the quality of the data is not good enough to extract a covariance matrix), so we will do the same and use the S07 errors assuming that the covariance matrix for the data is diagonal. We use the “good fields” data (see S07 for the definition) as they are supposed to be more reliable and – since they show stronger clustering – have proven to be the hardest to reproduce theoretically (e.g., Shankar et al. 2010b). As already mentioned, we exclude the data at very large scales ($r > 100$ cMpc) from our analysis because they are particularly challenging to measure both in observations (e.g., Eftekharzadeh et al. 2015) and in simulations (see the end of the last Section).

In the subsequent analysis, we are also interested in reproducing the quasar clustering at lower redshift. For this purpose, we use data from the Baryon Oscillation Spectroscopic Survey (BOSS, Eftekharzadeh et al. 2015; hereafter, E15). We focus on the redshift range $z = 2.2 - 2.8$, where the majority of the BOSS quasars reside. We use the data for the projected correlation function, $w_p(r_p)$, in the radial range $4 \text{ cMpc } h^{-1} < r_p < 25 \text{ cMpc } h^{-1}$. In this region, the E15 data are considered more reliable by the authors and an estimate for the error covariance matrix is available.

One of the key points of our analysis is that, while the QLF includes all quasars known in a given redshift bin, the quasar auto-correlation function is usually measured by considering only quasars above a given luminosity threshold L_{thr} . This is an important point to take into account in our model

(see eq. 2.6-2.10), as the presence of such a threshold may bias the inferred clustering significantly. The flux limit employed for the S07 measurements is $m_i = 20.2$ (where m_i is the apparent magnitude in the i band). In order to convert this to a value of L_{thr} , we first convert the apparent magnitude m_i to an absolute magnitude, M_{1450} , using the $K(z)$ correction⁵ (see, e.g., Kulkarni et al. 2019 and references therein). We obtain that $m_i = 20.2$ corresponds to $M_{1450} = -25.72$ at $z = 4$. We then convert this value to a bolometric luminosity by applying the bolometric corrections provided by Runnoe et al. (2012a)⁶. We get a value for the S07 luminosity threshold equal to $\log_{10} L_{\text{thr}}/\text{erg s}^{-1} = 46.7$.

As for the E15 clustering data at $z \approx 2.5$, the luminosity threshold that we should employ is more subtle. While the authors consider the entirety of the BOSS sample (Ross et al. 2013) for their clustering analysis, they also show that this sample is highly incomplete at low luminosities. This is an issue in the context of our model, as, when setting a threshold L_{thr} , we are implicitly assuming that the sample is complete above the threshold. Given that properly modeling completeness in the E15 sample is outside the scope of this work, we set the value of L_{thr} to the 25th percentile of the luminosity distribution of the observed quasars at $z = 2.5$. This value represents a compromise between taking into account part of the highly incomplete sample of faint quasars that are included in the clustering analysis and minimizing the bias that these quasars generate in the predicted clustering. By considering Figure 3 in E15, we set this threshold value to a $M_i(z = 2)$ magnitude of -25.3 . Following Lusso et al. (2015), we convert this to $M_{1450} = M_i(z = 2) + 1.28 = -24.02$, and finally to a bolometric threshold of $\log_{10} L_{\text{thr}}/\text{erg s}^{-1} = 46.1$.

As for the QLF, there are many different estimates available. For the sake of consistency with the clustering measurements, we choose to employ the UV-bright quasar catalogue compiled by Kulkarni et al. (2019, hereafter K19). These authors provide a homogenised catalogue of 80,000 color-selected AGN from redshift $z = 0$ to 7.5, together with MCMC-based estimates of the QLF at all redshifts. We employ this dataset and select quasars at different redshifts according to our models. For the model at $z = 4$, we set $3.5 < z < 4.5$ (largely consistent with the S07 high- z sample); in this range, the bright end of the QLF is determined by the same SDSS quasars that are used to compute the clustering (Schneider et al. 2010), whereas the low-luminosity quasars are presented in Glikman et al. (2011). The model

⁵The conversion between m_i and M_{1450} can be made using $K_{i,1450}(z)$, which is defined as: $M_{1450}(z) = m_i - 5 \log_{10}(d_L(z)/\text{Mpc}) - 25 - K_{i,1450}(z)$, with $d_L(z)$ being the luminosity distance at redshift z . Following Lusso et al. (2015), we set $K_{i,1450}(z = 4) \approx -1.9$.

⁶The bolometric correction for $\lambda = 1450 \text{ \AA}$ is $\log_{10} L_{\text{iso}}/\text{erg s}^{-1} = 4.745 + 0.910 \log_{10} \lambda L_{\lambda}/\text{erg s}^{-1}$. L_{iso} is the bolometric luminosity calculated under the assumption of isotropy, and it is related to the real bolometric luminosity L via a factor that accounts for the viewing angle, $L = 0.75 L_{\text{iso}}$.

at $z = 2.5$, instead, is entirely determined by quasars observed by the BOSS survey (Ross et al. 2013).

Quasars in the K19 dataset are binned according to their M_{1450} magnitude, and the uncertainties are computed using Poisson statistics. However, as also discussed in K19, the QLF data always present significant systematic errors due to, e.g., uncertainties in the quasar selection. This implies that the quoted uncertainties on the QLF data may be significantly underestimated, as is also evident from the large scatter (up to ≈ 1 dex) between different estimates of the QLF that are available in the literature (e.g., Shen et al. 2020; Grazian et al. 2023). These issues are particularly problematic in our framework, as our goal is to perform statistical inference by simultaneously matching the quasar luminosity and auto-correlation functions, and that can only be done properly if the associated uncertainties are well understood and treated. Therefore, in order to avoid biases in our inference analysis owing to very small formal statistical uncertainties on the QLF, we add a systematic error to every QLF measurement in quadrature to the Poisson ones determined by K19. That is, the uncertainties on our QLF data points are set to be $\sigma^2 = \sigma_{\text{sys}}^2 + \sigma_{\text{count}}^2$, where σ_{sys}^2 (σ_{count}^2) stands for the systematic (statistical) uncertainty. We adopt a constant systematic uncertainty of 0.2 dex for the $z \approx 4$ dataset and of 0.05 dex for the $z \approx 2.5$ one. This implies a systematic relative uncertainty of $\approx 45\%$ ($\approx 10\%$) for $z = 4$ ($z = 2.5$). These values are chosen to be similar to the average relative statistical uncertainties at the two redshifts considered ($\approx 40\%$ and $\approx 5\%$ at $z = 4$ and $z = 2.5$, respectively).

As the final step, we convert the values of the quasars' absolute magnitudes in K19, M_{1450} , to bolometric luminosities using the Runnoe et al. (2012a) bolometric corrections. We stress the fact that our results are independent of the adopted bolometric corrections, as our model could easily be expressed in terms of quasars' UV magnitudes only. However, as discussed in Sec. 2.2.1, we choose to convert everything to bolometric luminosities for consistency with previous work on the subject.

2.3.2 Likelihood functions

We employ a Bayesian framework to write the posterior distributions for our model parameters. As described in Sec. 2.2.1, the model consists of a log-normal CLF centered on a power-law dependence of the quasar luminosity on halo mass. The free parameters are the normalization and slope of the quasar luminosity-halo mass relation (L_{ref} and γ , respectively), the logarithmic scatter around this relation (σ), and the fraction of quasars that are active at any given moment (f_{on}). The final set of parameters, Θ , is then: $(\sigma, L_{\text{ref}}, \gamma, f_{\text{on}})$.

We set flat priors on σ and γ , and flat priors on the logarithm of L_{ref} and f_{on} (see e.g. Jeffreys 1946). Our priors span the following parameter

ranges: $\sigma \in (0.1 \text{ dex}, 1.5 \text{ dex})$; $\log_{10} L_{\text{ref}}/\text{erg s}^{-1} \in (44.0, 46.6)$; $\gamma \in (0.5, 3)$; $\log_{10} f_{\text{on}} \in (-3, 0)$. These limits are chosen in order to focus on the region of the parameter space where models are physically motivated (e.g., the scatter in the $L_c - M$ relation is unlikely to be smaller than 0.1 dex).

In what follows, we want to fit the QLF and the auto-correlation function both independently and simultaneously. We can get constraints on these two observables by setting the following likelihood functions:

$$\mathcal{L}^{(k)}(\mathbf{d}^{(k)} | \Theta) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{y} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{y} - \boldsymbol{\mu}) \right), \quad (2.17)$$

where $k \in \{\text{QLF}, \text{corr}\}$ stands either for the quasar luminosity function data or for the auto-correlation function data. As for the other variables, $\mathbf{d}^{(k)}$ stands for the set of n data points with means $\mathbf{y}$ and covariance Σ coming from observations, whereas $\boldsymbol{\mu}$ stands for the set of values predicted by our models.

With the above likelihood, the results for the correlation function ("corr") are found to not be very constraining, as there is a large set of models that produce the correct clustering but substantially under(over)-estimate the number density of bright quasars. Therefore, when quoting results for the correlation function only, we provide an additional integral constraint by imposing that the model matches the observed number density of bright quasars. We integrate the QLF above the luminosity limit used for the clustering measurements (see Sec. 2.3.1), and obtain an estimate for the number density of bright quasars, n_{bright} . The associated uncertainty, σ_{bright} , is determined by using different realizations of the QLF fits from K19. Then, we predict the number of quasars with a luminosity above this threshold, L_{thr} , based on our model (n_{model}), and use the following likelihood:

$$\mathcal{L}^{(\text{corr}+\text{nden})} = \frac{\exp^{-(n_{\text{bright}} - n_{\text{model}})^2 / \sigma_{\text{bright}}^2}}{\sqrt{2\pi} \sigma_{\text{bright}}} \mathcal{L}^{(\text{corr})}. \quad (2.18)$$

Note that we do not fit to the shape of the QLF, but only to the total abundance of quasars above L_{thr} . This is an integral constraint that favors models producing a physically reasonable total number of bright quasars.

Finally, we provide joint constraints on the parameters by fitting the QLF and the auto-correlation function simultaneously. In other words, we write the joint likelihood distribution as the product of the two likelihoods (we assume that the two measurements are independent, and weigh the two dataset equally):

$$\mathcal{L}^{(\text{joint})} = \mathcal{L}^{(\text{QLF})} \mathcal{L}^{(\text{corr})}. \quad (2.19)$$

Note that for the joint likelihood distribution, we consider $\mathcal{L}^{(\text{corr})}$ (rather than $\mathcal{L}^{(\text{corr}+\text{nden})}$), as the QLF already provides an implicit constraint on the total abundance of luminous sources.

Table 2.2: Best-fitting parameter values for our model-data comparison at $z = 4$ (see the main text for definitions of the different parameters, as well as eq. 2.17-2.19). “QLF only” refers to the quasar luminosity function data only, “corr+nden” refers to the auto-correlation function data in conjunction with the number density of bright quasars, and “joint” refers to the combined QLF+auto-correlation function data. The last column shows the minimum value of the normalized chi-squared (see text for details).

Quantity	σ	$\log_{10} L_{\text{ref}} [\text{erg s}^{-1}]$	γ	$f_{\text{on}} [\%]$	χ^2_{norm}
QLF only	0.38	46.4	0.78	0.2	2.2/5
corr+nden	0.10	44.4	2.99	100	4.6/4
joint	0.11	45.1	2.07	66	12.9/12

2.4 Results

In this section, we describe the results we obtain by fitting our model to the observed quasar luminosity and auto-correlation functions, both independently (“QLF” and “corr+nden” cases) and simultaneously (“joint” case; see Sec. 2.3.2 for the definitions). Henceforth, we will refer to the “QLF” model as “QLF only” in order to distinguish our model from the QLF itself. We first consider the $z = 4$ case – which is the main focus of this paper – and then discuss the results at lower redshift ($z = 2.5$) as well.

2.4.1 Analysis at $z \approx 4$

As a first step, we are interested to know whether our model can reproduce the two observables. We can answer this question by employing a simple optimization algorithm to find the maximum of the likelihood distributions (or, equivalently, of the posterior distributions) for the three cases of interest: quasar luminosity function only (“QLF only”), correlation function + number density of bright quasars (“corr+nden”), and quasar luminosity and correlation functions together (“joint”). The maxima of the likelihoods represent our best-fitting models, which we can then compare directly with observations (see Section 2.4.1.1 for the results of the full parameter inference).

In Table 2.2, we report these best-fitting parameters for the cases mentioned above. Figure 2.3 shows our model predictions at the maximum likelihood parameter values for the CLF, the HOD, the QHMF, the QLF, and the projected quasar autocorrelation function (w_p/r_p); see Fig. 2.1 for a schematic overview of these quantities.

In the top right panel of Fig. 2.3, we show the conditional luminosity functions, $\text{CLF}(L|M)$, as a function of the quasar luminosity L and the halo mass M . The three cases “QLF only”, “corr+nden”, and “joint” are shown with different colors (blue, orange, and green, respectively). The associated




Figure 2.3: Overview of our model-data comparison at $z = 4$. The blue, orange, and green colors refer to the best-fitting models for the “QLF only”, “corr+nden”, and “joint” likelihood distributions, respectively (see text for details and Table 2.2 for the parameters’ values). The upper right panel shows the Conditional Luminosity Function (CLF($L|M$)), with the associated color bars at the bottom representing the probability density. The horizontal dot-dashed gray line in the same panel (and the vertical one in the upper left panel) refers to the luminosity threshold for clustering measurements, $\log_{10} L_{\text{thr}}$. Integrating the CLF along the luminosity axis above this threshold gives the Halo Occupation Distribution (HOD; middle right panel), which can be combined with the Halo Mass Function (QHMF), to give the Quasar-Host Mass Function (QHMF; coloured lines in the lower right panel). Vertical dot-dashed lines in the lower right panel are the median values of the QHMF distribution, M_{med} (see eq. 2.12). The QHMF is then used to predict the projected quasar auto-correlation function, $w_p(r_p)$ (lower left panel). The $z \approx 4$, S07 data for the auto-correlation function are also shown in the same panel (data outside the fitting range, $3 < r_p/\text{cMpc} < 100$, are shown as semi-transparent points). The upper left panel shows our predictions for the Quasar Luminosity Function (QLF), together with the $z \approx 4$ data from K19.




Figure 2.4: *Left:* Corner plots of the 4-d posterior distributions for the different cases described in Sec. 2.4.1 (blue for the “QLF” model, orange for “corr+nden”, and green for “joint”). Contours in the 2-d histograms highlight the 1σ and 2σ regions, whereas the dashed lines in the 1-d histograms represent the median values of the parameters (with 1σ errors shown as shaded regions). Best-fitting parameters from Table 2.2 (see also Fig. 2.3) are shown with star symbols in each corner plot. *Right:* Comparison of the predicted quasar luminosity (top) and auto-correlation (bottom) functions with the observational data from K19 and S07, respectively. The color coding is the same as in the left panel. Median values (solid lines) and 1σ uncertainty regions (shaded areas) are obtained by randomly sampling the Markov chains for the posterior distribution 100 times. Data points for the auto-correlation function that are outside of our fitting range (see Sec. 2.3.1) are shown as semi-transparent points in the bottom right panel. The vertical dot-dashed line in the upper right panel is the luminosity threshold for quasar clustering, L_{thr} (see Sec. 2.3).

color bars at the bottom of the Figure represent the probability densities for the different CLF cases. Integrating the CLF above the luminosity threshold L_{thr} (gray dashed-dotted line in the CLF panel), we obtain the halo occupation distribution (HOD; middle right panel; eq. 2.10). Combining the HOD with the halo mass function (HMF), we get the Quasar-Host Mass Function (QHMF; eq. 2.6); this is shown in the bottom right panel, together with the $z = 4$ HMF (gray line). The two left panels show the predictions for the observable quantities: the auto-correlation function is shown on the bottom left, together with data from S07; the quasar luminosity function (eq. 2.2) is shown on top (data are from K19). While the auto-correlation function is obtained from the QHMF via eq. 2.7-2.9, the QLF is the result of integrating along the mass axis of the CLF weighted by the HMF (eq. 2.2).

Overall, looking at the two left panels of Figure 2.3, we conclude that in all cases the models constitute very good fits to the data they are meant to reproduce (see below for the caveat on the “QLF only” case). In order to quantify this, we use reduced chi-squared statistics, $\chi^2_{\text{norm}} = \chi^2/\nu_{\text{ndof}}$, where ν_{ndof} is the number of degrees of freedom (i.e., the number of data points minus the number of parameters). We find $\chi^2_{\text{norm}} = 2.2/5$, $\chi^2_{\text{norm}} = 4.6/4$, and $\chi^2_{\text{norm}} = 12.9/12$ for the “QLF only”, “corr+nden”, and “joint” cases, respectively. These values are also shown in Table 2.2 for reference.

One striking feature of the best-fitting models is that they have very different properties, as can be seen in the top right panel of Fig. 2.3 and the best-fitting parameters shown in Table 2.2. All of them are characterized by low values of the scatter in the quasar luminosity-halo mass relation, σ , but the offset, slope, and normalization of this relation vary significantly between the models.

The “QLF only” model shows an approximately linear relation with a high value of the reference luminosity L_{ref} . As a result, the characteristic mass of halos hosting quasars with a luminosity above L_{thr} is low ($\log_{10} M/M_{\odot} \approx 12.35$, lower right panel of Fig. 2.3). This has two consequences. Firstly, halos with $\log_{10} M/M_{\odot} \approx 12 - 12.5$ are much more abundant than the number of observed quasars, and thus a very low active fraction ($f_{\text{on}} \approx 0.1\%$) is needed to match the QLF. Secondly, such a low characteristic mass for the halos hosting luminous quasars implies a low value for the quasar auto-correlation function, in conflict with the S07 measurements (lower left panel). In fact, we see that the best-fitting model for the “QLF only” case does not fare well when compared with the clustering data.

The “corr+nden” model, instead, finds a much larger characteristic host mass for bright quasars ($\log_{10} M/M_{\odot} \approx 13 - 13.5$). Such a large mass is achieved by packing quasars in almost all the most massive halos. This is done thanks to a few key ingredients (upper right panel of Fig. 2.3): a low value of the quasar luminosity at the reference mass of $\log_{10} M_{\text{ref}}/M_{\odot} = 12.5$, a highly non-linear relation between quasar luminosity and halo mass ($\gamma \approx 3$) and a very low scatter in this relation $\sigma \approx 0.1$. The first two parameters

determine the mass, $\log_{10} \tilde{M}$, at which the quasar luminosity – halo mass relation crosses the luminosity limit L_{thr} . The second and third parameters, instead, determine how sharply the HOD drops at masses lower than $\log_{10} \tilde{M}$ (middle right panel of Fig. 2.3). The extreme scenario implied by our best-fitting model is needed to reproduce the measured auto-correlation function. Indeed, the shape and normalization of the S07 data are very well reproduced by our model (Fig. 2.3, lower panel) at the scales considered in the analysis ($3 \text{ cMpc} \lesssim r_p \lesssim 100 \text{ cMpc}$).

Besides fitting the auto-correlation function, the “corr+nden” model aims to reproduce the number density of quasars above the luminosity threshold $\log_{10} L_{\text{thr}}$. This is also achieved by the best-fitting model, which predicts a number density $n_{\text{model}} = 3.18 \times 10^{-8} \text{ cMpc}^{-3}$, 0.5 standard deviations higher than the observational value of $n_{\text{bright}} = 2.73 \times 10^{-8} \text{ cMpc}^{-3}$. The shape of the QLF, however, is not well reproduced by the model, because it overpredicts the abundance of very bright systems and underpredicts the abundance of $\log_{10} L/\text{erg s}^{-1} \approx 46 - 47$ quasars. This is due to the fact, despite the very low value of σ , the strong non-linearity in the quasar luminosity-halo mass relation ($\gamma \approx 3$) associates a large fraction of the massive halos to the brightest observable quasars.

When we simultaneously fit both the quasar auto-correlation and the luminosity function (“joint” model), we obtain results that are quite similar to the “corr+nden” case, and are compatible with the same extreme scenario in which quasars are packed in the most massive halos, i.e., a non-linear quasar luminosity-halo mass relation with a steep slope and very small scatter, low value of the quasar luminosity at the reference mass, and a large active fraction of quasars. The quasar luminosity-halo mass relation for the “joint” model is however not as extreme as the one for the “corr+nden” model, as it is characterized by a lower value of the power-law exponent, $\gamma \approx 2$. This has very little impact on the auto-correlation function, as the quasar-host mass functions (lower right panel of Fig. 2.3) are very similar in the two cases. It does have an effect, however, on the shape of the QLF, with the “joint” model providing a better fit, especially at the very bright end.

Overall, the QLF is very well reproduced by the “joint” model, with the exception of the low-luminosity end ($\log L/\text{erg s}^{-1} \approx 45.5$). In this region, the largest differences between the “QLF only” and the “joint” model appear, with the “QLF only” model faring better at predicting a flattening of the shape of the QLF. This flattening, however, is an artificial feature of our model, originating from the prior assumption that halos with a mass lower than $\log_{10} M/M_{\odot} = 11.5$ do not host quasars. We consider this issue not worthy of further investigation, as the faint-end of the QLF is still largely unconstrained by data, and deeper observations are needed to probe its behavior at the high redshift (e.g., Akiyama et al. 2018; Parsa et al. 2018; Giallongo et al. 2019; Harikane et al. 2023; Grazian et al. 2023). Furthermore,

our primary focus here is to interpret the bright quasars that are also probed by clustering surveys. It is possible that a more flexible quasar luminosity-halo mass relation is necessary to account for the abundance of low-luminosity systems.

2.4.1.1 MCMC analysis

Given that our models are a good representation of the observational data, we can proceed further with inference and determine how well the data constrain the model parameters. We explore the posterior distributions using a Markov-Chain Monte Carlo (MCMC) approach. We employ the Python package `emcee` (Foreman-Mackey et al. 2013) to sample the posteriors using the affine-invariant ensemble prescription (Goodman & Weare 2010). We place $m = 48$ walkers distributed randomly in the parameter space and evolve them for $N > 10^5$ steps. We set the final number of steps so that our chains are at least 100 times longer than the auto-correlation time τ (see e.g., Sharma 2017), and thin the chains considering only one element every τ steps in order to account for auto-correlations. We also discard the first 10^3 elements of every chain to account for the burn-in phase.

Figure 2.4 (left panel) shows the corner plot for the 4-d posterior distributions (as a function of $\sigma, L_{\text{ref}}, \gamma, f_{\text{on}}$) for the three cases considered in the analysis (“QLF”, “corr+nden”, and “joint”). The best-fitting model for each of these cases, which was discussed above and shown in Fig. 2.3, is highlighted with a star symbol in the corner plots. The samples of the posterior distributions obtained by the Markov Chains are then used to obtain predictions for the quasar luminosity and the auto-correlation functions; we compare these quantities with the data in the right panels of Figure 2.4.

As expected, the “QLF only” and “corr+nden” models peak in very different regions of the parameter space. The “corr+nden” model constrains the parameters to the region with $\sigma \lesssim 0.5$, $\gamma \gtrsim 2$, $\log_{10} L_{\text{ref}}/\text{erg s}^{-1} \approx 44.5 - 45.5$, and f_{on} close to unity. This region of the parameter space is the only one that is compatible with the above-mentioned scenario in which bright quasars are active only in the most massive halos. This is also the reason why there are no models predicting stronger quasar clustering than observed (Figure 2.4, right panel), as our models are already predicting the strongest possible clustering compatible with the observed abundance of bright systems.

The “QLF only” model, on the other hand, peaks at lower γ and f_{on} , larger $\log_{10} L_{\text{ref}}$, and a value of σ which is larger than the “corr+nden” but still moderately low ($\sigma \approx 0.3 - 0.5$). However, the distribution for the “QLF only” case is much more complex, and therefore the resulting constraints on the parameters are not as straightforward. In particular, there is a region of the parameter space that is well within the constraints given by the auto-correlation function, and for which the “QLF only” model also has a

good match with the QLF data (at the $\lesssim 2\sigma$ level). Unsurprisingly, this is the region where the “joint” posterior distribution is located (green contours in Fig. 2.4).

The reason why this region of parameter space can reproduce both the QLF and the auto-correlation function can be understood as follows. As mentioned in the introduction, the behavior of the QLF at the bright end is very different from the one of the HMF, with the latter being characterized by an exponential cutoff that is not present in the QLF. This is usually explained by assuming that the more abundant population of lower-mass halos can also contribute to the population of luminous quasars (see e.g., Ren et al. 2020). Indeed, this is what our fiducial “QLF only” model seems to suggest (see also Fig. 2.3), as the very high quasar luminosity predicted for the $10^{12.5} M_\odot$ halo population implies that even with ≈ 0.4 dex of scatter the correct shape of the luminosity function can be reproduced. In this picture, quasars are relatively common phenomena arising in the bulk of the halo population at that redshift, with a very low duty cycle of $\varepsilon_{\text{DC}} \approx 0.1\%$. However, even a scenario in which only the most massive halos are active as bright quasars (with a duty cycle $\varepsilon_{\text{DC}} \gtrsim 50\%$) can be compatible with the observed shape of the QLF. In this second case, the non-linearity in the quasar luminosity–halo mass relation plays a key role in mapping the exponential cutoff of the HMF into the power-law bright end slope of the QLF, while at the same time packing bright quasars only in the most massive hosts. While both of these scenarios provide a good description of the QLF – differing significantly only at lower luminosities – only the latter is also compatible with a very large clustering length of quasars.

In conclusion, despite the fact that the quasar luminosity and auto-correlation functions alone provide relatively loose constraints on the shape of the Conditional Luminosity Function (CLF), when considered in conjunction they are able to determine a very well-defined region in the parameter space for which a good agreement with all observational data is achieved (right panel of Fig. 2.4). This is the most significant conclusion of our analysis, and we will discuss it further in Sec 2.5.

2.4.2 Comparison with $z \approx 2.5$

Having applied our model to $z \approx 4$ data, it is also important to test whether the model is flexible enough to reproduce observations at lower redshifts, where the observed strength of quasar clustering is not as extreme. We note that our goal in this paper is not to provide a complete and self-consistent evolutionary description of quasar properties across cosmic time, but simply to strengthen the conclusions we have drawn in the previous section by showing that the same framework can also be applied to describe the spatial and luminosity distributions of quasars at different epochs. In particular, we focus on the redshift range $z = 2.2 - 2.8$, where the BOSS survey (Ross




Figure 2.5: Same as Fig. 2.4, but for different redshifts ($z = 4$ in green and $z = 2.5$ in purple). The results always refer to the “joint” model (Sec. 2.4.1). The vertical dot-dashed lines in the upper right panel are the luminosity thresholds, L_{thr} , used to measure quasar clustering at the two redshifts. Data points for the auto-correlation function that are outside of our fitting range because they are considered not reliable (see Sec. 2.3.1) are shown with semi-transparent colours in the bottom right panel.

et al. 2013; Eftekharzadeh et al. 2015) has provided solid measurements of the quasar luminosity and auto-correlation functions. We choose this data set because it is sufficiently different from the one at $z \approx 4$ to suggest that the properties of quasars may have varied significantly in a relatively short amount of time.

For simplicity, we focus here on the “joint” models only. In other words, we run the MCMC-based algorithm with the same setup as in Sec. 2.4.1.1 fitting the quasar luminosity function and the auto-correlation function simultaneously. Figure 2.5 shows the resulting corner plots for the posterior distribution of the “joint” model at $z = 2.5$ (purple), together with the one at $z = 4.0$ (green; same as Fig. 2.4) for comparison. We report the values of the resulting 1-d constraints on the model parameters for both redshifts in Table 2.3. In the right panel of Figure 2.5, we show the predictions of our models based on randomly sampling the Markov chains for the posterior distributions, together with the data that we aim to reproduce. We note that, as mentioned in Sec. 2.3.1, we only include the data points for the E15 auto-correlation function at $z \approx 2.5$ in the range $6 \text{ cMpc} \lesssim r_p \lesssim 40 \text{ cMpc } h^{-1}$. This is because data outside this range are not considered reliable and not included in the covariance matrix estimation (see E15). Indeed, we find that our model provides a good match to the E15 data within the fitting range, but it is significantly lower than the measured data at larger scales. Given the strong biases that may be associated with large-scale estimates of the correlation function, we do not consider this issue worthy of further investigation.

The corner plots in Figure 2.5 show that the regions of the parameter space constrained by the two redshifts are quite different. Interestingly, the shape of the $z = 2.5$ posterior distribution exhibits non-trivial behavior in the 2-d projections, yielding tight constraints on the γ parameter, but also strong degeneracies between σ , $\log_{10} L_{\text{ref}}$, and f_{on} . In general, however, the different parameters are well constrained, even better than at $z = 4$ due to the higher sensitivity of the data. The resulting 1-d posteriors for $z = 2.5$ and $z = 4$ peak at a similar value of $\log_{10} L_{\text{ref}}$, but they are quite different for the other parameters. Lower- z results are characterized by a lower value of γ (≈ 1.15) and f_{on} (≈ 0.01), and a higher value of the scatter in the $L - M$ relation, σ .

The top panel of Figure 2.6 shows how these posteriors translate into distributions for the QHMFs (eq. 2.6). In this plot, the QHMFs for $z = 2.5$ and $z = 4$ are shown, together with the HMFs at the same redshifts (semi-transparent lines). Uncertainties on the QHMFs are computed by randomly subsampling the Markov chains for the posterior distributions. The two QHMFs are quite different, reflecting the differences in the level of clustering measured at the two redshifts. In the $z = 4$ case, quasars only reside in the most massive systems ($\log_{10} M/M_{\odot} \gtrsim 13$), with the QHMF distribution tightly following the HMF (see also Fig. 2.3 for the best-fitting model).

At $z = 2.5$, instead, the QHMF distribution has a lower median value ($\log_{10} M/M_{\odot} \approx 12.5$) and it is much broader, with a large range of halos of different masses capable of hosting quasars.

The differences in the QHMF translate directly into different measurements for the quasar duty cycle. As discussed in Sec. 2.2.1, we define the quasar duty cycle (eq. 2.13) as the ratio between the QHMF integrated above the median value of its distribution, M_{med} (eq. 2.12), and the HMF integrated above the same threshold. For $z = 2.5$, we find a value of the duty cycle equal to $\varepsilon_{\text{DC}} = 0.4 \pm 0.1 \%$, whereas for $z = 4$ we find $\varepsilon_{\text{DC}} = 33^{+34}_{-23} \%$. We note that these values are closely related to the values of the f_{on} parameter (Table 2.3), which describes the active fraction of quasars at any given moment. Only in the case of a perfectly deterministic $L - M$ relation (i.e., with zero scatter), however, would we find a duty cycle exactly equal to f_{on} . In the presence of scatter in the $L - M$ relation, the shape of the QHMF can vary significantly with respect to the one of the HMF, and this changes the fraction of quasars that are above the threshold luminosity, L_{thr} , at any given mass, and hence the quasar duty cycle.

However, we should mention the caveat that these results are obtained by setting two different luminosity thresholds, L_{thr} , at the two redshifts considered, according to the minimum luminosities imposed in the respective clustering measurements. As shown in the top right panel of Figure 2.5, the $z = 2.5$ luminosity threshold is ≈ 0.6 dex lower than the one at $z = 4$ ($L_{\text{thr}} = 46.1 \text{ erg s}^{-1}$ and $L_{\text{thr}} = 46.7 \text{ erg s}^{-1}$ at $z = 2.5$ and $z = 4$, respectively). Changing the value of L_{thr} may have direct consequences for the QHMF, HOD, and quasar duty cycle, since all these quantities have an explicit dependence on L_{thr} (eq. 2.6-2.13).

In order to provide a fair comparison between these quantities at the two redshifts considered in the analysis, we impose the same L_{thr} at both redshifts by using the $z = 4$ luminosity threshold (i.e., $L_{\text{thr}} = 46.7 \text{ erg s}^{-1}$) to recompute the above-mentioned quantities at $z = 2.5$. While the duty cycle remains unchanged (within uncertainties), we find that although the QHMF is still very broad, its normalization and median value are lower and higher, respectively. In particular, the median value of the QHMF, M_{med} (eq. 2.12) shifts from $\log_{10} M_{\text{med}}/M_{\odot} \approx 12.5$ to $\log_{10} M_{\text{med}}/M_{\odot} \approx 12.8$. This suggests a mild dependence of clustering on luminosity, as more luminous quasars tend to be hosted by more massive halos. However, given that the QHMF distribution is very broad in both cases, there is a strong overlap between the populations of very bright ($\log_{10} L/\text{erg s}^{-1} \gtrsim 46.5$) and moderately luminous ($\log_{10} L/\text{erg s}^{-1} \approx 46 - 46.5$) quasars in terms of their host halo masses.

We leave a detailed analysis of the implications of our model in terms of the luminosity dependence of quasar clustering for future work. Here, we simply note that even when adopting the same luminosity threshold, we find a remarkable difference between $z = 4$ and $z = 2.5$ quasars. The former are very extreme objects, hosted only by the most massive halos that are

Table 2.3: Constraints on the model parameters based on the corner plots shown in Figure 2.5.

Redshift	σ	$\log_{10} L_{\text{ref}} [\text{erg s}^{-1}]$	γ	$f_{\text{on}} [\%]$
$z = 2.5$	$0.52^{+0.18}_{-0.22}$	$45.7^{+0.46}_{-0.61}$	$1.15^{+0.06}_{-0.07}$	$1.01^{+1.59}_{-0.04}$
$z = 4$	$0.20^{+0.13}_{-0.08}$	$45.2^{+0.3}_{-0.3}$	$2.00^{+0.22}_{-0.23}$	51^{+32}_{-31}

present at that redshift, representing $4 - 5\sigma$ peaks in Gaussian random fields (Diemer 2018; see also Sec. 2.1). The latter, instead, are hosted by much more common halos at $z = 2.5$, which are only slightly over-massive with respect to the bulk of the halo population at that redshift ($2 - 3\sigma$ peaks). For this reason, despite the increase in the quasar number density between $z = 4$ and $z = 2.5$, the quasar duty cycle – which measures how abundant quasars are with respect to their host population – decreases by two orders of magnitudes between the same two redshifts.

In conclusion, our data-model comparison reveals that the same parametrization of the CLF employed at $z = 4$ is also able to reproduce the data at lower- z , with a significant evolution of the CLF parameters reflecting a remarkable change in the physical properties of quasars with cosmic time. In the following, we further discuss the implications of these findings.

2.5 Discussion

In the analysis performed above, we could successfully match the quasar luminosity and auto-correlation functions at two different redshifts provided that: (a) there exists a non-linear relation between quasar luminosity and halo mass, and the non-linearity increases with redshift; (b) the scatter in this relation is fairly small ($\sigma \lesssim 0.3 - 0.6$) and decreases significantly with redshift; (c) in accordance with this relation, luminous quasars ($\log_{10} L / \text{erg s}^{-1} \gtrsim 46.5$) are hosted by halos with mass $\log_{10} M / M_{\odot} \approx 13 - 13.5$ ($\log_{10} M / M_{\odot} \approx 12.5 - 13$) at $z = 4$ ($z = 2.5$); (d) the quasar duty cycle is a strong function of redshift, with a very low $\varepsilon_{\text{DC}} \approx 0.4\%$ at low- z that increases to $\varepsilon_{\text{DC}} \approx 30\%$ at $z = 4$. In the following, we further elaborate on this picture by investigating its implications for SMBH accretion and growth and by placing it in the context of previous work on the subject. We end the section by highlighting the main strengths and weaknesses of our analysis.

2.5.1 Implications for quasars’ physical properties

2.5.1.1 Black hole mass and accretion efficiency

The CLF posits an empirical relation between quasar luminosity and halo mass. However, many quasar population models (e.g., [Conroy & White 2013](#); [Veale et al. 2014](#); [Zhang et al. 2023b](#)) built this relation on physical grounds by relating the quasar luminosity to the mass of the central black hole, and this latter mass to the mass of either the host halo or the host galaxy/bulge. We can relate these two approaches by introducing the Eddington ratio η , which is defined by the following relation:

$$L = \zeta \eta M_{\text{BH}}, \quad (2.20)$$

where M_{BH} is the mass of the black hole, and $\zeta = 3.67 \times 10^4 L_{\odot}/M_{\odot}$ is a constant factor.

Then, we assume, e.g., that the mass of a black hole is determined solely by the mass of the host halo. In other words, we introduce a probability $P(M_{\text{BH}}|M)$ for the mass of the black hole given the halo mass. If we also write the “Eddington ratio distribution” $\text{ERDF}(\eta|M_{\text{BH}}, M)$ in terms of the other quantities considered, the conditional luminosity function reads:

$$\text{CLF}(L|M) = \int \frac{dM_{\text{BH}}}{\xi M_{\text{BH}}} \text{ERDF}\left(\frac{L}{\xi M_{\text{BH}}} \middle| M_{\text{BH}}, M\right) P(M_{\text{BH}}|M). \quad (2.21)$$

In this way, we have related the CLF – which is an empirically determined stochastic relationship between quasar luminosity and halo mass – to two other distribution functions (the ERDF and the black hole mass distribution) that have a clear physical meaning, being related to the physics of black hole accretion and growth.

In order to make this relationship explicit in our analysis, we can simply rewrite the quasar luminosity as the product of the Eddington ratio and the black hole mass (eq. 2.20). In this way, we can explicitly study how these two parameters – albeit completely degenerate – depend on the mass of the host halo, M , according to our model. The middle panel of Figure 2.6 illustrates this dependence. In this panel, we employ the $\text{CLF}(L|M)$ relation given by our model to write the probability distribution for the product of the Eddington ratio and the black hole mass-halo mass ratio, $P(\eta M_{\text{BH}}/M|M)$. Note that we divide the product $\eta M_{\text{BH}} = L/\zeta$ by the halo mass, M , because we expect black hole mass and halo mass to be approximately proportional based on local scaling relations (e.g., [Efstathiou & Rees 1988](#); [White et al. 2008](#); [Booth & Schaye 2010](#); [Marasco et al. 2021](#)) and because we can then work with a dimensionless quantity. Redshifts in the middle panel of Figure 2.6 are color-coded as in the top panel and in Figure 2.5. Median values and uncertainties for $P(\eta M_{\text{BH}}/M|M)$ are extracted by randomly sampling the

Markov chains for the posterior distributions, as well as the $L - M$ relation given by our model (see the caption for details).

While at $z = 2.5$ the median value of $\eta M_{\text{BH}}/M$ shows only a weak trend with halo mass, the situation is much different at $z = 4$, with the product $\eta M_{\text{BH}}/M$ strongly correlating with M . This can be achieved by assuming that either the black hole mass, the Eddington ratio, or both increase with halo mass. In other words, for very massive hosts black holes are either particularly massive (with a black hole mass-halo mass ratio higher than for lower-mass counterparts) or efficiently accreting (i.e., with large Eddington ratios). This trend is driven by the fact that the measured strong clustering at $z \approx 4$ requires that the most luminous quasar population is completely dominated by high-mass hosts.

It is also useful to cast these constraints in terms of galaxy stellar masses. This can be done by exploiting one of the parameterizations of the halo mass-stellar mass relation that are available in the literature. Here, we use the redshift-dependent halo mass-stellar mass relation from [Behroozi et al. \(2013\)](#) to rewrite $P(\eta M_{\text{BH}}/M|M)$ in terms of the galaxy mass M_* , i.e., $P(\eta M_{\text{BH}}/M_*|M)$. For simplicity, we neglect the scatter between stellar mass and halo mass in this conversion. The bottom panel of Figure 2.6 shows how the product of the Eddington ratio and the black hole mass-galaxy mass ratio varies as a function of halo mass. This is especially interesting in light of the fact that there is long-standing evidence in favor of a linear (or quasi-linear) relation between black hole and galaxy masses in the local universe (the so-called $M_{\text{BH}} - M_*$ relation, see e.g., [Magorrian et al. 1998](#); [Kormendy & Ho 2013](#); [Reines & Volonteri 2015](#)). In the same panel (Figure 2.6), we plot with a dashed red line the expectation for the product $\eta M_{\text{BH}}/M_*$ as a function of M , based on assuming the local $M_{\text{BH}} - M_*$ relation as measured by [Reines & Volonteri \(2015\)](#), converting galaxy masses to halo masses according to [Behroozi et al. \(2013\)](#), and setting a fixed Eddington ratio of $\eta = 1$. The scatter around this quantity (red shaded region) only considers the scatter in the $M_{\text{BH}} - M_*$ relation as quoted by [Reines & Volonteri \(2015\)](#). Due to the fact that the $M_{\text{BH}} - M_*$ relation is almost linear, the product between the Eddington ratio and the black hole mass-galaxy mass ratio is almost independent of halo mass, with an average constant value of $\eta M_{\text{BH}}/M_* \approx -3.5$.

Comparing this expectation based on local relations to the predictions of our model for $z = 2.5, 4$, we find that for low halo masses ($\log_{10} M/M_\odot \lesssim 12.5$) the predictions tend to agree within the uncertainties (see below for the caveat on extrapolating below $\log_{10} M/M_\odot \approx 12$). For larger masses, however, the difference between local predictions and our models becomes quite significant. At $z = 2.5$, there is a mild but significant positive trend of increasing $\eta M_{\text{BH}}/M_*$ with M . This trend becomes even steeper and tighter at $z = 4$, with the product $\eta M_{\text{BH}}/M_*$ ranging from ≈ -3.5 for $\log_{10} M/M_\odot \lesssim 12.5$ to ≈ -2 for $\log_{10} M/M_\odot \lesssim 13.5$. The trend can be

interpreted considering that the galaxy formation efficiency at $z = 2-4$ peaks at halo masses around $\log_{10} M_{\text{peak}}/M_{\odot} \approx 12 - 12.5$; as a consequence, the stellar mass-halo mass relation flattens at masses higher than $\log_{10} M_{\text{peak}}$ (Behroozi et al. 2013, 2019). Our model, on the other hand, does not predict any flattening in the quasar luminosity-halo mass relation for masses $\log_{10} M > \log_{10} M_{\text{peak}}$. Hence, the product $\eta M_{\text{BH}}/M_*$ becomes a steep function of halo mass for $\log_{10} M/M_{\odot} \gtrsim 12.5$. We believe the absence of a flattening in the quasar luminosity-halo mass relation at the high mass end is not simply a consequence of the chosen parametrization for the CLF. By experimenting with different parameterizations of the CLF, we found that any departure from a steep, non-linear relation between quasar luminosity and halo mass is incompatible with the measured value of the clustering (especially at $z = 4$), as a flattening of this relation would lower the characteristic halo mass of bright-quasar hosts. Therefore, we conclude that while galaxy growth appears to be quenched at the high mass end, even at high redshifts (e.g., Behroozi et al. 2019), this does not seem to be the case for black hole growth, as black holes in very massive halos need to be very massive and/or accreting efficiently. Indeed, observational evidence for an evolution of the $M_{\text{BH}} - M_*$ relation has been found repeatedly at high- z (implying over-massive black holes) together with signs of an increase in the median value of the ERDF with redshift (e.g., Vestergaard & Osmer 2009; Wu et al. 2022; Maiolino et al. 2024; Pacucci et al. 2023; Stone et al. 2023; however, see Li et al. 2022; Zhang et al. 2023c for a discussion of selection biases).

We conclude by noting that, in our analysis, the shape of the CLF is actually constrained by data only in a limited range of halo masses. For low halo masses, the corresponding quasar luminosities fall in a range where quasar clustering has never been measured and estimates for the QLF are not available (or highly uncertain). On the other hand, for very high halo masses (and hence very high quasar luminosities), quasars become so rare that estimates for the QLF are once again very uncertain. Moreover, if the quasars are luminous enough to be completely above the luminosity threshold for clustering, then the exact behavior of the luminosity as a function of halo mass becomes irrelevant. Therefore, in all panels of Figure 2.6, we show the regions in halo mass where our constraints on the CLF are based purely on extrapolations as dotted lines. This mainly concerns low halo masses ($\log_{10} M/M_{\odot} \lesssim 12.5$) at $z = 4$, and both very low ($\log_{10} M/M_{\odot} \lesssim 12$) and very high ($\log_{10} M/M_{\odot} \gtrsim 13.5$) masses at $z = 2.5$.

2.5.1.2 Quasar lifetime and the growth of high- z black holes

While the empirical relation between quasar luminosity and halo mass gives valuable information on the connection between black holes, their accretion efficiency, and their host halos/galaxies, another key piece of the puzzle

resides in the inferred values of the quasar duty cycle, ε_{DC} . This quantity is defined as the fraction of halos that are hosting bright quasars at any given time (see Sec. 2.2.1). If quasar activity is a stochastic process, however, the duty cycle is also equal to the total fraction of time in which a black hole is active as a bright quasar during the lifetime of an average host halo. In other words, the duty cycle is an average constraint on the total lifetime of a quasar, t_{Q} . Following [Martini & Weinberg \(2001\)](#) (see also [Martini 2004](#), [Haiman & Hui 2001](#)), we can simply assume that the characteristic lifetime of a halo is roughly equal to the age of the universe $t_{\text{U}}(z)$, and get an estimate for the quasar lifetime by writing $t_{\text{Q}} = t_{\text{U}}(z) \varepsilon_{\text{DC}}$. Using the values of ε_{DC} obtained in Sec. 2.4.2, we get $t_{\text{Q}} \approx 0.1 - 1$ Gyr at $z \approx 4$, and $t_{\text{Q}} \approx 10 - 15$ Myr at $z \approx 2.5$.

The quasar lifetime is one of the most fundamental quantities for understanding the role that SMBHs play in a cosmological context. According to the standard picture of SMBH growth (e.g., [Lynden-Bell 1969](#)), luminous quasars are powered by gas accretion onto a SMBH, and the rest mass energy of this material is divided between the small fraction ($\approx 10\%$) of radiation that we observe, and the growth of the black hole. In this picture, a phase of luminous quasar activity translates directly into a buildup of mass for the central SMBH. This provides a direct connection between the total luminosity emitted by quasars over cosmic time and the total mass residing in SMBHs in the local Universe (the so-called ‘‘Soltan argument’’, [Soltan 1982](#)). If the quasar lifetime is long compared to the Hubble timescale (i.e., the duty cycle is large), then the buildup of the total SMBH mass has taken place in only a small fraction of host galaxies that were active as bright quasars for a large fraction of their lifetimes. A short quasar lifetime, on the other hand, implies that most galaxies have undergone a brief bright quasar phase during their evolution history. The results of our analysis suggest that the latter scenario is valid at cosmic noon ($z \approx 1 - 3$), when most of the SMBH growth has taken place (e.g., [Shen et al. 2020](#)). The short quasar lifetime we find at $z \approx 2.5$ is, in fact, a direct consequence of the fact that quasar activity at that redshift takes place in relatively common halos with a broad distribution of host halo masses (top panel of Figure 2.6). Opposite conclusions can be obtained by considering our $z \approx 4$ results. In this case, we find that the large duty cycle translates into a quasar lifetime that is a large fraction of the Hubble time ($t_{\text{Q}} \approx 0.1 - 1$ Gyr). This implies that SMBH growth may be radically different in the young Universe as compared to cosmic noon. As suggested by the $z \approx 4$ QHMF in Figure 2.6 (top panel), quasar activity at high z takes place only in the few most massive halos that are present at that redshift, and hence these systems are active as bright quasars for a large fraction of cosmic time.

Estimating the quasar lifetime at high z is even more compelling in light of the fact that observations of very massive black holes powering luminous quasars at $z \gtrsim 5$ challenge our standard paradigm for black hole

formation and growth (Mazzucchelli et al. 2017b; Farina et al. 2022). In the standard picture, black holes follow an Eddington-limited exponential growth with a timescale that is equal to the “Salpeter time”, $t_S \approx 40$ Myr. At high z , models suggest that there is just enough cosmic time to grow the observed SMBH masses starting from massive seeds of $\approx 10^3 - 10^5 M_\odot$ (e.g., Inayoshi et al. 2020). For this reason, gauging the quasar lifetime is important because it offers an indirect probe of whether sustained accretion on SMBHs can take place at high z in the form of bright quasar activity. The long lifetime we infer at $z \approx 4$ is indeed consistent with this picture, providing an argument in support of models for Eddington-limited growth of high- z black holes. We can provide a rough estimate for this argument by considering as a characteristic host halo mass the median value of the $z \approx 4$ QHMF (Figure 2.6, top panel), $\log_{10} M_{\text{med}}/M_\odot \approx 13.3$. If we assume accretion at the Eddington rate ($\eta = 1$), we can translate this characteristic halo mass into a black hole mass using the relation between $\eta M_{\text{BH}}/M$ and M (middle panel of Figure 2.6): we get $\log_{10} M_{\text{BH}}/M_\odot \approx 9$ (Kollmeier et al. 2006). By assuming a seed mass of $10^2 M_\odot$ ($10^5 M_\odot$), we find that a total quasar lifetime of ≈ 600 Myr (≈ 350 Myr) is required to grow the black holes under the assumption of Eddington-limited accretion. This is in good agreement with the estimate for t_Q obtained above⁷. This simple argument shows how studying the demographic properties of quasars (such as their abundance and clustering) can place indirect constraints on the formation and evolution history of SMBHs.

Alternative estimates for the quasar lifetime can be obtained by a number of other methods (for an overview, see Martini 2004). Interestingly, results from studies of the quasar proximity effect at high z (Khrykin et al. 2016, 2019) paint a rather different picture than the one suggested here, finding values of the quasar lifetime that are several orders of magnitudes smaller (see also Davies et al. 2019; Eilers et al. 2021). Khrykin et al. (2021) compiled a set of HeII proximity zone measurements for $z \approx 3 - 4$ quasars, and inferred a log-normal quasar lifetime distribution with a mean of $t_Q \approx 0.2$ Myr and a standard deviation of ≈ 0.8 dex. It is important to note, however, that proximity zone measurements are sensitive only to a fraction of the past quasar lightcurve (up to ≈ 30 Myr for HeII). Clustering measurements, on the other hand, provide integral constraints on the total lightcurve emitted by quasars over the entire history of the Universe. In other words, they are

⁷This estimate assumes that black holes grow at the Eddington limit for their entire history. The demographic properties of quasars at the present time, however, do not constrain black hole growth on a timescale larger than the inferred value of t_Q . We can provide an alternative argument to link the quasar duty cycle to the growth of black hole mass by considering the characteristic luminosity of our quasar sample $L \gtrsim L_{\text{thr}}$, and convert that to an accreted black hole mass by assuming a radiative efficiency of $\approx 10\%$ and a total lifetime t_Q . We get $\log_{10} M_{\text{BH}}/M_\odot \approx 8.7 - 9.7$ for $t_Q = 0.1 - 1$ Gyr, which again points to the fact that black holes can grow out to very high masses based on our inferred duty cycle.

only sensitive to the zeroth moment of the quasar lightcurve distribution (i.e., the aggregate probability of the lightcurve). The discrepancy between lifetime estimates for proximity zone sizes and clustering measurements, then, may suggest that quasar lightcurves exhibit non-trivial variations on timescales close to the ones probed by proximity zones. With this respect, exploring the full probability distribution associated with quasar lightcurves in the context of our quasar demographic model would provide a way to connect these very different observational probes of quasar activity in a single consistent picture. We will investigate this point in future work.

2.5.2 Comparison with previous work

The model presented in this work builds on a long-standing tradition of interpreting quasar observables via population modeling, i.e., by linking quasars to the well-known population of halos (or sometimes galaxies) according to some empirical/phenomenological prescriptions. Explaining the observed relative abundance of quasars at different luminosities (i.e., the QLF) within such frameworks has been achieved many times, with a large variety of empirical models and physical prescriptions employed (e.g., [Efstathiou & Rees 1988](#); [Wyithe & Loeb 2003](#); [Croton 2009](#); [Conroy & White 2013](#); [Fanidakis et al. 2013](#); [Veale et al. 2014](#); [Caplar et al. 2015](#); [Weigel et al. 2017](#); [Ren & Trenti 2021](#); [Zhang et al. 2023b](#)). The bottom line is that the QLF is pretty straightforward to model starting from the hierarchical growth of structures predicted in the Λ CDM framework. On the other hand, the QLF alone does not place tight constraints on key properties of quasars such as their black hole mass, accretion rate, lifetime, and host halo mass, not even in the context of redshift-dependent models (e.g., [Wyithe & Padmanabhan 2006](#); [Wyithe & Loeb 2009](#); [Veale et al. 2014](#)). Indeed, our analysis in Sec. 2.4.1 (Fig. 2.4) suggests that a very wide variety of model parameters can be in good agreement with the QLF. As shown by [Veale et al. \(2014\)](#), alternative parametrizations would fare nearly equally well at all redshifts. The large uncertainties on the actual shape and normalization of the QLF that are due to the significant systematics involved in these measurements ([Kulkarni et al. 2019](#)) exacerbate this issue, especially at high redshift.

For this reason, considering the independent constraints coming from quasar clustering is extremely useful, as they provide constraints on the masses of the halos that are capable of hosting quasars. Reproducing the clustering of low-redshift ($z \lesssim 2.5$) quasars has been shown to be possible both in empirical models (e.g., [Kauffmann & Haehnelt 2002](#); [Hopkins et al. 2007b](#); [Croton 2009](#); [Shankar et al. 2010a](#); [Conroy & White 2013](#); [Aversa et al. 2015](#); [Shankar et al. 2020](#)), semi-analytic models (e.g., [Bonoli et al. 2009](#); [Fanidakis et al. 2013](#); [Oogi et al. 2016](#)) and cosmological hydrodynamical simulations (e.g., [DeGraf & Sijacki 2017](#)). All of these studies, however, show a significant tension with the clustering measurements at redshift $z \gtrsim 3$.




Figure 2.6: *Top:* Quasar-host mass function (QHMF) at $z = 2.5$ (solid purple line) and $z = 4$ (green), according to our model. These functions and their respective uncertainties are the median and the 16th and 84th percentiles of the distributions obtained by randomly subsampling the Markov chains of the posteriors shown in Figure 2.5. The halo mass functions (HMFs) for both redshifts are shown with semi-transparent lines, whereas the dashed-dotted lines indicate the median values of the QHMF distributions. In all panels, regions in the halo mass spectrum where the behavior of the conditional luminosity function (CLF) is purely extrapolated and not explicitly constrained by data are shown with dotted lines. *Middle:* Same as the top panel, but showing the dependence on halo mass of the product between the Eddington ratio (η) and the black hole-halo mass ratio (M_{BH}/M). In this case, there are two sources of scatter: the uncertainty on the model given by the posterior distribution and the intrinsic scatter coming from the σ parameter in the CLF. We plot the former with a darker shading, whereas the total contribution of the two sources of scatter is shown with a lighter shading. *Bottom:* Same as the middle panel, but showing the quantity $\eta M_{\text{BH}}/M_*$ instead (with M_* being the galaxy mass). The relation between halo mass and galaxy mass is taken from Behroozi et al. (2013). The red dashed line shows the prediction for $\eta M_{\text{BH}}/M_*$ assuming the Reines & Volonteri (2015) relation between black hole and galaxy masses (with the shading showing the scatter in the relation), and setting $\eta = 1$.

The implications for the strong clustering measured by S07 at $z \approx 3 - 4$, in particular, have been discussed by White et al. 2008, Wyithe & Loeb 2009, and Shankar et al. 2010b. White et al. (2008) assume that the quasar luminosity-halo mass relation is linear, and find that the total number density of bright quasars can be reconciled with the linear bias measured by S07 only if the scatter about this linear relation is very small ($\sigma \lesssim 0.3$ dex). They also find that their conclusions are strongly dependent on the specific functional form assumed for the linear bias-halo mass relation $b(M)$: the Mo & White bias (Mo & White 1996; Jing 1998) is found to be marginally compatible with data, whereas the Sheth, Mo, & Tormen one (Sheth et al. 2001) is inconsistent with the measured bias at the $\approx 2\sigma$ level. Adopting the Sheth, Mo, & Tormen functional form of $b(M)$ after showing that it is a better fit to N-body simulations, Shankar et al. (2010b) interpret the S07 data in the context of an evolutionary model for supermassive black holes, and they strengthen the conclusion that there is tension between the measured bias and the theoretical predictions at $z = 4$. Similar results are found by Wyithe & Loeb (2009), who advocate for a contribution of a merger-driven bias to the $z = 4$ clustering (see also Bonoli et al. 2010; Cen & Safarzadeh 2015 for a discussion of the impact of assembly bias on quasar clustering).

Our work shares some similarities with the three studies mentioned above: we also assume a direct relation between quasar luminosity and halo mass and use the quasar clustering data to infer the specific shape of this relation. Key conclusions of our analysis can also be found in these former attempts to explain the S07 observations: in Sec. 2.4.2, we find that bright quasars need to be hosted by very massive ($\log_{10} M/M_{\odot} \gtrsim 13$) halos, and, as a consequence, the quasar duty cycle is a significant fraction of unity ($\epsilon_{\text{DC}} \approx 10 - 100\%$). In agreement with White et al. (2008) and Shankar et al. (2010b), we conclude that a relatively small scatter ($\sigma \lesssim 0.3$; Table 2.3) in the quasar luminosity-halo mass relation is necessary to explain the S07 measurement. As also done by Wyithe & Loeb (2009), we adopt a more flexible parametrization of this relation by assuming that it can be non-linear, and find that a steep slope ($\gamma \gtrsim 2$) achieves a much better fit to the data.

The major novelty that our work brings to the understanding of this problem, however, does not reside in the interpretation of the results, but rather in the framework we use to build our model. As explained in Sec. 2.2, we extract the correlation function and the relative abundance of quasars directly from extremely large-volume cosmological N-body simulations, using a novel method to quickly compute the quasar auto-correlation function for any quasar-host mass distribution. In this way, we can directly compare our predictions for the quasar projected correlation function with the S07 observational data. Our model – being based on N-body simulations – naturally accounts for the non-linear contributions to quasar clustering that are essential to interpret the S07 clustering measurements correctly,

especially at scales $r \lesssim 10 - 15$ cMpc. Using this approach, we thus achieve a much more solid data-model comparison, as we do not have to resort to the notion of large-scale linear bias, which is significantly uncertain for strongly biased systems (Diemer 2018) and which discards the information about the shape and the physical features that lie in the S07 data points.

Indeed, according to the statistical analysis performed in Sec. 2.4.1, we find our model can match the data with a satisfactory level of accuracy (i.e., reduced chi-squared ≈ 1), suggesting that, despite being rather extreme, the S07 data can be explained in the context of the standard framework in which clustering of dark matter halos is solely dictated by their mass, without the need to invoke any contributions from assembly/merger bias. We note that we cannot exclude, of course, that such a contribution is present. If that is the case, it would imply that the mass function of $z = 4$ bright-quasar hosts may be somewhat less skewed towards very large halo masses ($\log_{10} M/M_{\odot} \gtrsim 13$). We leave an assessment of the role that merger bias plays in cosmological simulations to future work.

Finally, we note that several measurements of the characteristic mass of quasar-hosting halos at $z = 2 - 4$ are available in the literature. They employ quasar-quasar (S07; E15; Timlin et al. 2018) and quasar-galaxy (Trainor & Steidel 2012; Ikeda et al. 2015; García-Vergara et al. 2017; He et al. 2018; García-Vergara et al. 2019) clustering, as well as gas kinematics in the circumgalactic medium (CGM) of quasar-hosting galaxies (Fossati et al. 2021; de Beer et al. 2023). While the host halo masses predicted by these studies vary significantly, in the present work we have decided to focus on the S07 and E15 measurements of SDSS/BOSS quasars only, because these quasar samples are entirely spectroscopic and thus free from any low-redshift contaminants. However, the same analysis described in this paper could also be performed by taking into account the other clustering measurements mentioned above.

2.5.3 Caveats and final remarks

As shown schematically in Fig. 2.1, the results presented in this work depend on two key ingredients: the choice of the CLF, and the extraction of the halo mass function and the halo (cross-)correlation functions from cosmological N-body simulations. In the following, we will discuss the strengths and weaknesses of our method by considering these components in turn. Let us start with the latter: there are multiple sources of uncertainty in the final estimates we obtain for the halo mass function and the halo cross-correlation functions. First, despite the fact that the box sizes of the simulations employed here are among the largest ever run (Angulo & Hahn 2022), halos are so rare at the very massive end ($4 - 7\sigma$ peaks in the density field) that the results of simulations at these masses suffer from significant noise. In order to circumvent this issue – and to extrapolate the results of simulations

to the highest mass possible – we used analytical functions to fit the data extracted from the simulations (Sec. 2.2.2). These analytical fits, however, are not perfectly accurate and contribute some systematic errors to our final model predictions.

Nonetheless, we believe that these sources of error can be neglected in our data-model comparison (Sec. 2.4). This is because – as also noted in Sec. 2.2.2 – the observables we are trying to reproduce, i.e., the QLF and the projected correlation function, suffer from significant statistical and systematic uncertainties (as high as $\approx 100\%$ at $z = 4$ and $\approx 30\%$ at $z = 2.5$). In Sec. 2.2.2.1–2.2.2.2 and Appendix 2.B, we assess how well our fitting functions reproduce simulations, and show that their relative accuracy is generally $\lesssim 5 - 10\%$ for both the halo mass function and the cross-correlation functions. For very small ($r \lesssim 5 \text{ cMpc}$) and very large ($r \gtrsim 100 \text{ cMpc}$) scales, measuring the cross-correlation functions in simulations is particularly challenging, especially at the high mass end. The small-scale behavior is highly affected by halo exclusion effects, whereas at large scales the finite size of the simulated boxes reduces the number of pairs, and the baryon acoustic oscillation (BAO) peak makes the shape of correlation functions difficult to capture with our coarse radial bins. As a consequence, our fitting functions are also subject to larger errors at both of these scales. However, these errors do not have a significant impact on our final results, as observational data are also very uncertain at the same scales; for this very same reason, we have excluded the S07 measurements at very large scales ($r > 100 \text{ cMpc}$) from our analysis (see Sec. 2.3.1).

As for the extrapolation of cross-correlations functions to masses higher than the ones that we can probe with our simulations, we have argued that such extrapolation is well motivated by considering the case of $z = 2.5$ in Appendix 2.C. Furthermore, we note that the accuracy of this extrapolation does not have a significant impact on our results: this can be determined by looking at the QHMFs in Figure 2.6 (top panel). At both redshifts, the quasars hosted by halos whose mass is not well represented in simulations are only a small fraction of the total number of quasars (e.g., $\lesssim 5\%$ at $z = 4$). This implies that their actual contribution to the quasar auto-correlation function is negligible compared to the uncertainty in the data.

Other possible sources of uncertainty in our model predictions that we have not discussed yet are the cosmology assumed in the simulations and the exclusion of sub-halos in the creation of the halo cross-correlation functions. Cosmological parameters such as σ_8 and Ω_m are predicted to have a significant effect on the collapse of structures in the standard ΛCDM model, and consequently on the spatial distribution of very massive halos at all redshifts. Studying the impact of these parameters on our final predictions for the clustering of quasars is beyond the scope of this work. Given the current large relative uncertainty on the data, however, we believe that including

variations of the cosmological parameters in our inference procedure would have little effect on our final results.

The exclusion of sub-halos is motivated by the fact that we are only considering clustering measurements at medium-large scales, which are not affected by the distribution of quasars inside a single halo – the so-called one-halo term in HOD models (e.g., Cooray & Sheth 2002). In principle, including the contribution of sub-haloes may boost the large-scale clustering too, as having multiple quasars living in the same dark matter halo implies a larger number of large-scale pairs. In practice, however, sub-haloes are less massive than centrals, and thus they do not tend to host very bright quasars according to the CLF found in Sec. 2.4. We have included sub-haloes in some test runs and verified that large-scale clustering changes only at the percent level, and significant differences are only present at $r \lesssim 1 - 2$ cMpc even for the most massive halos. On top of that, we note that all of the effects discussed here go in the direction of an enhancement of the predicted clustering, and do not affect the main conclusion of this paper, i.e., that the very strong clustering measured at $z = 4$ can be reproduced with standard assumptions of bright quasars inhabiting massive halos.

In this work, we have assumed one, very simple parametrization for the CLF. We believe that this simple framework is a strength of our model, as it provides very clear physical insight into the formation of quasars and their connection with the hierarchical growth of structures in the context of a Λ CDM universe. On the other hand, we have tested this basic parametrization on a relatively small amount of (very uncertain) data. We have done this on purpose: the main focus of this paper is on reproducing quasar clustering at $z = 4$, and given the quality of the data we have at the present moment, a more sophisticated choice for the CLF would likely have been too flexible to be constrained. However, it is possible that extending our model to a larger/higher signal-to-noise ratio (S/N) dataset, e.g. at $z = 0 - 2$, would be possible only with a more sophisticated parametrization for the CLF. We leave a thorough examination of different prescriptions for the CLF for future work. In particular, we plan to apply our model to the multiple measurements of quasar clustering available at low- z (e.g., Porciani et al. 2004; Croom et al. 2005; Ross et al. 2009) as well as to the analyses of the dependence of clustering on luminosity at the same redshifts (e.g., Porciani & Norberg 2006; Shen et al. 2009; Eftekharzadeh et al. 2015).

2.6 Summary

We have introduced a novel framework that makes use of multiple cosmological N-body simulations to efficiently predict quasar observables such as the quasar luminosity function (QLF) and the quasar auto-correlation function. The halo mass function and the cross-correlation functions of halos with

different masses are extracted from the dark-matter-only (DMO) versions of the FLAMINGO simulations and used to inform analytical fitting functions. These form the backbone of the model, which is then completed by the choice of a conditional luminosity function (CLF) that links halo masses to quasar luminosities. With these ingredients, we are able to predict the clustering and the luminosity function of quasars, as well as other key properties such as the mass distribution of quasar-hosting halos and the quasar duty cycle (Figure 2.1).

We focus our analysis on the extremely strong clustering measured by Shen et al. (2007) at $z \approx 4$, with the goal of determining whether we can reproduce this measurement in the context of our model. We use a simple parametrization for the CLF, assuming a power-law dependence of quasar luminosity on halo mass ($L \propto M^\gamma$) with a log-normal scatter σ . We fit the $z = 4$ QLF and projected correlation function both independently and jointly, in order to gain insight into the best-fitting parameters for each of the cases considered. We then turn our attention to lower- z data, and show that our model can also match the measurements of the same quantities at $z \approx 2.5$ (Ross et al. 2009; Eftekharzadeh et al. 2015), albeit with significantly different values of the model parameters.

We summarise here the main findings of the analysis described above:

- Quasar clustering and abundance measurements at $z \approx 4$ require quasars to reside in the most massive halos at that redshift, with a characteristic mass of $\log_{10} M/M_\odot \gtrsim 13$ (Figure 2.3). This implies that the relation between quasar luminosity and halo mass ($L - M$) is highly non-linear ($\gamma \gtrsim 2$) with a very small amount of scatter ($\sigma \lesssim 0.3$ dex).
- Many different combinations of model parameters can achieve a good fit to the measured QLF at $z \approx 4$ (Figure 2.4). This is because very different empirical prescriptions for the quasar luminosity-halo mass relation (e.g., large scatter and shallow slope or vice-versa) are able to map the exponentially declining end of the halo mass function into the shallower bright end of the QLF. However, the only set of parameters which is also compatible with clustering measurements is the one mentioned above (i.e, a highly non-linear $L - M$ relation with very small scatter), as an increase in the scatter would lower the characteristic mass of quasar-hosting halos, and thus decrease the clustering predicted by our model.
- In order to match the total number density of bright $z \approx 4$ quasars in models in which quasars reside in sufficiently high halo masses to reproduce the observed clustering, the active fraction of quasars (f_{on}) has to be close to unity. This implies that high- z quasars shine for a large fraction of the Hubble time, with a duty cycle in the range $\varepsilon_{\text{DC}} = 10 - 60\%$. In turn, this duty cycle results in a large total quasar

lifetime $t_Q \approx 10^8 - 10^9$ yr, consistent with the standard picture of black hole growth in the young universe.

- The steep $z \approx 4$ relation between quasar luminosity and halo mass contrasts with the well-known prediction of a flattening in the stellar mass-halo mass relation at high mass at every epoch (e.g. [Behroozi et al. 2013](#)). This implies that in very massive high- z halos – while the star formation may have been quenched already – the supermassive black hole at the center of the galaxy needs to be either over-massive and/or highly accreting. This may have an impact on the shape and normalization of the black hole mass-galaxy mass relation at high redshift (see e.g., [Maiolino et al. 2024](#); [Stone et al. 2023](#)).
- Furthermore, the extremely small scatter ($\sigma \approx 0.1 - 0.3$ dex) inferred for the $L - M$ relation at $z \approx 4$ points to some physical processes enforcing a tight relationship between quasars and their dark matter halo hosts. In other words, the relation between black hole mass and stellar and/or halo mass, together with the distribution of Eddington ratios, all conspire to yield a remarkably low scatter.
- The clustering and relative abundance of quasars at lower redshift ($z \approx 2.5$) can be explained by the same parametric relation between quasar luminosity and halo mass. However, the parameters describing this relation show a significant evolution with redshift (Figure 2.5): the slope of the $L - M$ is significantly shallower ($\gamma \approx 1.15$) than at $z \approx 4$, and the scatter larger ($\sigma \approx 0.5$ dex).
- Overall, our comparison between $z \approx 2.5$ and $z \approx 4$ reveals two radically different pictures in terms of the connection between quasars and their host halo population (Figure 2.6). High- z ($z \approx 4$) quasars are hosted by very massive halos, with a very large occupation fraction (i.e., a large fraction of these halos host bright quasars at any given time). At lower redshift ($z \approx 2.5$), instead, quasars reside in halos with a broad range of masses, with the bulk of the population being characterized by relatively common, $\log_{10} M/M_\odot \approx 12.5$ mass halos. As a consequence, only a small fraction of low- z quasars are actively shining at any given moment, with a quasar duty cycle of $\varepsilon_{\text{DC}} \approx 0.5\%$. These conclusions are consistent with the standard picture of “cosmic downsizing” of quasars and AGN (e.g., [Merloni 2004](#); [Scannapieco et al. 2005](#); [Fanidakis et al. 2012](#)), as the bulk of the quasar population is hosted by progressively smaller halos as redshift decreases.

The framework presented here can be readily applied to interpret quasar clustering measurements at all redshifts. In particular, focusing on very high redshift is especially interesting in light of the fact that the large-scale environment of very bright quasars has been proven hard to pinpoint in

the early universe (e.g., [Fan et al. 2023](#)). For example, [Arita et al. \(2023\)](#) recently measured the quasar auto-correlation function at $z \approx 6$ for the first time, finding results broadly consistent with the very strong clustering measured at $z \approx 4$. On top of that, several JWST programs such as ASPIRE ([Wang et al. 2023](#)) and EIGER ([Kashino et al. 2023](#); [Eilers et al. 2023](#)) are starting to deliver measurements of quasar clustering by probing the distribution of line emitters around bright, $z \approx 6-7$ quasars. Connecting the framework presented here to the upcoming quasar-galaxy cross-correlation measurements from JWST will offer a clear and comprehensive picture of the large-scale environments in which the first quasars formed (see also [Costa 2024](#) for an alternative approach).

As suggested by the results obtained in this work, interpreting quasar properties within a consistent framework that takes into account both their demographics and their spatial distribution can give great insight into the relationship between the hierarchical growth of structures in the Universe and the evolution of supermassive black holes over cosmic time.

2.A Appendix: Obtaining the quasar auto-correlation from the halo cross-correlation functions

Let us consider a stochastic process $\mathcal{N}^{(q)}$, describing – in our case of interest – the spatial distribution of quasars. This distribution is discrete: following [Peebles \(1980\)](#), we divide the volume of interest into infinitesimal elements δV_i , and – given the average quasar density $\bar{n}_q$ – we can write the probability of having a quasar in the volume element δV_1 as:

$$\delta P_1 = \langle \mathcal{N}_1^{(q)} \rangle = \bar{n}_q \delta V_1. \quad (2.22)$$

Similarly, we define the two-point correlation function, $\xi(r_{12}) \equiv \xi_{12}$, via the probability of having a quasar in the volume element δV_1 and another one in the volume element δV_2 :

$$\delta P_{12} = \langle \mathcal{N}_1^{(q)} \mathcal{N}_2^{(q)} \rangle \equiv \bar{n}_q^2 \delta V_1 \delta V_2 (1 + \xi_{12}). \quad (2.23)$$

We introduce now the continuous number density field $n^{(q)}(\mathbf{x})$, which we define via the expression:

$$\delta P_{12} = \langle \mathcal{N}_1^{(q)} \mathcal{N}_2^{(q)} \rangle \equiv \langle n^{(q)}(\mathbf{x}_1) n^{(q)}(\mathbf{x}_2) \rangle \delta V_1 \delta V_2, \quad (2.24)$$

and write this equation in terms of the density contrast field $\delta^{(q)}$ – defined as $n^{(q)}(\mathbf{x}) = \bar{n}_q (1 + \delta^{(q)}(\mathbf{x}))$:

$$\delta P_{12} = \bar{n}_q^2 \delta V_1 \delta V_2 \left(1 + \langle \delta^{(q)}(\mathbf{x}_1) \delta^{(q)}(\mathbf{x}_2) \rangle \right). \quad (2.25)$$

By comparing this to eq. 2.23, we find that the correlation function can also be expressed as:

$$\xi_{12} = \langle \delta^{(q)}(\mathbf{x}_1) \delta^{(q)}(\mathbf{x}_2) \rangle \equiv \langle \delta_1^{(q)} \delta_2^{(q)} \rangle. \quad (2.26)$$

We want to split the different contributions of quasars to the density field $n^{(q)}$ based on the mass of their host halos. We introduce a set of continuous fields $\{n^{(h)}(M_k)\}$, which represent the distributions of halos for different mass bins centered on M_k and with a width ΔM .

We now make the key hypothesis that the distribution of quasars with a host halo mass in the range $[M_k - \Delta M/2, M_k + \Delta M/2]$ is an unbiased tracer of the underlying distribution of halos, $n^{(h)}(M_k)$. In other words, the quasars at a given host halo mass are just undersampling the distribution of halos, and they are thus described by the same stochastic process. This is the case if the presence of a quasar depends solely on the mass of its host. Thus, we can write the quasar distribution, $n^{(q)}$ in terms of the distributions of halos with different masses by simply weighing them by the relative number of quasars at those masses:

$$n^{(q)} = \sum_k p_k n^{(h)}(M_k), \quad (2.27)$$

where p_k represents the probability that a quasar has a host-halo mass in the bin M_k . Using the “quasar-host mass function” (QHMF) introduced in Sec. 2.2.1 (eq. 2.6), we can express this probability as ($\bar{n}_{q,k}$ is the average number density of quasars in the host mass bin M_k):

$$p_k = \frac{\bar{n}_{q,k}}{\bar{n}_q} = \frac{n_{\text{QHMF}}(M_k) \Delta M}{\int_0^\infty n_{\text{QHMF}}(M) dM}. \quad (2.28)$$

Introducing the overdensity definitions for the distributions at different masses, $n^{(h)}(M_k) = \bar{n}_{q,k} (1 + \delta^{(h)}(M_k))$ we can write:

$$\begin{aligned} \langle n_1^{(q)} n_2^{(q)} \rangle &= \sum_j \sum_k p_j p_k \bar{n}_q^2 \left(1 + \langle \delta_1^{(h)}(M_j) \delta_2^{(h)}(M_k) \rangle \right) = \\ &= \bar{n}_q^2 \left(1 + \sum_j \sum_k p_j p_k \langle \delta_1^{(h)}(M_j) \delta_2^{(h)}(M_k) \rangle \right), \end{aligned} \quad (2.29)$$

where we have made use of the fact that $\sum_k p_k = 1$. Introducing the cross-correlation functions for halos of different masses, $\xi_{12}^{(h)}(M_j, M_k) = \langle \delta_1^{(h)}(M_j) \delta_2^{(h)}(M_k) \rangle$, we can express the quasar auto-correlation as:

$$\xi_{12} = \langle \delta_1^{(q)} \delta_2^{(q)} \rangle = \frac{\langle n_1^{(q)} n_2^{(q)} \rangle}{\bar{n}_q^2} - 1 = \sum_{j,k} p_j p_k \xi_{12}^{(h)}(M_j, M_k). \quad (2.30)$$

This proves eq. 2.7, which relates the quasar auto-correlation function $\xi(r_{12}) \equiv \xi_{12}$ to the cross-correlation functions of halos with different masses, $\xi_h(M_j, M_k; r_{12}) \equiv \xi_{12}^{(h)}(M_j, M_k)$.

2.B Appendix: Fitting the cross-correlation terms from simulations

In this Section, we provide details on the fitting we perform to the values of the halo cross-correlation functions, $\xi_h(M_j, M_k; r)$, extracted from simulations. As described in Sec. 2.2.2 and 2.2.2.2, we extract these values from two simulations with box sizes equal to $L = 2800$ cMpc and $L = 5600$ cMpc, respectively. For each simulation, we consider only halos in a specific range of masses, so that all the mass bins considered are populated by a sufficient number of well-resolved halos. In particular, we set the following ranges (Sec. 2.2.2): $\log_{10} M/M_\odot = 11.5 - 13.0$ for $L = 2800$ cMpc, and $\log_{10} M/M_\odot = 12.5 - 13.5$ for $L = 5600$ cMpc. We choose a bin width of 0.25 in $\log_{10} M$, so that we have 6×6 cross-correlation terms for $L = 2800$ cMpc, and 4×4 cross-correlation terms for $L = 5600$ cMpc. Note that the masses M_j and M_k in the expression $\xi_h(M_j, M_k; r)$ do not refer to the center of their respective bins, but rather to the median value of the halo mass function in those bins.

Our goal is then to find a single analytical description of these cross-correlation functions that can represent the two simulations simultaneously. In order to do that, we first divide all the cross-correlation functions $\xi_h(M_j, M_k; r)$ by a reference correlation $\xi_{\text{ref}}(r)$; in formulae, we define $\rho(M_j, M_k; r)$ to be:

$$\rho(M_j, M_k; r) = \xi_h(M_j, M_k; r) / \xi_{\text{ref}}(r). \quad (2.31)$$

In this way, we hope that $\rho(M_j, M_k; r)$ will be only marginally dependent on the scale r . We set $\xi_{\text{ref}}(r) \equiv \xi_h(\tilde{M}, \tilde{M}; r)$, with $\tilde{M}$ representing the $\log_{10} M = 12.5 - 12.75$ bin. This choice is arbitrary, but it is made to ensure that the mass bin sits in the overlap between the mass ranges of the two different simulations we use. We also attempt to minimize any dependences of the cross-correlation functions on cosmology and redshift by expressing all the masses in terms of peak heights $\nu(M)$ (see also Sec. 2.2.2.1).

Finally, we fit a 3-d polynomial $\rho_{\text{fit}}(\nu_j, \nu_k, r)$ to the values extracted from the simulations. We empirically find that a second-degree polynomial in mass and third-degree in the radial dimension fits the data points well enough and at the same time attains a smooth behavior with respect to all three variables (i.e., we prevent overfitting). The errors on the data points are assigned based on the Poisson statistics of the pair counts (eq. 2.16). As also done in Sec. 2.2.2.1, we weigh the errors associated with the two

simulations differently in order to achieve a better fit. In particular, we double the values of the Poisson errors for the $L = 2800$ cMpc simulation, and we halve the ones associated with the $L = 5600$ cMpc box.

Figures 2.7–2.8 show the results of the fitting for the two redshifts considered in this work: $z = 4$ (Fig. 2.7) and $z = 2.5$ (Fig. 2.8). The first row of each plot displays the resulting fitting function $\rho_{\text{fit}}(\nu(M_j), \nu(M_k), r)$. Each panel in this row shows the values of $\rho_{\text{fit}}(\nu(M_j), \nu(M_k), \bar{r})$ as a function of the two masses M_j and M_k at a different scale $\bar{r}$. The second and third rows show the relative differences $(\rho/\rho_{\text{fit}} - 1)$ between our fit and the two simulations considered. The mass ranges that are selected in each simulation are shown as grey boxes in the 2-d mass planes. According to these figures, our simple analytical framework can describe the behavior of cross-correlation functions in a wide mass range with a good degree of accuracy ($\lesssim 5 - 10\%$). Notable exceptions to this can be found for very high masses ($\log_{10} M/M_\odot \gtrsim 13.2 - 13.3$) and very large scales ($r \gtrsim 100$ cMpc). However, these are both expected, as both at large masses and large scales correlation functions are difficult to measure in simulations. Further discussion of this can be found in Sec. 2.2.2.2 and Sec. 2.5.3. Similar conclusions on the quality of our fits can be drawn by looking at the right panels of Figure 2.2 and Figure 2.9, where predictions from our fitting functions are compared to auto-correlation functions extracted from the simulations.

2.C Appendix: Halo mass function and correlation functions for redshift $z = 2.5$

In the main text (Sec. 2.2.2.1–2.2.2.2), we discussed our predictions for the halo mass function and the halo (cross-)correlation functions at $z = 4$. We show here the same results for the other redshift that we consider in our analysis, $z = 2.5$. In Figure 2.9 (left panel), we show the halo mass function extracted from the two simulations ($L = 2800$ cMpc and $L = 5600$ cMpc in teal and red, respectively), as well as our fitting function (eq. 2.14–2.15, gray line). The best-fitting parameter values for $z = 2.5$ are: $A = 0.464$; $a = 3.43$; $b = 0.847$; $c = 1.31$. Note that for the fitting we employ the same mass ranges as we used for $z = 4$ (Table 2.1). This choice is clearly sub-optimal, as halos are much more abundant at lower- z , and therefore mass bins with $\log_{10} M/M_\odot > 13.5$ are well populated even for the smallest box considered ($L = 2800$ cMpc). However, we choose to not take into account masses larger than $\log_{10} M/M_\odot > 13.5$ for the fitting so that we can benchmark how well our fitting function fares if extrapolated to masses larger than this limit. In this way, we can test whether our fitting framework is valid to interpret the behavior of the halo mass function up to masses higher than the ones we can simulate. As shown in Fig. 2.C (left panel), the trend of the halo mass

function at large masses is well described by the extrapolation of our fitting up to $\log_{10} M/M_{\odot} \gtrsim 14$; at higher masses, halos become very rare even at $z = 2.5$, and the halo mass function becomes quite noisy and its behavior highly uncertain.

The right panel of Figure 2.9 shows the halo auto-correlation functions for different mass bins obtained both from simulations (colored points) and from our fitting function (gray lines; see Appendix 2.B). We use 8 mass bins ranging from $\log_{10} M/M_{\odot} = 11.5$ to $\log_{10} M/M_{\odot} = 13.5$ and with a width of 0.25 dex. Lower mass bins correspond in Fig. 2.C to lower values of the correlation functions, and vice-versa. Once again, the mass ranges employed for our fitting are the same for both $z = 2.5$ and $z = 4$, and do not go higher than $\log_{10} M/M_{\odot} = 13.5$. This gives us the possibility of testing how the extrapolation of $\xi_{h,\text{fit}}(M_j, M_k; r)$ fares at larger masses. As explained in Sec. 2.2.2.2, ensuring that our theoretical framework can be extended to very high masses ($\log_{10} M/M_{\odot} \approx 14$) is quite important, as – especially at $z = 4$ – a significant fraction of quasars are hosted by this population of massive halos that is not well represented in our simulations.

Extrapolations from our fits are shown in Fig. 2.9 (right panel) with dashed lines. We also extract halo auto-correlation functions from the $L = 5600$ cMpc box for the mass bin $\log_{10} M/M_{\odot} = 13.5 - 13.75$; we show these values with golden crosses in Fig. 2.9 (right panel). We see that the extrapolation agrees with simulations at the same level as the points that are used for fitting, with the only exceptions being very small ($r \lesssim 5$ cMpc) and very large ($r \gtrsim 100$ cMpc) scales. Further discussion on the implications of these results can be found in Sec. 2.5.3.

Acknowledgements

We are grateful to the FLAMINGO team for making their dark matter only simulations available. We acknowledge helpful conversations with the ENIGMA group at UC Santa Barbara and Leiden University. EP is grateful to Roi Kugel for helpful discussion, and to Molly Wolfson, Shane Bechtel, and Silvia Onorato for comments on an early draft of this paper. JFH and EP acknowledge support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 885301). This work is partly supported by funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860744 (BiD4BESt). This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1,




Figure 2.7: Results for the fitting of the cross-correlation terms $\rho(M_1, M_2, r)$ (see Appendix 2.B for definitions). The top row shows the fitting function $\rho_{\text{fit}}(M_1, M_2, r)$ as a function of the two masses M_1 and M_2 for different values of the distance r . The second and third rows show the relative difference between the fits and the values measured from the simulations ($L = 2800$ cMpc and $L = 5600$ cMpc, respectively). Mass ranges that are adopted for the fitting in each simulation are highlighted by the gray boxes in each mass-mass plane.




Figure 2.8: Same as Figure 2.7, but for redshift $z = 2.5$.




Figure 2.9: Same as Fig. 2.2 but for the snapshots at redshift $z = 2.5$. Golden crosses in the right panel represent the auto-correlation functions measured in the mass bin $\log_{10} M / M_{\odot} = 13.5 - 13.75$, in the $L = 5600 \text{ cMpc}$ simulation. This is used as a benchmark to assess how well our fits (dashed grey lines) can be extrapolated to higher masses.

ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

