

Tracing black hole and galaxy co-evolution in the ROMULUS simulations

Angelo Ricarte,^{1★} Michael Tremmel,^{1,2} Priyamvada Natarajan^{1,2} and Thomas Quinn³

¹*Department of Astronomy, Yale University, 52 Hillhouse Avenue, New Haven, CT 06511, USA*

²*Department of Physics, Yale University, PO Box 208121, New Haven, CT 06520, USA*

³*Department of Astronomy, University of Washington, PO Box 351580, Seattle, WA 98195, USA*

Accepted 2019 July 31. Received 2019 July 9; in original form 2019 April 22

ABSTRACT

We study the link between supermassive black hole growth and the stellar mass assembly of their host galaxies in the state-of-the-art ROMULUS suite of simulations. The cosmological simulations ROMULUS25 and ROMULUSC employ innovative recipes for the seeding, accretion, and dynamics of black holes in the field and cluster environments, respectively. We find that the black hole accretion rate traces the star formation rate among star-forming galaxies. This result holds for stellar masses between 10^8 and 10^{12} solar masses, with a very weak dependence on host halo mass or redshift. The inferred relation between accretion rate and star formation rate does not appear to depend on environment, as no difference is seen in the cluster/proto-cluster volume compared to the field. A model including the star formation rate, the black hole-to-stellar mass ratio, and the cold gas fraction can explain about 70 per cent of all variations in the black hole accretion rate among star-forming galaxies. Finally, bearing in mind the limited volume and resolution of these cosmological simulations, we find no evidence for a connection between black hole growth and galaxy mergers, on any time-scale and at any redshift. Black holes and their galaxies assemble in tandem in these simulations, regardless of the larger scale intergalactic environment, suggesting that black hole growth simply follows star formation on galactic scales.

Key words: black hole physics – galaxies: active – quasars: general.

1 INTRODUCTION

Every massive galaxy is believed to host a supermassive black hole (SMBH) at its centre (Kormendy & Richstone 1995). SMBH masses have been found to correlate with the stellar properties of their hosts, namely the bulge luminosity and velocity dispersion, suggestive of a co-evolutionary assembly history between these components (Haehnelt, Natarajan & Rees 1998; Magorrian et al. 1998; Ferrarese & Merritt 2000; Gebhardt et al. 2000; Kormendy & Ho 2013; Reines & Volonteri 2015; Saglia et al. 2016). SMBHs are thought to assemble their masses primarily through luminous gas accretion (Soltan 1982), which powers active galactic nuclei (AGNs) that shine at a wide range of wavelengths. In the standard galaxy evolution paradigm, galaxies must somehow provide the gas reservoir to fuel these AGN, which then impart feedback energy required to suppress star formation in massive haloes (Springel, Di Matteo & Hernquist 2005; Bower et al. 2006; Croton et al. 2006). The evolving AGN population is now constrained out to $z \lesssim 6$, allowing us to reconstruct the growth of SMBHs throughout cosmic time (Hopkins, Richards & Hernquist 2007; Ueda et al. 2014; Aird et al. 2015; Vito et al. 2018; Tasnim Ananna et al.

2019). Cosmological simulations have proven to be useful tools to better understand this evolution, and are now able to broadly reproduce observed relationships between SMBHs and their host galaxies, albeit with a variety of physical models (e.g. Di Matteo, Springel & Hernquist 2005; Hirschmann et al. 2014; Schaye et al. 2015; Sijacki et al. 2015; Volonteri et al. 2016; Tremmel et al. 2017).

However, the precise connection between AGNs and their host galaxies, exactly what ‘triggers’ an SMBH to grow and shine, is not yet fully understood. Clearly, accretion of gas from the inner regions of galaxy is the principal growth mode for most if not all black holes. To provide this gas, idealized galaxy merger simulations have implicated mergers as one of the drivers of AGN activity (e.g. Mihos & Hernquist 1996; Di Matteo et al. 2005; Capelo et al. 2015). Yet tests of this paradigm in more realistic environments, in both cosmological simulations and with observational data, have yielded mixed results. Many morphology studies of optical or X-ray AGN hosts have failed to find a connection with mergers (Cisternas et al. 2011; Mechtley et al. 2016; Villforth et al. 2017). On the other hand, mergers appear to be more important for infrared-selected samples (Treister et al. 2010; Glikman et al. 2015; Kocevski et al. 2015; Fan et al. 2016; Ricci et al. 2017; Donley et al. 2018), which are more complete to obscured accretion (e.g. Hickox & Alexander 2018). Targeted ALMA observations have also revealed statistically significant overabundances of companions among the

★ E-mail: angelo.ricarte@yale.edu

most luminous quasars at $z \sim 5$ (Trakhtenbrot et al. 2017, 2018). One possibility is that mergers only trigger the most luminous AGN, while secular processes power ‘typical’ moderate-luminosity Seyferts (Treister et al. 2012; Hickox et al. 2014). This question of the role that mergers play in the mass assembly history of black holes therefore appears to depend on the mass, luminosity, and selection of these sources.

Many cosmological simulations similarly point towards merger-independent channels of SMBH growth. Mergers are associated with AGN in the EAGLE cosmological simulations, with a more compelling link at low redshift, perhaps because gas is less readily available at late times (McAlpine et al. 2018). Yet analysing the MAGNETICUM pathfinder simulations, Steinborn et al. (2018) find that mergers only play a minor role in triggering AGN. In these simulations, the association of the most luminous AGNs with mergers can to some extent be explained by the fact that more massive galaxies are inherently more likely to be in merging systems. Similarly, only 35 per cent of SMBH growth is attributed to mergers in the HORIZON-AGN simulations, the rest originating from unknown secular processes (Martin et al. 2018). Using a series of zoom-in simulations, Pontzen et al. (2017) show that galaxy mergers are important for igniting AGN-driven outflows and quenching star formation in massive, high-redshift galaxies, but not for triggering large amounts of SMBH growth. These results of course hinge on the resolution of these simulations, as well as the sub-grid prescriptions for SMBH growth and star formation. Thus far, it is plausible that mergers, while important for shaping the properties of host galaxies, might not play a significant role for the SMBHs that they harbour.

More broadly, there is currently no consensus on what spatial scales are relevant for SMBH growth. Do AGN care about the intergalactic environment, only the gas within its sphere of influence, or some scales in between? At $z \sim 6$, the number density of luminous quasars detectable by SDSS corresponds to a typical host halo mass of $\sim 10^{13} M_{\odot}$ at that epoch, implying that only highly biased 5σ peaks in the density field could host the fastest growing SMBHs (Fan et al. 2003). AGN clustering studies at a wide range of redshifts and luminosities also yield approximately $10^{13} M_{\odot}$ as the characteristic host halo mass (see e.g. Cappelluti, Allevato & Finoguenov 2012, for a review). Consequently, it was expected that we should find galaxy overdensities in the vicinity of the most luminous quasars. However, no correlation has been found between luminous quasars and protocluster regions at $z \sim 3.8$ (Uchiyama et al. 2018). A similar negative result has been found for millimetre source overdensities at $6 < z < 7$, albeit within ALMA’s small field of view (Champagne et al. 2018). Using a large cosmological simulation, Di Matteo et al. (2017) found that the most massive SMBHs were found not in the largest overdensities, but rather in areas with the lowest tidal fields. Further complicating this picture, feedback from quasars may be capable of inhibiting star formation in nearby haloes, potentially hiding the overdensities in which they reside (Habouzit et al. 2018). How the intergalactic environment influences SMBH growth, if at all, is a rapidly evolving field of research.

One aspect of SMBH astrophysics that complicates all of these studies is that AGNs are variable on every time-scale (Hickox et al. 2014; Sartori et al. 2018), and in particular on shorter time-scales than that of star formation. Circumstantial evidence exists for AGN ‘flickering’ on time-scales of $\sim 10^4$ – 10^5 yr if optically elusive X-ray AGNs are interpreted as AGNs that have not yet had time to photoionize their host galaxies (Schawinski et al. 2015). This argument is supported by hydrodynamical simulations with very fine time-resolution (Novak, Ostriker & Ciotti 2011) as well as analytic considerations based on the maximum sizes of discs that

do not fragment under their self-gravity (King & Nixon 2015). Recently, several ‘changing look’ AGNs have also been identified, which have been observed to switch between unobscured (type I) and obscured (type II) accretion states along with similar changes in their levels of continuum emission on $\lesssim 10$ yr time-scales (e.g. LaMassa et al. 2015; MacLeod et al. 2016). AGN variability may wash out true correlations that exist on longer time-scales, and it is important to better understand this behaviour.

In this paper, we follow and study the demographics and assembly of SMBHs in the ROMULUS simulations, which implement state-of-the-art recipes for SMBH seeding, accretion, and dynamics. ROMULUS25 is a uniform 25 Mpc-per-side volume (Tremmel et al. 2017) that represents the field environment, while ROMULUSC is a zoom-in on a low-mass cluster of $1.5 \times 10^{14} M_{\odot}$ (Tremmel et al. 2019). Both are run to $z \approx 0$. This suite of simulations therefore brackets the range of environments in which black holes and galaxies can grow. Cosmological simulations such as the ROMULUS suite allow us to study SMBH-galaxy co-evolution in realistic large-scale environments. Although computational limitations restrict our ability to directly resolve star formation and SMBH accretion processes, these complex processes are approximated using innovative and physically motivated ‘sub-grid’ recipes. One unique advantage of the ROMULUS suite is the fact that the same sub-grid recipes have been implemented for both sets of environments – the field and a cluster.

Here, we explore the relations linking galactic and intergalactic properties to the accretion rates of their SMBHs. In particular, a causal explanation of the local correlation between SMBH and host stellar masses implies that there must exist some time-scale over which the black hole accretion rate and the star formation rate (SFR) trace each other. In the past decade, many observational studies have been performed to test this hypothesis, with mixed results (e.g. Mullaney et al. 2012; Stanley et al. 2015; Aird, Coil & Georgakakis 2019). We demonstrate in this work that the SMBH accretion rate follows the SFR well when smoothed to time-scales of ~ 300 Myr, and that the SFR, along with two other parameters, can explain up to 68 per cent of the variations in the SMBH growth rate among star-forming galaxies. Overall, these simulations suggest that SMBHs and their hosts grow in lockstep, but with different variability time-scales driven by the stochasticity in both the physics of star formation and black hole accretion. Mergers play no noticeable role for the AGN in ROMULUS, but we note that co-evolution might well proceed differently for rare objects that would not appear in its volume: the most massive SMBHs, luminous SDSS quasars, and transient objects such as ultra-luminous infrared galaxies (ULIRGs).

Our paper is organized as follows. In Section 2, we describe the pertinent details of the ROMULUS simulations. In Section 3, we report the results of our analysis, including derived local relations (Section 3.2), the SMBH accretion density (Section 3.3), the connection between AGN and SFRs (Section 3.4) followed by a statistical search for the most fundamental relation for SMBH growth (Section 3.5), Eddington ratio distributions (Section 3.7), and the potential link between AGN and galaxy mergers (Section 3.8). We discuss our findings in the context of other recent work in Section 4, and summarize our key findings in Section 5.

2 THE ROMULUS SIMULATIONS

In this paper, we examine SMBH assembly in ROMULUS25, a uniform 25 Mpc-per-side box, and ROMULUSC, a zoom-in simulation of a $10^{14} M_{\odot}$ cluster. At $z = 0.05$, the latest available redshift slice, ROMULUS25 contains five haloes of at least $10^{13} M_{\odot}$ and

39 haloes of at least $10^{12} M_{\odot}$. These are some of the highest resolution cosmological simulations that exist with comparable volume. SMBHs are treated in a more physical manner than in many previous approaches: they are seeded based on local gas properties, Bondi accretion is modified to account for angular momentum support, feedback is imparted via a single thermal mode, and dynamics are corrected to account for unresolved dynamical friction. All free parameters are set using a novel optimization technique to grade a small set of zoom-in simulations against empirical relations. Below, we summarize the pertinent aspects of the ROMULUS simulations, and more details on the simulation methodology can be found in Tremmel et al. (2017, 2019).

2.1 Numerics

The ROMULUS simulations were performed with the Tree+Smoothed Particle Hydrodynamics (SPH) code CHANGA (Menon et al. 2015). A standard Λ CDM cosmology was assumed with the Planck cosmological parameters (Planck Collaboration et al. 2016). Dark matter and gas particles have masses of $3.39 \times 10^5 M_{\odot}$ and $2.12 \times 10^5 M_{\odot}$, respectively. Gravity is resolved with a Plummer-equivalent softening length of 250 pc, while hydrodynamics are evaluated with a resolution of 70 pc. CHANGA includes an updated SPH implementation, allowing for more the accurate modelling of shearing flows with Kelvin–Helmholtz instabilities (Governato et al. 2015; Menon et al. 2015; Wadsley, Keller & Quinn 2017). The simulation also includes an updated implementation of turbulent diffusion, important for correctly modelling gas thermodynamics, and metal distribution (Wadsley, Veeravalli & Couchman 2008; Shen, Wadsley & Stinson 2010; Wadsley et al. 2017).

2.2 Star formation and cooling

Star formation is modelled with standard recipes, whereby gas particles with temperatures below 10^4 K and above densities of $0.2 m_p/\text{cc}$ may form stars with a characteristic time-scale of 10^6 yr and an efficiency of 15 percent, assuming a Kroupa initial mass function (IMF) (Kroupa 2001) and with supernova feedback proceeding via the ‘blastwave’ implementation (Stinson et al. 2006). Cooling in low-temperature ($<10^4$ K) gas is regulated by the metal abundance (Guedes et al. 2011) as well as thermal and metal diffusion (Shen et al. 2010; Governato et al. 2015), but high-temperature metal-line cooling is not included (see Tremmel et al. 2019 and references therein for further discussion and details). Molecular hydrogen abundance and cooling is not followed in the simulation, as the resolution is not high enough to properly resolve the multiphase interstellar medium at that level of detail.

2.3 SMBH seeding

The physics of seeding, the abundance of seeds and their initial mass function is one of the key unsolved problems in black hole physics today (see Ricarte & Natarajan 2018b, for a recent detailed discussion). Leveraging the resolution of the ROMULUS simulations, SMBHs are seeded using recipes based on the local gas properties, rather than simply imposing a halo mass threshold as is often implemented (e.g. Springel et al. 2005; Vogelsberger et al. 2013; Schaye et al. 2015). A gas particle selected to form a star will instead form an SMBH if all of the following criteria hold true:

- (i) Its metallicity is less than $3 \times 10^{-4} Z_{\odot}$.

- (ii) Its mass density is greater than $3 m_p/\text{cc}$.
- (iii) Its temperature is between 9500 and 10000 K.

These criteria ensure that SMBHs form in regions that are collapsing quickly, but also cooling slowly, limiting seed formation to high-density peaks in the early universe with high Jeans masses. Seeds in the ROMULUS suite are initialized with masses of $10^6 M_{\odot}$, immediately accreting from nearby gas particles if necessary, in order to conserve mass. Although this seed mass is somewhat high even for optimistic direct-collapse scenarios (Lodato & Natarajan 2007), it is important that seed masses be significantly greater than those of gas or dark matter particles in order to resolve dynamical friction and avoid spurious scattering events. Note that there are no limitations on either the number of seeds that can form in a single halo or how close together seeds can form. Consequently, some haloes can form multiple seeds that rapidly merge, resulting in ‘effective’ seed masses of a few times $10^6 M_{\odot}$.

2.4 SMBH accretion

SMBH accretion is estimated via a modified Bondi–Hoyle prescription (Bondi & Hoyle 1944), corrected to account for the rotational support of surrounding gas on resolved scales (see Tremmel et al. 2017, for more details). Such rotational support has been shown to be important for SMBH growth and feedback (Hopkins & Quataert 2010; Rosas-Guevara et al. 2015; Tremmel et al. 2017). An SMBH’s accretion rate is taken to be the minimum of the Eddington accretion rate and one of the following:

$$\dot{M}_{\bullet} = \left(\frac{n}{n_*}\right)^{\beta} \begin{cases} \frac{\pi(GM_{\bullet})^2 \rho}{(v_{\text{bulk}}^2 + c_s^2)^{3/2}} & \text{if } v_{\text{bulk}} > v_{\theta} \\ \frac{\pi(GM_{\bullet})^2 \rho c_s}{(v_{\theta}^2 + c_s^2)^2} & \text{if } v_{\text{bulk}} < v_{\theta}, \end{cases} \quad (1)$$

where G is the gravitational constant, n_* is the star formation threshold number density, v_{bulk} is the bulk velocity of the gas, v_{θ} is its rotational velocity, c_s is its sound speed, ρ is its mass density, and β is a free parameter set to 2 based on the optimization procedure described in Section 2.7. The prefactor $(n/n_*)^{\beta}$ is a boost factor commonly employed to correct for underestimates of the gas density and temperature due to the limited resolution of the simulation (Booth & Schaye 2009).

The ROMULUS simulations only employ one mode of AGN feedback, whereas other simulations sometimes implement two distinguishable modes – ‘quasar/radiative’ and ‘radio/mechanical’ modes, for example. In ROMULUS, if some mass $\dot{M}_{\bullet} \Delta t$ is accreted by an SMBH, then an amount of energy $E = \epsilon_r \epsilon_f \dot{M}_{\bullet} c^2 \Delta t$ is injected isotropically into the nearest 32 gas particles. Here, ϵ_r is the radiative efficiency of the accretion disc (assumed to be 0.1), ϵ_f is the feedback coupling efficiency (set to 0.02), and c is the speed of light. In order to avoid numerical overcooling due to limited mass and time-resolution, gas that receives AGN feedback energy is prevented from cooling for the duration of the SMBH’s time-step. This is meant to approximate a continuous transfer of energy throughout the SMBH’s time-step. In order to avoid long cooling shutoff times, and to more continuously sample the interaction between SMBHs and nearby gas, SMBHs are placed on the smallest global time-step of the simulation, typically 10^4 – 10^5 yr. With both the brief cooling shutoff and the short time-steps, ROMULUS is able to avoid gas overcooling, which often artificially suppresses the effects of thermal feedback models. As a result, SMBHs in ROMULUS are able to drive powerful outflows that successfully regulate and sometimes quench star formation in massive galaxies (Pontzen et al. 2017;

Tremmel et al. 2019), as well as enrich the circumgalactic medium (Sanchez et al. 2018).

2.5 SMBH dynamics

New techniques are employed to accurately estimate the dynamical friction of surrounding matter on to SMBHs in these simulations (Tremmel et al. 2015). Unless a correction factor is added, the dynamical friction force on to SMBHs is underestimated due to gravitational softening. This correction factor is obtained by integrating the Chandrasekhar (1943) formula within the gravitational softening length ϵ_g , resulting in

$$a_{\text{DF}} = -4\pi G^2 M_\bullet \rho(< v_\bullet) \ln \Lambda \frac{v_\bullet}{v^3}, \quad (2)$$

where $\rho(< v_\bullet)$ is mass density of particles moving slower than the black hole, and $\ln \Lambda = \ln(b_{\text{max}}/b_{\text{min}})$, where b_{max} is set to ϵ_g and b_{min} is set to the 90° deflection radius.

2.6 Halo finding and data analysis

Halo finding is performed with the AMIGA Halo Finder (Knollmann & Knebe 2009), key properties are calculated using PYNBODY (Pontzen et al. 2013), and these are organized into a TANGOS data base (Pontzen & Tremmel 2018). Note that ROMULUSC is a zoom-in simulation containing a high-resolution region within a low-resolution dark matter-only region. Due to the peculiarities of this technique, some haloes can be ‘contaminated’ by low-resolution particles. We therefore exclude from our analysis any haloes for which more than 5 per cent of the dark matter particles originate from the low-resolution region.

While fitting the stellar-to-halo mass relation, Munshi et al. (2013) determine that hydrodynamical simulations systematically overestimate stellar masses and underestimate halo masses if these are derived from photometry and a critical density threshold, respectively. Stellar masses derived by simply summing star particles are systematically higher than masses estimated from single-band observations even when realistic surface brightness limitations are taken into account. Meanwhile, the addition of baryons causes virial radii defined by an overdensity threshold to decrease, shrinking halo masses relative to dark matter only simulations. Throughout this paper, we multiply all stellar masses by 0.6 and divide all halo masses by 0.8 in order to account for these two countervailing effects on the halo mass and stellar mass (see fig. 5 of Munshi et al. 2013). We include in our analysis any halo that contains at least 10^4 dark matter particles and that has a stellar mass of at least $10^8 M_\odot$ (post-correction), ensuring that the objects that we study are well resolved.

2.7 Systematic parameter optimization

One unique feature of the ROMULUS simulations is the systematic method by which optimal parameters for sub-grid recipes were honed. Four different sets of zoom-in simulations were performed with halo masses ranging from $10^{10.5}$ to $10^{12} M_\odot$, with dozens of different sub-grid parameter realizations. Outputs were then graded based on a combination of empirically measured $z = 0$ scaling relations: the stellar-to-halo mass relation (Moster, Naab & White 2013), the H I gas fraction as a function of stellar mass (Cannon et al. 2011; Haynes et al. 2011), the galaxy specific angular momentum as a function of stellar mass (Obreschkow & Glazebrook 2014), and the SMBH-to-stellar mass relation (Schramm & Silverman 2013).

The first set of simulations was run varying only star formation parameters, which were then fixed for the subsequent set which optimized SMBH accretion parameters.

Note that this parameter search was confined to halo masses below which AGNs are thought to be the dominant source of feedback energy. Results for haloes more massive than $10^{12} M_\odot$ were not directly optimized for, and are therefore purely predictions of the simulation, as they are not part of the calibration. Furthermore, only $z = 0$ properties were used to anchor and constrain the simulation, so the evolution of galaxies and SMBHs with time is also consequently a prediction of the simulation.

3 RESULTS

3.1 Galaxies in the field and cluster

To set the context of this work, we first begin by examining the galaxy populations of ROMULUS25 and ROMULUSC. In Fig. 1, we plot SFRs versus stellar masses of both simulations for redshifts $z \in \{0.05, 1, 3\}$. In order to illustrate how these galaxies might be selected in a flux-limited AGN survey, points are colour coded according to the bolometric luminosities of their most massive¹ SMBHs. Divisions between tiers are set at 10^{43} and $10^{45} \text{ erg s}^{-1}$. Both the SFR and black hole accretion rate (BHAR) are averaged over 300 Myr. Note that variability can increase the probability that an AGN passes a given flux limit, while also decreasing its duty cycle to conserve mass. An AGN shining on average at $10^{43} \text{ erg s}^{-1}$ can be interpreted as an AGN shining at $10^{44} \text{ erg s}^{-1}$ one tenth of the time, for example.

The green band corresponds to the observed star formation sequence from combined UV and IR SFRs (Whitaker et al. 2012). These observations report a shallow, mildly evolving slope of ≈ 0.6 and a constant scatter of 0.34 dex.² The blue band corresponds to the best-fitting star-forming main sequence derived from central, isolated galaxies in ROMULUS25 (Tremmel et al. 2017). We follow a similar procedure to the observations (e.g. Bluck et al. 2016) and fit the median values of the SFR within 0.1 dex bins of stellar mass between 10^8 and $10^{10} M_\odot$. Note that the ROMULUS simulations appear to underestimate the star formation in low-mass galaxies at high redshift. However, Whitaker et al. (2012) is only complete to stellar masses above $\gtrsim 10^{10} M_\odot$ for $z \gtrsim 1$, where the two bands do overlap. In the work that follows, a galaxy’s proximity to the main sequence is computed relative to the fits from the ROMULUS simulations rather than to the observed data, for internal consistency.

As one might expect, many more of the $z = 0.05$ galaxies are quenched in the overdense environment of ROMULUSC than in ROMULUS25. In particular, ROMULUSC exhibits a quenched fraction of 80–90 per cent at low masses, compared to 10 per cent in ROMULUS25 (see also fig. 14 in Tremmel et al. 2019). However, among those that are not quenched, there is no indication that their SFRs are different from those of ROMULUS25. That is, a star-forming galaxy in the cluster or proto-cluster appears to be no different from a star-forming galaxy in the field. Our colour coding reveals

¹Throughout this work, we select the most massive SMBH to represent the SMBH/AGN of each host galaxy (and exclude any secondary less massive black holes that might exist in the same halo). Checking at $z = 0.05$, the most massive, most luminous, and most central SMBHs are identical 95.5 per cent of the time.

²For the last panel, equations (1)–(3) in Whitaker et al. (2012) have been extrapolated slightly from $z = 2.5$ to $z = 3$.




Figure 1. SFRs and stellar masses for ROMULUS25 and ROMULUSC at redshifts 0.05, 1, and 3. To illustrate how these galaxies might be selected by a flux-limited AGN survey, points are colour coded according to the bolometric luminosities of their most massive SMBHs that are hosted in these galaxies. Both the black hole accretion rates and SFRs are averaged over 300 Myr, and non-unity duty cycles over this time period would affect selection. The blue region corresponds to the star-forming sequence inferred from the ROMULUS25 data points, while the green region corresponds to the empirical relation determined by Whitaker et al. (2012).

that the resolutions of the ROMULUS simulations allow us to probe much lower galaxy masses than those typically accessible to flux-limited surveys. X-ray or quasar surveys can typically only reveal the most massive subset of the population we analyse, especially at high redshift. It is important to point out that the work we present applies more to typical galaxies than to luminous $> 10^{45} \text{ erg s}^{-1}$ quasars, rare-objects for which ROMULUS25 lacks the volume for exploration.

3.2 Local SMBH–galaxy relations

We now examine the derived $z = 0$ relations between SMBH mass and the host galaxy stellar mass in the ROMULUS suite. These local relations serve as a boundary condition for determining the assembly history of SMBHs. There exist well-established observed relations between SMBH masses and the stellar contents of their hosts. SMBH mass has been shown to correlate with an intrinsic scatter of $\sim 0.3\text{--}0.5$ dex with both the luminosity and velocity dispersion of the galactic bulge (Beifiori et al. 2012; Kormendy & Ho 2013; Saglia et al. 2016; van den Bosch 2016). Reines & Volonteri (2015) investigate the relationship between SMBH mass and *total* stellar mass, and find that AGNs are offset below the relation for bulge-dominated galaxies by over an order of magnitude. Shankar et al. (2016) argue that the offset AGN relation may in fact be closer to the true one, and the relation defined by quiescents is significantly biased towards galaxies for which the SMBH sphere of influence can be resolved. These correlations are important for gaining insight into both fuelling and feedback as a function of host mass. They are currently adopted to calibrate models of SMBH growth over cosmic time, and many predictions for as yet undetected SMBH populations rest on its inferred slope and amplitude.

3.2.1 $M_\bullet$ to M_* and M_{200}

We plot the derived relationships between SMBH mass and total stellar mass and SMBH mass and virial mass for the ROMULUS simulations in Fig. 2. Here, once again, $M_\bullet$ corresponds to the most massive SMBH associated with each halo. Only galaxies whose most massive SMBHs are within 2 kpc of galactic centre are shown. We note that at low-stellar masses, seed masses can comprise the majority of SMBH mass. Sometimes, multiple seeds merge at high redshift, since ROMULUS does not impose any minimum distance between seeds. To compensate for this phenomenon and particulars arising from the seeding methodology adopted, here $M_{\bullet, \text{acc}}$ in these plots corresponds to the amount of SMBH mass that is obtained via accretion. This, therefore excludes all of the seed masses that contribute to the final mass, but includes the accreted portion of every SMBH’s mass that may have merged to create the final SMBH. We confirm that for most SMBHs with mass $> 10^7 M_\odot$ (or a factor of 10 larger than their initial mass), most of the total mass is produced through accretion and their final masses are insensitive to the seed mass assumed. Points are colour coded according to each SMBH’s Eddington ratio averaged over the past 300 Myr (a measure of its specific growth rate further discussed in Section 3.7). Those with Eddington ratios above 10^{-2} are marked with stars instead of circles. We overplot relationships between SMBH and stellar mass from Kormendy & Ho (2013) and Reines & Volonteri (2015). To derive the correlations with halo mass, we convert these observed relationships via the stellar-to-halo mass relation obtained from abundance matching (Moster et al. 2013). Note that the stellar-to-halo mass relation is very uncertain below halo masses of $\sim 10^{11} M_\odot$, and the SMBH-to-stellar mass relation is not observationally measured for stellar masses below $\sim 10^{10} M_\odot$.




Figure 2. SMBH mass versus host stellar mass (*left*) and host virial mass (*right*) in the ROMULUS simulations. Only accreted mass is included, excluding the seed mass contribution. Relationships from Kormendy & Ho (2013) and Reines & Volonteri (2015) are shown for comparison. The ROMULUS simulations agree with the relationships observed for bulge-dominated galaxies as calibrated, but with significant scatter at low masses. Tidal stripping moves the ROMULUSC data points to the left relative to ROMULUS25 in the right-hand panel, but its effects are weak on the host’s stellar mass. Points are colour coded according to Eddington ratio, revealing a possible explanation for the different relations observed for bulge-dominated galaxies and BLAGN. If a broad-line region only appears above a certain Eddington ratio threshold, then at fixed accretion rate, lower mass SMBHs are more likely to be BLAGN.

Since ROMULUS was calibrated to obtain reasonable SMBH masses for low stellar mass host galaxies, it is encouraging that appropriate masses are obtained for the SMBHs that reside in the most massive haloes, including the $\sim 10^{10} M_{\odot}$ SMBH in the brightest cluster galaxy (BCG) of ROMULUSC. There is noticeably more scatter at low masses than at high masses, as we would expect from the Central Limit Theorem (Peng 2007; Jahnke & Macciò 2011). Interestingly, there is no indication that the accreted mass departs from a linear relation even below the seed mass of $10^6 M_{\odot}$, suggesting that galactic inflows on larger scales than the SMBH sphere of influence regulate accretion in low-mass galaxies, irrespective of the SMBH mass.

Reines & Volonteri (2015) report separate relationships for elliptical galaxies and broad-line AGNs. By separating ROMULUS galaxies into high- and low-Eddington ratio populations, we find a possible explanation. Using a multiwavelength sample of AGN, Trump et al. (2011) find that a broad-line region is only present among AGN with $f_{\text{Edd}} > 10^{-2}$, which may mark a boundary between radiatively efficient to inefficient accretion disc structures.

As shown in Fig. 2, the SMBHs at $z = 0$ that fulfil this criterion (marked as stars) have systematically lower masses. Examining their luminosities, we find that this is not because lower mass SMBHs have higher accretion rates, but simply because lower mass SMBHs have higher Eddington ratios for a *given* accretion rate.

Turning to halo mass, the $M_{\bullet}-M_{200}$ relation for ROMULUSC is noticeably offset to the left with respect to ROMULUS25. However, there is no noticeable difference in the $M_{\bullet}-M_{*}$ relationship between the cluster and the field. This is likely due to tidal stripping, which removes the outer regions of dark matter haloes of cluster members, but is not usually strong enough to impact the total stellar mass. One might have expected an elevated $M_{\bullet}-M_{*}$ relation in ROMULUSC compared to ROMULUS25, since both components assemble earlier in the Universe when the ratio $M_{\bullet}/M_{*}$ is thought to be higher (e.g. Yang et al. 2018). Yet as we later show in Section 3.4, in ROMULUS the $M_{\bullet}-M_{*}$ relation is established earlier in the universe and it does not evolve strongly with redshift. Furthermore, the main halo of ROMULUSC does not assemble half of its final mass until $z = 0.54$, which is relatively recent to expect significant evolution.

In contrast with several other hydrodynamical simulations, there is no indication that SMBH growth is stunted in the lowest mass galaxies due to supernova feedback. The EAGLE simulations find that a critical host virial temperature is required for runaway SMBH growth (McAlpine et al. 2018). Bower et al. (2017) suggest that this threshold represents the point beyond which supernova feedback cannot prevent the build-up of gas in the central regions of galaxies. Dubois et al. (2015) and Habouzit, Volonteri & Dubois (2017) similarly find that feedback strong enough to curtail overproduction of stars stunts early SMBH growth in the Seth and SuperChunky simulations, respectively. In the absence of AGN feedback, supernova feedback also regulates SMBH growth in the FIRE simulations (Anglés-Alcázar et al. 2017). We speculate that the different recipes for accretion and/or feedback used in these simulations may be responsible for these mixed findings. The relatively high seed mass in ROMULUS, and the propensity for several immediate mergers between seeds, may also play a role. We further discuss potential explanations of this disagreement later in the Section 4.

3.2.2 $M_\bullet$ to σ

Some authors have argued that the host velocity dispersion, σ , is the more fundamental quantity in determining SMBH mass growth (Volonteri, Natarajan & Gültekin 2011; Beifiori et al. 2012; van den Bosch 2016). This quantity reflects not only the host’s mass but also the depth of its potential well. Indeed, AGNs appear to be more common among galaxies that are not only massive but also preferentially compact for their masses (Kocevski et al. 2017; Powell et al. 2017). An $M_\bullet$ – σ relation can also be theoretically generated by SMBH feedback, with momentum- and energy-driven winds yielding $M_\bullet \propto \sigma^4$ (King 2003) and $M_\bullet \propto \sigma^5$ (Haehnelt et al. 1998) relations, respectively (see also Natarajan & Treister 2009; Zubovas & King 2012; King & Pounds 2015).

We plot the $M_\bullet$ – σ relation for ROMULUS galaxies in Fig. 3 along with the relations observed by Kormendy & Ho (2013), van den Bosch (2016), and Martín-Navarro & Mezcua (2018) for which we have provided the slope of the $M_\bullet$ – σ relation, commonly written as β . Note that these first two samples are limited to $\sigma \gtrsim 100 \text{ km s}^{-1}$, while Martín-Navarro & Mezcua (2018) consider only Seyfert I galaxies. We estimate σ directly from the star particles in the each galaxy. First, we calculate the effective radius, R_{eff} , based on the surface brightness profiles of each galaxy. This is performed in the ‘i’ band using cylindrical annuli with the galaxy rotated so that it would be face-on to the observer. The rotation was performed based on the angular momentum of gas, stars, and dark matter within the inner 5 kpc of the halo. A single Sérsic profile was fit to each galaxy assuming surface brightness cut-off of 32 mag arcsec $^{-2}$ and a maximum radius of $5 \times$ the half-light radius. Then, we compute $\sigma = \sqrt{\langle v^2 \rangle - \langle v \rangle^2}$. In order to potentially isolate the velocity dispersion of a bulge that is smaller than the effective radius, we define $\sigma_{\text{max}} (< R_{\text{eff}})$ as the maximum velocity dispersion that an observer could infer by enclosing star particles within a maximum radius that varies between 1 kpc (due to our limited resolution) and the effective radius. In practice, we notice little difference between computing σ in this manner compared to calculating σ for stars enclosed within the effective radius.

We find that for galaxies with stellar mass above $10^{10} M_\odot$, ROMULUS agrees well with the observed relations for high-mass galaxies. Even the most massive SMBH in ROMULUSC, that at the centre of the BCG, falls on the relations observed at high masses. At lower masses, ROMULUS departs from these relations,

yielding higher SMBH masses for their hosts’ velocity dispersions. For Seyfert galaxies, Martín-Navarro & Mezcua (2018) report that the $M_\bullet$ – σ relation flattens at these lower host masses, but with a much lower normalization than occurs in ROMULUS (Martín-Navarro & Mezcua 2018). In ROMULUS, the $M_\bullet$ – σ relation appears to be less fundamental than the $M_\bullet$ – M_* relation. As we later show in Section 3.4, this is consistent with the phenomenon that SMBHs grow in tandem with the stellar content of their entire galaxies in these simulations. Nevertheless, we caution that while the departure from local relations at low masses is plausible, it is likely quite sensitive to sub-grid recipes for both star formation and SMBH growth, and the resolution of the simulation (Anglés-Alcázar et al. 2017).

3.3 SMBH accretion density

The $(25 \text{ Mpc})^3$ volume of ROMULUS25 does not allow for the robust predictions of AGN luminosity functions, since the knee of the luminosity function falls below the minimum number density that can be probed in this box-size. In order to compare the assembly history of the ROMULUS25 SMBHs with the redshift-evolving AGN census, we instead calculate the bolometric luminosity density. This is the integral of the luminosity function and captures the total amount of accretion that occurs at a given epoch. This quantity is unaffected by any time-variability that is not resolved by the simulation. Variability may alter the shape of luminosity functions, but not its integral, making it a robust quantity to compare with observational data.

We plot the AGN bolometric luminosity density in ROMULUS25 in Fig. 4 for four different bolometric luminosity thresholds corresponding to 0, 10^{43} , 10^{44} , and $10^{45} \text{ erg s}^{-1}$. For readability, the solid curves have been smoothed using a boxcar filter with a full width of $\Delta \log(1+z) = 0.05$. In light red, we also plot the original $\Delta t = 10 \text{ Myr}$ resolution with no luminosity threshold. There is substantial stochasticity due to the relatively small box-size of ROMULUS25. Along the same lines, the $10^{45} \text{ erg s}^{-1}$ curve also reveals that a few rare, luminous objects drive the shape of the volume-averaged variability.

We then compare to the bolometric luminosity density estimated from two different population synthesis models based on X-ray observations (Aird et al. 2015; Tasnim Ananna et al. 2019). These measurements are sensitive to absorbed AGN and include an estimate of the Compton-thick contribution required in order to fit the cosmic X-ray background spectrum. The model of Tasnim Ananna et al. (2019) sits vertically offset from that of Aird et al. (2015) due to the higher inferred fraction of Compton thick sources. There is an interesting disagreement between these population synthesis models and ROMULUS25 in terms of both the shape and normalization of these curves. In ROMULUS25, the luminosity density matches the observations only during the epoch of peak AGN activity and stays remarkably constant throughout cosmic time. We offer two possible explanations. In Tremmel et al. (2017), it was shown that ROMULUS25 similarly overestimates the density of star formation at $z \lesssim 1$ and possibly $z \gtrsim 3$. The quenched fraction in ROMULUS is systematically low, particularly at high masses (Tremmel et al. 2019). As we shall show in Section 3.4, AGN activity follows the SFR in ROMULUS, and this is likely the cause of the overestimate of the AGN luminosity density as well. At high redshift, it is also possible that current surveys miss fainter and/or more obscured populations. As we show, the total luminosity density in ROMULUS25 is sensitive to the threshold above which we include SMBH activity.




Figure 3. SMBH mass versus host stellar velocity dispersion in ROMULUS25. Observed relations from van den Bosch (2016), Kormendy & Ho (2013), and Martín-Navarro & Mezcuca (2018) are shown, along with their slopes. Points are colour coded according to stellar mass, and more massive galaxies have higher velocity dispersions, as expected.




Figure 4. Luminosity density as a function of time in ROMULUS25. We compare estimates from the population synthesis models of Aird et al. (2015) and Tasnim Ananna et al. (2019). Different colours correspond to different bolometric luminosity thresholds as shown in $\log \text{erg s}^{-1}$. Solid lines are smoothed with a boxcar filter, while in light red we display the variability of the original 10 Myr time-resolution. Unlike the population synthesis models, the accretion density is relatively constant for $z < 5$ in ROMULUS25.

3.4 The AGN main sequence

We now explore the relationships between SMBHs and their hosts to determine how SMBHs and their galaxies co-evolve. By stacking X-ray observations of star-forming galaxies, Mullaney et al. (2012) reported a linear relation between the BHAR and the SFR, the so-called hidden AGN main sequence. Subsequent studies stacking SFRs of X-ray AGN host galaxies appeared to refute this picture, except for perhaps the most luminous AGN (Rosario et al. 2012;

Azadi et al. 2015; Stanley et al. 2015; Ramasawmy et al. 2019). It is now thought that an AGN varies on shorter time-scales than its host’s star formation (Hickox et al. 2014), washing out correlations if one stacks on AGN luminosity instead of SFR (Azadi et al. 2015; Lanzuisi et al. 2017). In addition, the relation obtained depends on the shape of the bivariate distribution, how objects are selected, and how they are binned (Volonteri et al. 2015a; McAlpine et al. 2017). Recently, some authors have claimed that stellar mass, rather than SFR, is the more fundamental quantity governing the BHAR (Yang et al. 2017; Fornasini et al. 2018), except perhaps among bulge-dominated galaxies (Yang et al. 2019). Aird et al. (2019) found relationships between *specific* accretion rates and the SFR, which was elevated for quiescent galaxies.

We plot the relationship between BHAR and SFR for the star-forming galaxies of ROMULUS25 in Fig. 5. We include only galaxies with stellar masses of at least $10^8 M_\odot$ that are no more than 1 dex below the star-forming sequence. The BHAR of the most massive SMBH is shown, and we exclude cases where the SMBH is greater than 2 kpc from the halo centre. Finally, in order to mitigate AGN variability and compare both quantities on the same time-scale, we average both the BHAR and the SFR over the past 300 Myr. We plot a simple linear regression in logarithmic space (a power-law fit) as a black dotted line. In grey, we overplot the relation observed at $z = 1$ by Mullaney et al. (2012). ROMULUS25 agrees remarkably well with this relation, even at much higher and lower redshifts. Points are colour coded according to the ratio between SMBH mass and stellar mass. This reveals a clear vertical gradient, indicating that some of the scatter is correlated with SMBH-to-stellar mass ratio. Note that in the least massive galaxies, mass ratios are artificially large due to the large seed mass. We will return to this correlation in Section 3.5.

The relationship between BHAR and SFR does not appear to change as a function of redshift or stellar mass. This is better illustrated by Fig. 6. Here, we represent the same data in Fig. 5 and compute the moving average $\dot{M}_\bullet / \dot{M}_*$ and its standard deviation using a boxcar filter with total width of $\Delta \log(M_*/M_\odot) = 0.6$. The black region represents the Mullaney et al. (2012) $z \sim 1$ relation, a constant ratio. ROMULUS25 is consistent with this relation at all redshifts and stellar masses. At the highest masses, there might be a




Figure 5. BHAR versus SFR for star-forming galaxies in ROMULUS25. Both quantities are averaged over the past 300 Myr. In grey, we overplot the relation observed at $z = 1$ from stacked star-forming galaxies (Mullaney et al. 2012). This relation agrees well with a linear regression, shown as a black dotted line, regardless of redshift. Points are colour coded according to the SMBH-to-stellar mass ratio.




Figure 6. The ratio of BHAR to SFR in ROMULUS25 as a function of stellar mass. There is no trend as a function of either stellar mass or redshift. The black region corresponds to the $z \sim 1$ relation observed by Mullaney et al. (2012).

slight upturn, especially at lower redshifts, albeit within the scatter of these relations. This may be related to the fact that there do not exist galaxies with low values of $M_\bullet/M_*$ for stellar masses above $10^{11} M_\odot$ (see Fig. 2).

In Fig. 7, we repeat this analysis for ROMULUSC to search for differences between the field and cluster. Note that there are many

fewer galaxies included at low redshift because the majority of cluster members become quenched. Remarkably, despite the very different intergalactic environments, the galaxies which remain on the star-forming sequence are still consistent with the Mullaney et al. (2012) relation. Hence, there appears to be no difference between the cluster and the field. Previous work on ROMULUSC has determined that although AGNs are suppressed overall in the cluster, there still exists a high-Eddington ratio population at low redshift (see fig. 18 in Tremmel et al. 2019). Our interpretation is that quenching is demonstrably more efficient in the cluster environment, but before it is quenched, a galaxy and its SMBH assemble independently of the larger intergalactic environment.

Finally, we also test the hypothesis that the BHAR depends on the stellar mass more strongly than the SFR (Yang et al. 2017, 2018). In Fig. 8, we repeat our analysis of the AGN main-sequence in ROMULUS25, replacing the independent variable with stellar mass. It is visually apparent that the scatter of the relation is smaller with stellar mass than with SFR, but this is true only at fixed redshift. In ROMULUS, the normalization of the BHAR– M_* correlation increases with increasing redshift, but there is no evidence of redshift evolution in the relationship between BHAR and SFR (see Fig. 6). We therefore arrive at the nuanced conclusion that stellar mass is a better predictor of the BHAR at fixed redshift, but SFR is a better predictor of the BHAR across a large range of redshifts. This motivates a more statistically rigorous multivariable analysis, which we undertake in the following section.

3.5 A Statistical search for the most fundamental relation

To search for the most fundamental relations between the BHAR and global galaxy properties, we perform multilinear regressions




Figure 7. As in Fig. 5, but for ROMULUSC. There are fewer points at low redshift because only star-forming galaxies are included. Since the same relation from Mullaney et al. (2012) can describe both these cluster galaxies and the field galaxies in Fig. 5, there appears to be no difference between the field and the cluster environments.




Figure 8. As in Fig. 5, but for stellar mass instead of SFR. Although the relationship is tighter than with SFR, a redshift dependence is required. A relationship derived from 18 000 galaxies in the CANDELS/GOODS-South field is overplotted in grey (Yang et al. 2017), while a redshift-evolving relationship based on CANDELS and COSMOS galaxies is shown in black (Yang et al. 2018). Note that ROMULUS tends to overpredict the BHAR compared to observational data, as we previously reported in Fig. 4.

in logarithmic space for a variety of models. Our metric of choice to discriminate between models is the Bayesian Information Criterion (BIC), which penalizes models for adding unnecessary parameters (Schwarz 1978). Making the assumption that

scatter about these relations is Gaussian, we express the BIC as

$$\text{BIC} = n \ln \widehat{\sigma_e^2} + k \ln(n), \quad (3)$$

Table 1. Best-fitting values for multilinear regressions to the relationship between $\dot{M}_\bullet$ and global galaxy properties, defined in equation (4). In these simulations, $\dot{M}_*$ is a better predictor of $\dot{M}_\bullet$ than M_* . Adding q as a second parameter greatly improves the model, while adding gas-related third parameters mildly improve the model as well.

	$\log_{10}(\dot{M}_\bullet, 0)$	C_{M_*}	C_{M_*}	C_q	$C_{(1+z)}$	$C_{f_{\text{gas}}}$	$C_{f_{\text{cold}}}$	$\sqrt{\sigma_e^2}$	R^2	BIC
Model A	−9.78		0.709					0.741	0.193	−606.5
Model B	−3.20	0.763						0.687	0.308	−765.3
Model C	−12.8		0.948		2.15			0.606	0.462	−1018.4
Model D	−0.0712	0.946		1.21				0.490	0.648	−1458.7
Model E	0.0418	0.977		1.29		−0.630		0.482	0.660	−1489.4
Model F	−0.386	0.904		1.17			−0.434	0.464	0.684	−1564.9

where n is the number of data points, $\widehat{\sigma_e^2}$ is the error variance, and k is the number of free parameters. The raw value of the BIC does not carry significant meaning, but a change in its value does. Preferred models have smaller (in our case, more negative) values of the BIC.

Our multilinear models take the form

$$\log_{10}(\dot{M}_\bullet) = \log_{10}(\dot{M}_\bullet, 0) + \sum C_x \log_{10}(x), \quad (4)$$

where x represents any independent variable, such as $(1+z)$ or $\dot{M}_*$. Both BHAR and SFR have units of $\text{M}_\odot \text{yr}^{-1}$. We fit the values obtained in Section 3.4 for snapshots $z \in \{0.05, 0.5, 1.0, 2.0, 3.0, 4.0\}$, epochs that are sufficiently separated in time to be treated as independent points despite our 300 Myr smoothing. In addition to the processing described in the previous section, we also restrict our fits to galaxies with stellar masses of at least $10^9 \text{M}_\odot$. This is done because the seed mass leads to artificially high values of the SMBH-to-stellar mass ratio in low-mass galaxies.

We test several different models, whose independent variables are listed here:

- (i) Model A: $\dot{M}_*$
- (ii) Model B: $\dot{M}_*$
- (iii) Model C: $\dot{M}_*, (1+z)$
- (iv) Model D: $\dot{M}_*, q$
- (v) Model E: $\dot{M}_*, q, f_{\text{gas}}$
- (vi) Model F: $\dot{M}_*, q, f_{\text{cold}}$

Here, we define $q = M_\bullet/M_*, f_{\text{gas}} = M_{\text{gas}}/(M_{\text{gas}} + M_*)$, and $f_{\text{cold}} = M_{\text{gas, cold}}/M_{\text{gas}}$. Gas is designated as ‘cold’ if it is below $2 \times 10^4 \text{K}$. Best-fitting values are provided in Table 1, along with several statistics of fit. As we show, these models are sorted in order of increasing preference.

3.5.1 Stellar mass, SFR, and other parameters

Models A and B attempt to relate the BHAR to solely the stellar mass and the SFR, respectively, with no redshift evolution. Referring to Table 1, these models have R^2 values of 0.19 and 0.31, respectively. As we have seen in Fig. 8, the poor R^2 value in Model A is due to the redshift evolution. Model C includes a $(1+z)$ dependence to this $\dot{M}_*$ -dependent model, which significantly increases its R^2 value to 0.46. However, a competing model with the same number of parameters, Model D, yields a much higher R^2 value of 0.65 with two physical parameters: $\dot{M}_*$ and q . We conclude that $\dot{M}_*$ is a more fundamental parameter than M_* , and that q provides a significant improvement on top of that single parameter. To be more succinct, we find that the models of increasing quality are $\dot{M}_* < \dot{M}_* < M_*, (1+z) < \dot{M}_*, q$. The BIC decreases by at least 100 for each of these steps, while a difference of only 10 is considered strong evidence in favour of the model with the lower BIC.

Table 2. Scatter obtained for Model D ($\dot{M}_*$ and q) with different methodologies. Averaging over longer time periods and removing quenched galaxies substantially tightens the relationship.

Method	$\sqrt{\sigma_e^2}$
Smooth over 30 Myr	0.759
Smooth over 300 Myr	0.628
As above, and remove quenched galaxies (final)	0.490

Although q has been shown to be a statistically significant additional parameter, we note that this relationship is not necessarily causal. There are physical reasons to expect a near-linear relationship with q . Recall that under the assumption of modified Bondi accretion, the accretion rate should increase with SMBH mass, as we observe. In addition, higher values of SMBH mass permit higher values of the BHAR due to the Eddington limit, although we show in Section 3.7 that the Eddington limit is probably not a major limitation in these simulations. On the other hand, the relationship with q may simply reflect the fact that galaxies which have successfully fuelled their SMBHs in the past are more likely to continue doing so. In this latter interpretation, the correlation with q is a reflection of more fundamental, unknown galactic conditions that remain stable over time-scales of at least hundreds of millions of years.

Having established a model based on $\dot{M}_*$ and q (Model D), we search for a third parameter. Two more quantities that one might expect to correlate with $\dot{M}_\bullet$ are the baryonic gas fraction, f_{gas} , and the fraction of gas which is cold, f_{cold} , which we explore in Models E and F. The addition of either of these parameters does significantly improve upon Model D on the basis of the BIC, with f_{cold} preferred over f_{gas} . Interestingly, we find that the coefficients for either of these parameters is *negative*. That is, for a given SFR and mass ratio, larger values of either the gas fraction or the cold gas fraction correspond to *less* black hole accretion. These anticorrelations persist even if all satellite galaxies are removed, defined in this case to be galaxies within the virial radius of another galaxy with higher stellar mass. This disfavors environmental processes as the drivers of these anticorrelations. Instead, this may reflect SMBH feedback: SMBHs that are accreting more efficiently can either heat the gas or evacuate it from their host galaxies.

3.5.2 A closer look at the scatter

In Table 2, we report how the scatter in Model D ($\dot{M}_*$ and q) decreases by both averaging over a longer time-scale and removing quenched galaxies. Both of these steps are important in obtaining a tighter relation. The first of these steps helps to mitigate AGN

variability. As we have illustrated in Fig. 4, there is significant stochasticity on the 10 Myr time-scales over which BHAR data are saved. The second step, removing quenched galaxies, tends to remove outliers. We comment that these quenched galaxies preferentially have high BHARs for their SFRs.

3.6 Cross-correlation analysis

In previous sections, we have established a preferred ratio between the BHAR and SFR. Here, we search for a cross-correlation between the BHAR and SFR with a variable time-lag. This would help determine if the *shapes* of these curves are interdependent. To this end, we calculate the Pearson correlation coefficient ‘ r ,’ defined for a sample of paired data $\{(x_i, y_i)\}$ as

$$r_{x,y} = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2} \sqrt{\sum_i (y_i - \bar{y})^2}}, \quad (5)$$

where $\bar{x}$ and $\bar{y}$ denote the mean of $\{x\}$ and $\{y\}$. Perfect correlation corresponds to $r = 1$, while perfect anticorrelation corresponds to $r = -1$. We emphasize that this calculation removes the means of each sample, isolating for each galaxy how the *shape* of the BHAR curves influence the SFR curves, or vice versa.

Starting at a given redshift, we first compute the BHAR and SFR for each galaxy above a stellar mass of $10^8 M_\odot$ going back 1 Gyr. To mimic the selections employed in Section 3.4, we then remove all galaxies that fall 1 dex below the star-forming sequence. Next, we calculate this value between BHAR and SFR, while adding a variable time-lag between the two quantities. To avoid adding spurious features, we ensure that we do not assume that these time-series are periodic.

In Fig. 9, we plot the results for six different starting redshifts. Cross-correlation functions for each galaxy are overlaid in the background with colours encoding the stellar mass. The x -axis corresponds to the value of the variable time-lag, with positive values representing the BHAR leading the SFR. In black, we plot the 1σ region among all galaxies, while the solid black curve represents the median.

On average, there is no cross-correlation between the BHAR and SFR with time-lags up to 0.5 Gyr. This means that although there is a preferred ratio between the BHAR and SFR as established in Sections 3.4 and 3.5, a fluctuation in one value does not predict a similar fluctuation in the other. At high redshift, it appears as if positive cross-correlations are preferred at all time-scales. However, we believe this is likely due to the fact the universe changes significantly over 1 Gyr at these redshifts, increasing in both BHAR and SFR, which naturally leads to a positive cross-correlation. This negative result does not change if we first take the logarithm of both quantities, if we restrict to narrower selections of stellar mass, or if we first smooth both quantities to a common variability time-scale of 30 Myr. These results imply that while there is a preferred ratio between the BHAR and SFR, each quantity varies independently.

3.7 Eddington ratio distributions

Like most cosmological simulations, the ROMULUS simulations cap the SMBH accretion rate at the Eddington limit. The Eddington limit is the maximum luminosity at which an accreting object can shine, assuming local force-balance between gravity and radiation pressure (Eddington 1926). This is given by

$$L_{\text{Edd}}(M_\bullet) = \frac{4\pi G M_\bullet m_p c}{\sigma_T} \approx 1.3 \times 10^{44} \left(\frac{M_\bullet}{10^6 M_\odot} \right) \text{erg s}^{-1}, \quad (6)$$

where m_p is the proton mass and σ_T is the Thompson cross-section. An SMBH’s Eddington ratio is defined $f_{\text{Edd}} = L_\bullet / L_{\text{Edd}}$.

The extent to which the Eddington limit influences SMBH assembly remains an important topic, with profound implications for the growth of the highest redshift quasars. The earliest quasars at $z \sim 6-7$ have broad-line masses that require uninterrupted Eddington-limited accretion since the seeding epoch for stellar mass seeds, or nearly uninterrupted Eddington-limited accretion for heavy seeds (Haiman & Loeb 2001; Volonteri, Silk & Dubus 2015b; Pezzulli, Valiante & Schneider 2016). The allowance of super-Eddington accretion helps assemble these masses with shorter duty cycles of feeding events. State-of-the-art general relativistic radiative magnetohydrodynamical (GRRMHD) simulations have shown that super-Eddington mass accretion rates are sustainable at the accretion disc scale (Jiang, Stone & Davis 2014; McKinney, Dai & Avara 2015; Dai et al. 2018). Yet despite this, the observed population of AGN does appear to strictly obey the Eddington limit (Wu et al. 2015). Perhaps in reality, galactic conditions on larger scales than those accessible to GRRMHD simulations do not permit super-Eddington inflows.

We investigate, therefore, the distribution of Eddington ratios in ROMULUS25, plotted in Fig. 10. We include the most massive SMBH in each halo above $10^8 M_\odot$, even if its accretion rate is zero. We plot Eddington ratio distributions where the BHAR is averaged over three different time-scales: 1.6, 10, and 100 Myr. Our distributions are derived using a Gaussian kernel density estimation technique. (Note that the shape of the high-Eddington ratio fall-off is due entirely to assuming a Gaussian kernel.) The dashed coloured lines in each panel correspond to the $z = 0$ values to guide the eye, while the black dotted line demarcates the mean log Eddington ratio, $\langle f_{\text{Edd}} \rangle$, on 1.6 Myr time-scales.

As redshift increases, the Eddington ratio distribution shifts towards higher values. This is as expected, since specific SFRs also increase with redshift, likely reflecting the fact that these galaxies are relatively more gas-rich overall. The mean Eddington ratio rises from $10^{-2.6}$ at $z = 0.05$ to $10^{-1.7}$ at $z = 6$. Overall, this behaviour is consistent with a trend $\langle f_{\text{Edd}} \rangle \propto (1+z)$. We find that the Eddington ratio distributions we derive depend somewhat on the time-scale over which they are averaged. The shorter the time-scale, the more frequent $f_{\text{Edd}} > 0.1$ events are (see the $z = 4, 6$ panels). In other words, ROMULUS25 does not sustain Eddington-limited accretion flows for 100 Myr periods. AGN Variability between 1 and 100 Myr produces more AGNs at both higher and lower accretion rates than the mean of $10^{-2.5}$. Unresolved variability on time-scales shorter than 1.6 Myr, including ‘flicker,’ may continue to modify these distributions when comparing simulation values to instantaneous values that can be observed.

Even at $z = 6$, these distributions begin to fall off well before the Eddington limit. Only 2.5 ± 1.0 per cent of SMBHs are accreting at $f_{\text{Edd}} > 0.1$ on 1.6 Myr time-scales at $z = 6$. This suggests that for the typical galaxy, the Eddington limit is not a relevant barrier to SMBH growth. Rather, intragalactic astrophysics operating on larger scales sets the accretion rate. We emphasize, however, that this may not be the case for quasars selected in flux-limited samples.

3.8 Mergers do not affect the accretion rate

High-resolution, idealized merger simulations indicate that galaxy mergers can trigger AGN by disturbing the gravitational potential, causing gas to shock against itself and lose angular momentum (Di Matteo et al. 2005; Capelo & Dotti 2017). Cosmological simulations have mostly found little to no connection with mergers (Martin et al.




Figure 9. Cross-correlations between BHAR and SFR. The black solid line corresponds to the median among all galaxies, the black filled region encloses the 1σ region, and each overlaid curve corresponds to an individual galaxy that is colour coded according to its stellar mass. There is no cross-correlation between the two quantities either instantaneously or with a time-lag. Although there is a preferred ratio between BHAR and SFR, variations in one do not predict variations in the other. Note that the mild preference for positive correlations at high redshift on all time-scales is likely due to the tendency for both BHAR and SFR to significantly increase over 1 Gyr at this epoch.




Figure 10. Eddington ratio distributions as a function of redshift and time-interval. The dotted line in each panel shows the $z = 0$ distributions. Eddington ratios rise as redshift increases. The time-scale over which the Eddington ratio is computed also matters, with shorter time-scales producing more frequent extreme low and high values.




Figure 11. Eddington ratios versus time since a merger (*blue*) or time to a merger (*red*) of mass ratio $> 1:4$ in ROMULUS25. No trend is evident at any time-scale or any redshift.

2018; McAlpine et al. 2018; Steinborn et al. 2018). That said, high-resolution studies indicate that much of the angular momentum loss occurs within 50 pc of the SMBH (Capelo et al. 2015), which is currently inaccessible to cosmological simulations. Morphological studies of the hosts of X-ray or optically selected AGNs have also struggled to find a significant connection between merger and AGN (Cisternas et al. 2011; Mechtley et al. 2016; Villforth et al. 2017). Infra-red selected AGNs, however, tend to be found more often in merging systems (Glikman et al. 2015; Kocevski et al. 2015; Fan et al. 2016; Ricci et al. 2017; Donley et al. 2018; Trakhtenbrot et al. 2018), suggesting that late-stage merger-driven SMBH growth could be heavily obscured.

To investigate the AGN-merger connection, we compute merger histories for all galaxies in ROMULUS25 at $z = 0.05$ that have a mass of at least $10^8 M_\odot$, following the main progenitor branch. Then, for different redshift slices, we compute for each progenitor (i) its Eddington ratio, (ii) the time since its most recent major merger, and (iii) the time to its next major merger. Here, a major merger is defined as a stellar mass ratio of at least $1:4$, and we find that our results do not change if this ratio is lowered to $1:10$. Although we claim in Section 3.7 that the Eddington limit does not play a large role for SMBHs in ROMULUS, we employ it in this section in order to treat SMBHs of all masses equally. To help mitigate AGN variability, we convolve all Eddington ratio time-series with a Gaussian kernel with a standard deviation of 30 Myr.

We plot the dependence of Eddington ratio on time since major merger in Fig. 11. Horizontal error bars indicate the spacing between saved simulations snapshots, since mergers can only be detected in between them. There is no evidence for a trend between the time since merger and Eddington ratio, regardless of the time-scale or the redshift. We also plot future mergers, in case there is a delay between merger-triggering and AGN activity. In particular, idealized galaxy merger simulations have found that AGN should be triggered at second pericentre passage, which may precede the point at which haloes are considered merged by the halo finder (Capelo et al. 2015; Volonteri et al. 2015a). We similarly find no correlation with Eddington ratio and the time to future merger.

Our results diminish the role of the intergalactic environment in determining the SMBH accretion rate. That said, the limited

resolution of these simulations, both spatially and temporally, may hamper our ability to detect a small-scale AGN-merger connection that might be critical to the fuelling problem. For example, our simulations would not be able to detect an AGN-merger connection if in reality it is driven by angular momentum losses on scales $\lesssim 250$ pc, or if AGNs rely on large quantities of cold, molecular gas which we do not track. We may also miss very brief bursts of merger-triggered AGN activity, which could be selected in a flux-limited survey. Finally, as mentioned in Section 3.3, ROMULUS25 lacks the volume to test the hypothesis that only the most luminous AGN are merger triggered. ULIRGs have $z \sim 2$ number densities of $\sim 10^{-5} \text{ Mpc}^{-3}$, for example Magnelli et al. (2011), the kind of abundance that is inaccessible in these simulations.

4 DISCUSSION

4.1 SMBH-galaxy co-evolution

The existence of a universal correlation between SFR and BHAR in the ROMULUS suite implies that SMBHs and their hosts tend to grow in tandem. The evolution of galaxies in ROMULUS25 on the $M_\bullet - M_*$ plane is illustrated in Fig. 12. Here, we begin with every galaxy at $z = 0.05$ that has a stellar mass of at least $5 \times 10^9 M_\odot$. For each galaxy, we track back the mass of its most massive SMBH and its stellar mass from one redshift to another, and plot the resulting line segment. For ease of readability, each segment is colour coded according to the angle it makes with the x -axis ($\tan \theta = \Delta \log M_\bullet / \Delta \log M_*$). The dashed black line corresponds to the relation reported by Kormendy & Ho (2013).

SMBHs which start off overmassive with respect to their hosts (due to the high seed mass) at high redshift do not overgrow. As we have shown in Section 3.4, their growth is capped according to their host's SFR, even if this results in low Eddington ratios. Note that a constant $\Delta M_\bullet / \Delta M_*$ does not imply a constant $\Delta \log M_\bullet / \Delta \log M_*$. Instead, an SMBH which is overmassive with respect to its host will grow at a lower Eddington ratio than one which is undermassive. This is because SMBH growth is limited by their host's gas supply rather than by its own mass (Pacucci et al. 2017). The result is that SMBHs and their hosts are attracted towards a line of constant




Figure 12. SMBH–galaxy co-evolution out to $z = 4$. The evolution of galaxies that have stellar masses above $5 \times 10^9 M_\odot$ at $z = 0$ are shown. The dashed black line displays the local M_b – M_* relation reported by Kormendy & Ho (2013). Line segments are colour coded according to the angle $\theta = \arctan(\Delta \log M_b / \Delta \log M_*)$. SMBHs and their hosts co-evolve, attracted towards a line of constant M_b/M_* .

SMBH-to-stellar mass ratio, moving along this line once they have reached it. Self-regulation via feedback might be the source for this confinement (Natarajan & Treister 2009). At the highest stellar masses, the SMBH growth may proceed in galaxies which have quenched their star formation. In addition, stripping causes some galaxies to exhibit backwards-facing gradients.

At low redshift, there are galaxies with low-mass SMBHs whose growth does not keep up with the SFR. This population is reflected in the low- M_b scatter in Fig. 2 as well as the low- M_b/M_* galaxies in Fig. 5. We will investigate the properties of these galaxies in future work.

4.2 Implications for future modelling

The ROMULUS simulations indicate that regardless of stellar mass, redshift, or intergalactic environment, the BHAR follows the SFR for star-forming galaxies. If the SMBH-to-stellar mass ratio and the cold gas fraction are included, the R^2 values of these models in Table 1 indicate that up to 66 or 68 per cent of the variations in the BHAR can be explained. However, while these correlations are significant, the underlying causal connection with these additional variables is still unclear. We have shown that AGN variability significantly increases the scatter about the mean relation, even on time-scales between 1.6 and 300 Myr. AGN ‘flicker’ on yet shorter time-scales may continue to increase the scatter.

At the same time, it is insufficient to universally prescribe SMBH growth to equal some fraction of the SFR plus some independent scatter. This has often been the assumption in several early

semi-analytic models that trace black hole mass assembly over cosmic time. ROMULUS provides two similar lines of evidence against such a simple picture. First, the M_b – M_* relation at $z = 0$ has substantial scatter at low masses, while under the simple picture all of the scatter would average out. Such SMBH accretion prescriptions were explored in the semi-analytic model of Ricarte & Natarajan (2018a), which indeed yielded little scatter in the M_b – σ relation at low masses. Secondly, the BHAR is shown to have a positive dependence on the SMBH-to-galaxy mass ratio (see Section 3.5). This means that the BHAR is autocorrelated on time-scales of at least 300 Myr. An SMBH that has grown efficiently in the past is more likely to continue to do so.

SMBHs in ROMULUS can grow efficiently in hosts of any stellar mass $> 10^8 M_\odot$. This is in contrast with many other similar cosmological simulations in which the growth of SMBHs is suppressed by supernova feedback. These differences may be due to any combination of resolution, accretion prescriptions, supernova feedback prescriptions, and AGN feedback prescriptions. One of the key differences in the ROMULUS simulations is a unique fuelling prescription in which rotational support suppresses the BHAR. We speculate that if for some reason the rotational support is higher in high-mass galaxies than in low-mass ones, recalibrating to the local M_b – M_* relations could promote more SMBH growth in low-mass galaxies. In addition, ROMULUS only contains one mode of AGN feedback (thermal), which may again alter calibrations. There is no observational evidence for a break in the M_b – M_* relation down to stellar masses of $10^9 M_\odot$ from broad-line relations, although this area of parameter space is sparsely sampled observationally at the

present time (Reines & Volonteri 2015). It is not clear whether ROMULUS or other simulations better represent the growth of BHs in low-mass galaxies.

We find no link between the BHAR and galaxy mergers, joining a body of other state-of-the-art cosmological simulations which have called into question their role in triggering AGNs (Martin et al. 2018; McAlpine et al. 2018; Steinborn et al. 2018). This work favours a picture in which secular processes steadily assemble both stars and SMBHs until the galaxy is quenched. Nevertheless, we caution that cosmological simulations currently lack the resolution to definitively settle this question. High-resolution studies reveal that the assembly histories of SMBHs can change dramatically as resolution is increased (Anglés-Alcázar et al. 2017). It is possible, for example, that AGNs are only triggered by extreme gas densities that cannot be resolved in current cosmological simulations that can only be feasibly provided by galaxy mergers. In addition, we cannot rule out the possibility that the rare, luminous quasars or the growth of the most massive SMBHs are triggered by mergers. Such rare objects occur too infrequently to be found in the volume of ROMULUS25, motivating future high-redshift zoom-in simulations to test this picture in more extreme environments. Although more galaxies are quenched in ROMULUSC than in ROMULUS25, the same AGN main-sequence appears in both environments. This suggests that the typical SMBH does not care about the intergalactic environment and is only affected by physical processes that are internal to the galaxy.

5 CONCLUSION

ROMULUS25 and ROMULUSC are state-of-the-art cosmological simulations that deploy novel models for SMBH seeding, accretion, and dynamics, making them ideal for probing the drivers of SMBH growth in realistic environments. The main results of our study are as follows:

- (i) In star-forming galaxies, the BHAR follows the SFR. There is no dependence on redshift, stellar mass, or even large-scale environment.
- (ii) The SMBH-to-stellar mass ratio and cold gas fraction are shown statistically to be secondary and tertiary parameters in determining the accretion rate for star-forming galaxies. By including these extra parameters, and by averaging over 300 Myr, up to 68 per cent of the variations in the BHAR can be explained. However, AGN variability on shorter timescales reduces instantaneously observable correlations.
- (iii) SMBHs in ROMULUS grow efficiently even in low-mass galaxies, in contrast with other similar cosmological simulations.
- (iv) The ROMULUS simulations do not exhibit any link between galaxy mergers and the SMBH accretion rate for the typical galaxy, although we cannot rule out a connection for rare transient objects such as ULIRGS or $z = 6$ quasars, which do not occur in the simulation's volume.
- (v) AGN are variable enough on timescales between 10 and 100 Myr to complicate studies of the link between the BHAR and local conditions in the host galaxy.

These results suggest that SMBHs simply consume a fraction of the gas which forms stars, regardless of large-scale environment or even the host's stellar mass. One major takeaway for demographic modelling is that setting the SMBH accretion rate to be proportional to the total SFR appears to be very well supported. Crucially, however, the degree to which SMBH accretion rates are above or below the average relationship is auto-correlated in time. The

significant scatter in the $M_{\bullet}-M_{\star}$ relationship at $z = 0$ implies that the scatter between BHAR and SFR does not simply average out with time. Rather, some as of yet unknown parameters favour the growth of SMBHs in some galaxies over others. In future work, we will further investigate the scatter about the mean relation: what drives some SMBHs to accrete more or less efficiently than the mean.

ACKNOWLEDGEMENTS

We thank our referee, Kastytis Zubovas, for his insightful comments that substantially improved the content of this paper. We thank Daisuke Nagai, Nico Cappelluti, Daniel Angles-Alcázar and Mila Chadayammuri for illuminating discussions. We thank Tonima Tasnim Ananna for providing the luminosity densities from her model. AR is supported by NASA Headquarters under the NASA Earth and Space Science Fellowship Program – Grant 80NSSC17K0459. MT gratefully acknowledges support from the YCAA Prize Postdoctoral Fellowship. TQ and MT were partially supported by NSF award AST-1514868. This research is part of the Blue Waters sustained petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. This work is also part of a Petascale Computing Resource Allocations allocation support by the National Science Foundation (award number OAC-1613674). This work also used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Resources supporting this work were also provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

REFERENCES

- Aird J., Coil A. L., Georgakakis A., Nandra K., Barro G., Pérez-González P. G., 2015, *MNRAS*, 451, 1892
- Aird J., Coil A. L., Georgakakis A., 2019, *MNRAS*, 484, 4360
- Anglés-Alcázar D., Faucher-Giguère C.-A., Quataert E., Hopkins P. F., Feldmann R., Torrey P., Wetzel A., Kereš D., 2017, *MNRAS*, 472, L109
- Azadi M. et al., 2015, *ApJ*, 806, 187
- Beifiori A., Courteau S., Corsini E. M., Zhu Y., 2012, *MNRAS*, 419, 2497
- Bluck A. F. L. et al., 2016, *MNRAS*, 462, 2559
- Bondi H., Hoyle F., 1944, *MNRAS*, 104, 273
- Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
- Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, *MNRAS*, 370, 645
- Bower R. G., Schaye J., Frenk C. S., Theuns T., Schaller M., Crain R. A., McAlpine S., 2017, *MNRAS*, 465, 32
- Cannon J. M. et al., 2011, *ApJ*, 739, L22
- Capelo P. R., Dotti M., 2017, *MNRAS*, 465, 2643
- Capelo P. R., Volonteri M., Dotti M., Bellovary J. M., Mayer L., Governato F., 2015, *MNRAS*, 447, 2123
- Cappelluti N., Allevato V., Finoguenov A., 2012, *Adv. Astron.*, 2012, 853701
- Champagne J. B. et al., 2018, *ApJ*, 867, 153
- Chandrasekhar S., 1943, *ApJ*, 97, 255
- Cisternas M. et al., 2011, *ApJ*, 726, 57
- Croton D. J. et al., 2006, *MNRAS*, 365, 11
- Dai L., McKinney J. C., Roth N., Ramirez-Ruiz E., Miller M. C., 2018, *ApJ*, 859, L20

- Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
- Di Matteo T., Croft R. A. C., Feng Y., Waters D., Wilkins S., 2017, *MNRAS*, 467, 4243
- Donley J. L. et al., 2018, *ApJ*, 853, 63
- Dubois Y., Volonteri M., Silk J., Devriendt J., Slyz A., Teyssier R., 2015, *MNRAS*, 452, 1502
- Eddington A. S., 1926, *The Internal Constitution of the Stars*, Cambridge University Press, Cambridge, England
- Fan X. et al., 2003, *AJ*, 125, 1649
- Fan L. et al., 2016, *ApJ*, 822, L32
- Ferrarese L., Merritt D., 2000, *ApJ*, 539, L9
- Fornasini F. M., Civano F., Fabbiano G., Elvis M., Marchesi S., Miyaji T., Zezas A., 2018, *ApJ*, 865, 43
- Gebhardt K. et al., 2000, *ApJ*, 539, L13
- Glikman E., Simmons B., Mailly M., Schawinski K., Urry C. M., Lacy M., 2015, *ApJ*, 806, 218
- Governato F. et al., 2015, *MNRAS*, 448, 792
- Guedes J., Callegari S., Madau P., Mayer L., 2011, *ApJ*, 742, 76
- Habouzit M., Volonteri M., Dubois Y., 2017, *MNRAS*, 468, 3935
- Habouzit M., Volonteri M., Somerville R. S., Dubois Y., Peirani S., Pichon C., Devriendt J., 2019, *MNRAS*, preprint ([arXiv:1810.11535](https://arxiv.org/abs/1810.11535))
- Haehnelt M. G., Natarajan P., Rees M. J., 1998, *MNRAS*, 300, 817
- Haiman Z., Loeb A., 2001, *ApJ*, 552, 459
- Haynes M. P. et al., 2011, *AJ*, 142, 170
- Hickox R. C., Alexander D. M., 2018, *ARA&A*, 56, 625
- Hickox R. C., Mullaney J. R., Alexander D. M., Chen C.-T. J., Civano F. M., Goulding A. D., Hainline K. N., 2014, *ApJ*, 782, 9
- Hirschmann M., Dolag K., Saro A., Bachmann L., Borgani S., Burkert A., 2014, *MNRAS*, 442, 2304
- Hopkins P. F., Quataert E., 2010, *MNRAS*, 407, 1529
- Hopkins P. F., Richards G. T., Hernquist L., 2007, *ApJ*, 654, 731
- Jahnke K., Macciò A. V., 2011, *ApJ*, 734, 92
- Jiang Y.-F., Stone J. M., Davis S. W., 2014, *ApJ*, 796, 106
- King A., 2003, *ApJ*, 596, L27
- King A., Nixon C., 2015, *MNRAS*, 453, L46
- King A., Pounds K., 2015, *ARA&A*, 53, 115
- Knollmann S. R., Knebe A., 2009, *ApJS*, 182, 608
- Kocevski D. D. et al., 2015, *ApJ*, 814, 104
- Kocevski D. D. et al., 2017, *ApJ*, 846, 112
- Kormendy J., Ho L. C., 2013, *ARA&A*, 51, 511
- Kormendy J., Richstone D., 1995, *ARA&A*, 33, 581
- Kroupa P., 2001, *MNRAS*, 322, 231
- LaMassa S. M. et al., 2015, *ApJ*, 800, 144
- Lanzuisi G. et al., 2017, *A&A*, 602, A123
- Lodato G., Natarajan P., 2007, *MNRAS*, 377, L64
- MacLeod C. L. et al., 2016, *MNRAS*, 457, 389
- Magnelli B., Elbaz D., Chary R. R., Dickinson M., Le Borgne D., Frayer D. T., Willmer C. N. A., 2011, *A&A*, 528, A35
- Magorrian J. et al., 1998, *AJ*, 115, 2285
- Martin G. et al., 2018, *MNRAS*, 476, 2801
- Martín-Navarro I., Mezcuca M., 2018, *ApJ*, 855, L20
- McAlpine S., Bower R. G., Harrison C. M., Crain R. A., Schaller M., Schaye J., Theuns T., 2017, *MNRAS*, 468, 3395
- McAlpine S., Bower R. G., Rosario D. J., Crain R. A., Schaye J., Theuns T., 2018, *MNRAS*, 481, 3118
- McKinney J. C., Dai L., Avara M. J., 2015, *MNRAS*, 454, L6
- Mechtle M. et al., 2016, *ApJ*, 830, 156
- Menon H., Wesolowski L., Zheng G., Jetley P., Kale L., Quinn T., Governato F., 2015, *Comput. Astrophys. Cosmol.*, 2, 1
- Mihos J. C., Hernquist L., 1996, *ApJ*, 464, 641
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Mullaney J. R. et al., 2012, *ApJ*, 753, L30
- Munshi F. et al., 2013, *ApJ*, 766, 56
- Natarajan P., Treister E., 2009, *MNRAS*, 393, 838
- Novak G. S., Ostriker J. P., Ciotti L., 2011, *ApJ*, 737, 26
- Obreschkow D., Glazebrook K., 2014, *ApJ*, 784, 26
- Pacucci F., Natarajan P., Volonteri M., Cappelluti N., Urry C. M., 2017, *ApJ*, 850, L42
- Peng C. Y., 2007, *ApJ*, 671, 1098
- Pezzulli E., Valiante R., Schneider R., 2016, *MNRAS*, 458, 3047
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Pontzen A., Tremmel M., 2018, *ApJS*, 237, 23
- Pontzen A., Roškar R., Stinson G., Woods R., 2013, *Astrophysics Source Code Library*, record ascl:1305.002
- Pontzen A., Tremmel M., Roth N., Peiris H. V., Saintonge A., Volonteri M., Quinn T., Governato F., 2017, *MNRAS*, 465, 547
- Powell M. C., Urry C. M., Cardamone C. N., Simmons B. D., Schawinski K., Young S., Kawakatsu M., 2017, *ApJ*, 835, 22
- Ramasawmy J., Stevens J., Martin G., Geach J. E., 2019, *MNRAS*, 486, 4320
- Reines A. E., Volonteri M., 2015, *ApJ*, 813, 82
- Ricarte A., Natarajan P., 2018a, *MNRAS*, 474, 1995
- Ricarte A., Natarajan P., 2018b, *MNRAS*, 481, 3278
- Ricci C. et al., 2017, *MNRAS*, 468, 1273
- Rosario D. J. et al., 2012, *A&A*, 545, A45
- Rosas-Guevara Y. M. et al., 2015, *MNRAS*, 454, 1038
- Saglia R. P. et al., 2016, *ApJ*, 818, 47
- Sanchez N. N., Werk J. K., Tremmel M., Pontzen A., Christensen C., Quinn T., Cruz A., 2018, preprint ([arXiv:1810.12319](https://arxiv.org/abs/1810.12319))
- Sartori L. F., Schawinski K., Trakhtenbrot B., Caplar N., Treister E., Koss M. J., Urry C. M., Zhang C. E., 2018, *MNRAS*, 476, L34
- Schawinski K., Koss M., Berney S., Sartori L. F., 2015, *MNRAS*, 451, 2517
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schramm M., Silverman J. D., 2013, *ApJ*, 767, 13
- Schwarz G., 1978, *Ann. Stat.*, 6, 461
- Shankar F. et al., 2016, *MNRAS*, 460, 3119
- Shen S., Wadsley J., Stinson G., 2010, *MNRAS*, 407, 1581
- Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, 452, 575
- Soltan A., 1982, *MNRAS*, 200, 115
- Springel V., Di Matteo T., Hernquist L., 2005, *MNRAS*, 361, 776
- Stanley F., Harrison C. M., Alexander D. M., Swinbank A. M., Aird J. A., Del Moro A., Hickox R. C., Mullaney J. R., 2015, *MNRAS*, 453, 591
- Steinborn L. K., Hirschmann M., Dolag K., Shankar F., Juneau S., Krumpe M., Remus R.-S., Teklu A. F., 2018, *MNRAS*, 481, 341
- Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, *MNRAS*, 373, 1074
- Tasnim Ananna T. et al., 2019, *ApJ*, 871, 240
- Trakhtenbrot B., Lira P., Netzer H., Ciccone C., Maiolino R., Shemmer O., 2017, *ApJ*, 836, 8
- Trakhtenbrot B., Lira P., Netzer H., Ciccone C., Maiolino R., Shemmer O., 2018, preprint ([arXiv:1801.01508](https://arxiv.org/abs/1801.01508))
- Treister E., Natarajan P., Sanders D. B., Urry C. M., Schawinski K., Kartaltepe J., 2010, *Science*, 328, 600
- Treister E., Schawinski K., Urry C. M., Simmons B. D., 2012, *ApJ*, 758, L39
- Tremmel M. et al., 2019, *MNRAS*, 483, 3336
- Tremmel M., Governato F., Volonteri M., Quinn T. R., 2015, *MNRAS*, 451, 1868
- Tremmel M., Governato F., Volonteri M., Quinn T. R., Pontzen A., 2017, *MNRAS*, 470, 1121
- Trump J. R. et al., 2011, *ApJ*, 733, 60
- Uchiyama H. et al., 2018, *PASJ*, 70, S32
- Ueda Y., Akiyama M., Hasinger G., Miyaji T., Watson M. G., 2014, *ApJ*, 786, 104
- van den Bosch R. C. E., 2016, *ApJ*, 831, 134
- Villforth C. et al., 2017, *MNRAS*, 466, 812
- Vito F. et al., 2018, *MNRAS*, 473, 2378
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Volonteri M., Natarajan P., Gültekin K., 2011, *ApJ*, 737, 50
- Volonteri M., Capelo P. R., Netzer H., Bellovary J., Dotti M., Governato F., 2015a, *MNRAS*, 452, L6
- Volonteri M., Silk J., Dubus G., 2015b, *ApJ*, 804, 148

- Volonteri M., Dubois Y., Pichon C., Devriendt J., 2016, *MNRAS*, 460, 2979
- Wadsley J. W., Veeravalli G., Couchman H. M. P., 2008, *MNRAS*, 387, 427
- Wadsley J. W., Keller B. W., Quinn T. R., 2017, *MNRAS*, 471, 2357
- Whitaker K. E., van Dokkum P. G., Brammer G., Franx M., 2012, *ApJ*, 754, L29
- Wu X.-B. et al., 2015, *Nature*, 518, 512
- Yang G. et al., 2017, *ApJ*, 842, 72
- Yang G. et al., 2018, *MNRAS*, 475, 1887
- Yang G., Brandt W. N., Alexander D. M., Chen C.-T. J., Ni Q., Vito F., Zhu F.-F., 2019, *MNRAS*, 485, 3721
- Zubovas K., King A. R., 2012, *MNRAS*, 426, 2751

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