

Cosmic evolution of stellar quenching by AGN feedback: clues from the Horizon-AGN simulation

R. S. Beckmann,¹★ J. Devriendt,^{1,2} A. Slyz,¹ S. Peirani,^{3,4} M. L. A. Richardson,¹
Y. Dubois,^{3,4} C. Pichon,^{3,4,5} N. E. Chisari,¹ S. Kaviraj,⁶ C. Laigle¹ and M. Volonteri^{3,4}

¹Sub-department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

²Observatoire de Lyon, UMR 5574, 9 avenue Charles Andre, F-69561 Saint Genis Laval, France

³Sorbonne Universites, UPMC Univ Paris 06, UMR 7095, Institut d'Astrophysique de Paris, F-75005 Paris, France

⁴CNRS, UMR 7095, Institut d'Astrophysique de Paris, 98bis Boulevard Arago, F-75014 Paris, France

⁵Korea Institute for Advanced Study (KIAS), 85 Hoegiro, Dongdaemun-gu, Seoul 02455, Republic of Korea

⁶Centre for Astrophysics Research, University of Hertfordshire, College Lane, Hatfield, Herts AL10 9AB, UK

Accepted 2017 July 18. Received 2017 June 19; in original form 2017 January 24

ABSTRACT

The observed massive end of the galaxy stellar mass function is steeper than its predicted dark matter halo counterpart in the standard Λ cold dark matter paradigm. In this paper, we investigate the impact of active galactic nuclei (AGN) feedback on star formation in massive galaxies. We isolate the impact of AGN by comparing two simulations from the HORIZON suite, which are identical except that one also includes supermassive black holes (SMBHs) and related feedback models. This allows us to cross-identify individual galaxies between simulations and quantify the effect of AGN feedback on their properties, including stellar mass and gas outflows. We find that massive galaxies ($M_* \geq 10^{11} M_\odot$) are quenched by AGN feedback to the extent that their stellar masses decrease by up to 80 per cent at $z = 0$. SMBHs affect their host halo through a combination of outflows that reduce their baryonic mass, particularly for galaxies in the mass range $10^9 M_\odot \leq M_* \leq 10^{11} M_\odot$, and a disruption of central gas inflows, which limits *in situ* star formation. As a result, net gas inflows on to massive galaxies, $M_* \geq 10^{11} M_\odot$, drop by up to 70 per cent. We measure a redshift evolution in the stellar mass ratio of twin galaxies with and without AGN feedback, with galaxies of a given stellar mass showing stronger signs of quenching earlier on. This evolution is driven by a progressive flattening of the $M_{\text{SMBH}}-M_*$ relation with redshift, particularly for galaxies with $M_* \leq 10^{10} M_\odot$. M_{SMBH}/M_* ratios decrease over time, as falling average gas densities in galaxies curb SMBH growth.

Key words: methods: numerical – galaxies: active – galaxies: evolution – galaxies: high-redshift – quasars: supermassive black holes – galaxies: star formation.

1 INTRODUCTION

It has long been known that the hierarchical structure formation paradigm implied by the cold dark matter (CDM) model, whilst very successful overall, overproduces objects at the bright and faint end of the luminosity function (White & Frenk 1991; Kauffmann, White & Guiderdoni 1993). Observations show much more inefficient star formation in low- and high-mass haloes, with a peak in efficiency at the luminosity turnover (Bernardi et al. 2013; Moustakas et al. 2013; Davidzon et al. 2017). To avoid the overcooling prob-

lem and reproduce the observed luminosity function, an energetic feedback mechanism is required (Cole et al. 1994; Binney & Tabor 1995; Silk & Rees 1998; Springel, Di Matteo & Hernquist 2005; Dubois et al. 2013). Several different avenues have been suggested to provide the necessary energy input and quench star formation, including the extragalactic UV background, supernova (SN) feedback and feedback due to active galactic nuclei (AGN).

Photoionization by the extragalactic UV background and the first generations of stars suppresses gas accretion at high redshift, causing a number of smaller dark matter (DM) haloes to remain devoid of gas (Gnedin 2000; Somerville 2002; Okamoto, Gao & Theuns 2008; Geen, Slyz & Devriendt 2013). Whilst this mechanism provides a possible solution to the overabundance of very

* E-mail: ricarda.beckmann@physics.ox.ac.uk

low mass ($\leq 10^8 M_\odot$) substructures in Milky Way like haloes, it has been shown to have little effect on more massive objects that collapse later (Efstathiou 2000; Governato et al. 2007).

On the other hand, SN feedback is widely believed to play an important part in quenching star formation in haloes with masses below $10^{11} M_\odot$ (Dekel & Silk 1986; Benson et al. 2003; Hopkins et al. 2014). Their shallow potentials allow SN-driven winds with velocities comparable to the escape velocity (Mac Low & Ferrara 1999; Keller, Wadsley & Couchman 2015; Christensen et al. 2016) to empty the host galaxy of a significant amount of gas, thereby efficiently suppressing star formation. However, even stellar feedback concentrated in intense, compact starbursts caused by major mergers or violent disc instabilities cannot quench massive galaxies (Dekel & Burkert 2014).

It has been suggested that star formation halts in massive objects due to a slowdown in cold flows at low redshift (Feldmann & Mayer 2014). However, Mandelbaum et al. (2016) show that massive quiescent galaxies can have twice as much DM as star-forming galaxies, indicating that cosmic inflows probably continue long after star formation has ceased. Furthermore, the slowdown in cold flows is expected to take place over Gyr time-scales (Zolotov et al. 2015), which contradicts observational evidence that quenching of massive galaxies takes place on much shorter time-scales (Barro et al. 2016; Tomczak et al. 2016). Therefore, the reduction of star formation in massive galaxies is unlikely to occur solely because accretion is fading away.

Instead, AGN feedback could provide an effective route to quenching massive galaxies, as well as regulating the growth of supermassive black holes (SMBHs; Binney & Tabor 1995; Benson et al. 2003; Di Matteo, Springel & Hernquist 2005; Bower et al. 2006; Croton et al. 2006; Sijacki et al. 2007; Cattaneo et al. 2009; Fabian 2012). There are two main mechanisms through which this could proceed. One possibility is that black holes and stars feed from the same cold gas supply until it is depleted by AGN feedback, at which point both processes come to a halt. Observations of galaxies that simultaneously harbour both an AGN and an active starburst provide evidence that supports this claim (Heckman & Kauffmann 2006; Lutz et al. 2008; Nardini et al. 2008). In this scenario, AGN feedback might even accelerate star formation by further compressing the cold gas of the galaxy in a so-called positive feedback mode (Silk 2005; Gaibler et al. 2012; Santini et al. 2012; Bieri et al. 2015). However, its main role is to prevent the gas heated/expelled by SN winds to be re-accreted at a later stage, alongside more pristine material. This is the so-called maintenance mode, associated with powerful radio emission (Rafferty et al. 2006; Cattaneo & Teyssier 2007).

Alternatively, AGN feedback could act directly on the gas content of the galaxy. It could expel the interstellar medium (ISM) out of galaxies in massive galactic winds and/or prevent star formation by directly heating the ISM gas (Di Matteo et al. 2005; Murray, Quataert & Thompson 2005; Springel, Di Matteo & Hernquist 2005; Fabian, Celotti & Erlund 2006). This view is supported by observational evidence of frequent and fast outflows in massive galaxies (Tremonti, Moustakas & Diamond-Stanic 2007), able to drive a significant gas mass (Heckman et al. 2000; Veilleux, Cecil & Bland-Hawthorn 2005; Weiner et al. 2009; Sturm et al. 2011) using only 5–10 per cent of accretion power (Moe et al. 2009; Saez, Chartas & Brandt 2009; Dunn et al. 2010). Whilst such outflows are common in AGN with powerful radio jets (Arav et al. 2001; De Kool et al. 2001; Reeves, O’Brien & Ward 2003; Chartas et al. 2007; Moe et al. 2009; Dunn et al. 2010), careful anal-

ysis of higher redshift objects provides evidence that quasars can also launch powerful energy driven winds and thus cause a rapid star formation decline (Farrah et al. 2012; Maiolino et al. 2012; Page et al. 2012; Ciccone et al. 2014; Costa, Sijacki & Haehnelt 2015; Williams et al. 2017). Both modes of AGN feedback can reproduce the observed correlations between host galaxy and BH properties (Ferrarese & Merritt 2000; Gebhardt et al. 2000), as shown by Di Matteo et al. (2008), Booth & Schaye (2009), Dubois et al. (2012), Sijacki et al. (2015) and Volonteri et al. (2016). However, the time-scale over which quenching takes place is still a matter of debate, with evidence existing for both rapid quenching (Thomas et al. 2005) and much slower processes (Quintero et al. 2004). The time-scales are probably dependent on galaxy type (Schawinski et al. 2014).

In this work, we use state-of-the-art cosmological simulations to investigate when and how AGN feedback affects its host galaxy. We isolate the impact of such feedback on stellar masses and large-scale gas flows by comparing the evolution of a statistically representative sample of individual objects, identifying matching galaxies in two simulations, HORIZON-AGN (H-AGN) and HORIZON-noAGN (H-noAGN). As their names indicate, these simulations are identical in all aspects except one is run with and the other without AGN feedback. Following the evolution of twinned galaxies from redshift $z = 5$ down to $z = 0$ allows us to determine the epoch of quenching, and identify ensuing changes in the stellar masses of affected galaxies.

The paper is structured as follows. Section 2 briefly introduces the HORIZON simulation suite and Section 3 explains the procedure used to identify pairs of corresponding objects across both simulations. Section 4 presents the effect of AGN feedback on galaxy stellar masses throughout cosmic time and Section 5 determines the causes for the measured quenching by studying the evolution of the black hole population, the gas content of haloes and galaxies and gas inflow/outflow rates. Section 6 summarizes and discusses our results.

2 THE SIMULATIONS

This paper presents a comparative analysis of two simulations: HORIZON-AGN (H-AGN) and HORIZON-noAGN (H-noAGN). Both simulations are run from identical initial conditions and share the same technical specifications and implementations of physics. The only difference is that HORIZON-AGN also includes a subgrid modelling of SMBHs and the associated AGN feedback (see Section 2.5) whereas H-noAGN does not. More details can be found in Dubois et al. (2014).

2.1 Cosmology and initial conditions

Both simulations were run with RAMSES (Teyssier 2002), an adaptive mesh refinement code, using a second-order unsplit Godunov scheme to solve the Euler equations. An HLLC Riemann solver with a MinMod Total Variation Diminishing scheme was used to reconstruct interpolated variables. Initial conditions were produced using MP-GRAFC (Prunet et al. 2008) and both simulations were carried out until $z = 0.0$.

The initial conditions set-up is a standard Λ CDM cosmology consistent with the WMAP-7 data (Komatsu et al. 2011), with $\Omega_m = 0.272$, dark energy density $\Omega_\Lambda = 0.728$, baryon density $\Omega_b = 0.045$, Hubble constant $H_0 = 70.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$, amplitude of the matter power spectrum $\sigma_8 = 0.81$ and power-law index of the primordial power spectrum $n_s = 0.967$.

The simulated cube of $L_{\text{box}} = 100 \text{ Mpc } h^{-1}$ on a side is initially refined uniformly down to physical $\Delta x = 1 \text{ kpc}$, which requires a root grid with 1024^3 cells at $z \simeq 100$. Extra refinement levels are continuously triggered using a quasi-Lagrangian criterion: a grid cell is split into eight whenever its DM or baryonic mass exceeds eight times the initial DM or baryonic mass respectively. To keep the size of the smallest cells approximately constant in physical units, cells on the most refined level are split into eight every time the expansion factor doubles, if they fulfill the refinement criteria.

There are a total of 1024^3 DM particles in each simulation, leading to a DM mass resolution of $8 \times 10^7 M_\odot$. All collisionless particles, i.e. DM and stars, are evolved using a multigrid Poisson solver with a cloud-in-cell interpolation to assign particles to grid cells.

2.2 Cooling and heating

The gas is allowed to cool down to 10^4 K using H, He and atomic metal cooling, following Sutherland & Dopita (1993) and accounting for photon heating by a uniform UV background (Haardt & Madau 1996) from $z_{\text{reion}} = 10$ onwards. The ratio between elements is assumed to be solar in these cooling/heating calculations. The gas follows a mono-atomic equation of state, with adiabatic index $\gamma = 5/3$.

2.3 Stars and SNe

Star formation is modelled using a Kennicutt–Schmidt law $\dot{\rho}_* = \epsilon_* \rho / t_{\text{ff}}$, where $\dot{\rho}_*$ is the star formation rate (SFR) density, $\epsilon_* = 0.02$ is the (constant) star formation efficiency (Kennicutt 1998; Krumholz & Tan 2007), ρ is the gas density and t_{ff} is the local free fall time of the gas. Stars form when the gas number density exceeds $\rho_0 / (\mu m_H) = 0.1 \text{ H/cm}^3$, where m_H is the mass of a hydrogen atom and μ is the mean molecular weight, and star particles are generated according to a Poisson random process (Rasera & Teyssier 2006; Dubois & Teyssier 2008) with a stellar mass resolution of $M_* \simeq 2 \times 10^6 M_\odot$, kept constant throughout the simulation. To avoid numerical fragmentation and mimic the effect of stellar heating by young stars, a polytropic equation of state, $T = T_0 (\rho / \rho_0)^\kappa$, is used for gas above the star formation density threshold, with $\kappa = 4/3$.

For stellar feedback, a Salpeter initial mass function (Salpeter 1954) is assumed, with low- and high-mass cut-offs of $0.1 M_\odot$ and $100 M_\odot$, respectively. In an effort to account for stellar feedback as comprehensively as possible, the (subgrid) model implemented in this work includes stellar winds, Type II and Type Ia SNe. Mechanical feedback energies from Type II SNe and stellar winds are computed using STARBURST99 (Leitherer et al. 1999, 2010). Specifically, we use a Padova model (Girardi et al. 2000) with thermally pulsating asymptotic branch stars (Vassiliadis & Wood 1993), and stellar winds are calculated as in Leitherer, Robert & Drissen (1992). The frequency of Type I SN is estimated from Matteucci & Greggio (1986), assuming a binary fraction of 5 per cent. To reduce computational costs, stellar feedback is modelled as a source of kinetic energy during the first 50 Myr of the lifetime of star particles, and as a heat source after that. On top of energy, mass and metals injected into the interstellar medium (ISM) by stellar feedback, we also keep track a variety of chemical elements (O, Fe, C, N, Mg, Si) synthesized in stars, with stellar yields estimated according to the W7 model of Nomoto et al. (2007). More detailed discussions of

the stellar feedback model used can be found in Dubois & Teyssier (2008), Kimm et al. (2015) and Rosdahl et al. (2017).

2.4 SMBH formation and accretion

In H-AGN, black holes are seeded with an initial mass of $10^5 M_\odot$ in dense, star-forming regions, i.e. when a gas cell exceeds $\rho > \rho_0$ and is Jeans unstable, provided such regions are located more than 50 kpc away from a pre-existing black hole (Dubois et al. 2010). These black holes subsequently accrete gas at the Bondi–Hoyle–Lyttleton rate:

$$\dot{M}_{\text{BH}} = \frac{4\pi\alpha G^2 M_{\text{BH}}^2 \bar{\rho}}{(\bar{c}_s^2 + \bar{u}^2)^{3/2}}, \quad (1)$$

where M_{BH} is the black hole mass, $\bar{\rho}$ is the average gas density, \bar{c}_s is the average sound speed and \bar{u} is the average gas velocity relative to the BH. We emphasize that this model is based on rather crude assumptions about the hydrodynamical processes undergone by gas surrounding an *extremely small* accretor. In particular, it does not take possible density or velocity gradients on the scale of the (unresolved) accretion radius into account, nor does it provide any description of important hydrodynamical instabilities which develop on yet smaller scales (see e.g. Foglizzo, Galletti & Ruffert 2005). To somewhat mitigate resolution effects that make it difficult to capture cold, dense regions of the ISM, a boost factor α is used, following Booth & Schaye (2009) and Dubois et al. (2012):

$$\alpha = \begin{cases} (\rho / \rho_0)^2 & \text{if } \rho > \rho_0 \\ 1 & \text{otherwise.} \end{cases} \quad (2)$$

Accretion is capped at the Eddington rate

$$\dot{M}_{\text{Edd}} = 4\pi G M_{\text{BH}} m_p / (\epsilon_r \sigma_T c), \quad (3)$$

where σ_T is the Thompson cross-section, m_p is the proton mass and c is the speed of light. A standard radiative efficiency, typical of a Shakura–Sunyaev accretion disc around the BHs, of $\epsilon_r = 0.1$ is assumed (Shakura & Sunyaev 1973).

BHs are also allowed to merge with one another when they are closer than 4 kpc , and their relative velocity is smaller than the escape velocity of the binary.

2.5 AGN feedback

Two modes of AGN feedback are implemented in H-AGN, depending on the instantaneous accretion rate of the SMBH: the so-called radio and quasar modes.

At high accretion rates, i.e. for Eddington ratios $\chi = \dot{M}_{\text{BH}} / \dot{M}_{\text{Edd}} > 0.01$, the quasar mode deposits thermal energy isotropically into a sphere of radius Δx centred on the BH. This energy is deposited with an efficiency of $\epsilon_f = 0.15$ at a rate of

$$\dot{E}_{\text{AGN}} = \epsilon_f \epsilon_r \dot{M}_{\text{BH}} c^2. \quad (4)$$

The radio mode takes over at low accretion rates, $\chi = \dot{M}_{\text{BH}} / \dot{M}_{\text{Edd}} \leq 0.01$, and deposits kinetic energy into bipolar outflows with jet/wind velocities of 10^4 km s^{-1} , along an axis aligned with the angular momentum of the accreted material, following the model of Omma et al. (2004). The total rate of energy deposited is given by the previous equation for \dot{E}_{AGN} , albeit using a higher efficiency of $\epsilon_f = 1$ (see Dubois et al. 2010 for detail).

The radiative efficiencies of the two AGN feedback modes were chosen to reproduce the scaling relations between BH mass and galactic properties in the local Universe, $M_{\text{SMBH}} - M_*$ and $M_{\text{SMBH}} - \sigma_*$.

Table 1. Number of galaxy twins in each subsample, at the range of redshifts presented here.

z	Number of twins per mass category		
	Small $M_* \leq 10^9 M_\odot$	Medium $10^9 M_\odot < M_* < 10^{11} M_\odot$	Large $10^{11} M_\odot \leq M_*$
0.1	9508	51 652	2060
1	13 393	60 547	773
3	11 823	19 855	7
5	2 949	1 626	0

(Dubois et al. 2012). More generally, we refer the reader interested in the details of how accretion on to SMBHs, and subsequent feedback injection, depend on numerical resolution and subgrid model parameter choices to Dubois et al. (2012). Note that the two radiation efficiency parameters previously mentioned are the only ones which are tuned in the HORIZON simulations, in the sense that the other parameters (associated with the subgrid models of star formation and stellar feedback) were not allowed to vary in order to obtain a better match to bulk galaxy properties. For instance, even though our star formation efficiency choice ensures that galaxies will fall on the Kennicutt observational law by construction, it does not automatically guarantee that they will have the correct stellar/gas mass and/or size at any epoch.

2.6 Mass categories for galaxies

To facilitate the presentation of our results, we split our sample of galaxy twins into three subsamples, distinguished by the stellar mass of the H-AGN galaxy. We define *small* galaxies as twins with stellar masses $M_*^{\text{H-AGN}} < 10^9 M_\odot$ in H-AGN, *medium* galaxies as those with $10^9 M_\odot \leq M_*^{\text{H-AGN}} \leq 10^{11} M_\odot$ and *large* galaxies as those with $M_*^{\text{H-AGN}} > 10^{11} M_\odot$. See Table 1 for the number of twins in each mass category. As a visual guidance, these mass categories will be annotated by solid vertical lines on all relevant plots.

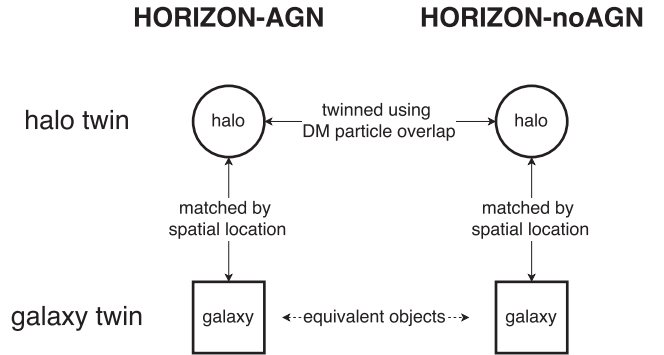
2.7 Nomenclature

For the purpose of this paper, ‘quenching’ refers to any reduction in galaxy SFRs when AGN feedback is included, compared to the case without it, not just a redshift-dependent specific SFR threshold of $0.3/t_{\text{Hubble}}$, as defined in Franx et al. (2008). The ‘quenching mass ratio’ refers to the stellar mass ratio of galaxy twins between the cases with and without AGN feedback, i.e. $M_*^{\text{H-AGN}}/M_*^{\text{H-noAGN}}$.

3 HALO MATCHING ACROSS SIMULATIONS

3.1 The twinning procedure

As usual, the first step consists in detecting objects of interest (haloes, subhaloes and galaxies) in each simulation, using the ADAP-TAHOP (sub)halofinder (Aubert, Pichon & Colombi 2004; Tweed et al. 2009). Having two simulations based on identical initial conditions allows the identification of corresponding objects between the two simulations, a procedure here referred to as ‘twinning’ (see also e.g. Geen et al. 2013). A pair of corresponding objects is called a twin and identifies two objects (one in each simulation) that have grown from the same overdensity in the initial conditions. These objects can either be (sub)haloes, for (sub)halo twins, or galaxies, for galaxy twins. If all the algorithms implemented to describe physical processes were identical in both simulations, the twins would be identical except for minor differences introduced by the stochastic

**Figure 1.** To twin galaxies across H-AGN and H-noAGN, DM (sub)haloes are first identified in each simulation, and (sub)haloes that show a 75 per cent or more overlap in the identities of their DM particles are twinned. Galaxies are then associated with a unique host (sub)halo in each simulation, based on how close they are to the centre of the (sub)halo (see the text for detail). Galaxies associated with (sub)halo twins are considered galaxy twins.

nature of the star formation algorithm. However, such seemingly innocuous differences would already prevent us from directly twinning galaxies. We therefore employ the more general method of Peirani et al. (2016) to perform this task, which is summarized in Fig. 1. Only DM (sub)haloes are twinned directly; to create galaxy twins, each galaxy is first associated with a host (sub)halo in the same simulation, before being twinned to the galaxy hosted by this (sub)halo’s twin in the other simulation.

More specifically, as both simulations start from identical initial conditions, with uniquely identified DM particles in identical positions, we can identify which of these particles cluster to form gravitationally bound (sub)haloes as the runs proceed. (Sub)haloes that grew from the same initial overdensities in both simulations should contain a large fraction of DM particles with identical identities at any time. In practice, for two DM (sub)haloes to be twinned, we require that at least 75 per cent of the DM particles present in a (sub)halo in H-AGN are also present in the H-noAGN (sub)halo. Note that in some cases, this choice will lead to a single (sub)halo in H-noAGN being associated with several (sub)haloes in H-AGN, as (sub)haloes mergers lead to the formation of (sub)subhaloes that are not necessarily disrupted at the same time in both simulations. In these cases, the object with the most similar mass is chosen as the twin (sub)halo, and the other matches are discarded.

Star particles stochastically form over the course of a simulation and their identifiers therefore reflect the detailed star formation history of that precise simulation, so it is not possible to identify galaxy twins directly through their star particle identities as is done for DM (sub)haloes. Instead, galaxies are considered to be twins if they are located within DM (sub)halo twins (see Fig. 1). We therefore begin by assigning a host (sub)halo to each galaxy if its centre is located within a distance $R_{\text{host}} = 0.05 \times R_{\text{vir}}$ of the centre of the (sub)halo. The centres of these (sub)haloes are computed using a shrinking sphere method (Power et al. 2003), and their precise location corresponds to the position of the most dense DM particle located in the final sphere identified with this method (Peirani et al. 2016). In case a (sub)halo contains more than one galaxy in its central region, we select the most massive one as being hosted by this (sub)halo and discard other matching objects. Note that proceeding in this way biases our results against so-called ‘orphan’ galaxies, i.e. galaxies whose host (sub)halo has been disrupted to the point that it falls below the particle detection threshold used by our halofinder. However, such orphans are quite rare (less than 1 per cent of our

Table 2. Number of galaxies, (sub)haloes and twins identified in H-AGN and H-noAGN. The sample includes all (sub)haloes resolved by at least 500 DM particles, and all galaxies resolved by at least 100 star particles, hosted by these (sub)haloes.

z	$N^{\text{H-AGN}}$	$N^{\text{H-noAGN}}$	$N^{\text{H-noAGN}}/N^{\text{H-AGN}}$	N_{twinned}
Number of (sub)haloes				
0.1	88 171	85 141	0.97	75 544
1	89 930	88 488	1.02	83 176
3	53 363	53 551	1.00	53 357
5	16 042	16 065	1.00	15 818
Number of galaxies				
0.1	75 651	73 450	0.97	63 220
1	80 588	82 238	1.02	72 713
3	34 128	47 428	1.39	31 685
5	6255	7757	1.24	4575

sample at any redshift) and almost exclusively belong to the category of small galaxies ($M_* < 10^9 M_\odot$), so our conclusions are unaffected by this bias. We also find that relying on DM (sub)halo host twinning to twin galaxies, rather than directly measuring the galaxy orbital properties, is a more robust process, as orbital parameters are sensitive to both internal changes in galaxy properties (in particular stellar mass) and host (sub)halo density profiles, which can differ quite significantly between simulations with and without AGN feedback (see e.g. Peirani et al. 2016).

3.2 Matched fractions

As we are interested in the most massive objects, the full sample considered here includes all (sub)haloes resolved by at least 500 DM particles, and all galaxies hosted in these (sub)haloes that contain at least 100 star particles. Note that this latter criterion, contrary to the one for DM (sub)haloes, does not correspond to a strict stellar mass threshold as star particles can have different masses (integer multiples of the minimal stellar mass). Across all redshifts, H-AGN and H-noAGN contain a comparable, albeit not identical number of both galaxies and (sub)haloes (see Table 2). For this reason, the following analysis considers the H-AGN sample as the reference to analyse the effectiveness of the twinning algorithm.

Fig. 2 shows that at high redshift, over 98 per cent of (sub)haloes present in H-AGN, corresponding to 15 818 out of 16 042 at redshift $z = 5$, are twinned successfully, with an even distribution across all mass bins. At lower redshift, the fraction of matched (sub)haloes decreases as the more and more different merger histories in the two simulations introduce larger discrepancies between individual objects. This fraction also decreases with (sub)halo mass, as (sub)haloes with smaller particle numbers are more sensitive to these merger history changes. However, with 75 544 (sub)haloes twinned at redshift $z = 0.1$ out of a sample of 88 171, i.e. an average matched fraction of 86 per cent, the resulting sample remains statistically representative.

The overall rates for matched galaxies are much lower than for (sub)haloes, as they require three steps to establish the twin link (galaxy to host (sub)halo, host (sub)halo to host (sub)halo twin, host (sub)halo twin to galaxy twin, see Fig. 1), with a number of objects dropping out of the sample at each step. Identifying a galaxy with its host (sub)halo can be challenging, especially in dense environments and at high redshift where interactions are more common. For example, increasing the size of the region within which a galaxy is associated with its host (sub)halo, R_{host} , from $R_{\text{host}} = 0.05 \times R_{\text{vir}}$




Figure 2. Fractions of matched (sub)haloes (solid lines) and galaxies (dashed lines) in the H-AGN and H-noAGN simulations. The dashed-dotted line represents the ideal case where all objects in H-AGN are uniquely associated with a corresponding object in H-noAGN. More than 86 per cent of (sub)haloes are twinned, but galaxy twinning is less efficient, particularly at stellar masses below $M_*^{\text{H-AGN}} = 10^9 M_\odot$ for all redshifts but also, albeit less significantly, at redshift $z \leq 1$ for galaxies with higher stellar masses.

to $R_{\text{host}} = 0.10 \times R_{\text{vir}}$, increases the fraction of matched galaxies at high redshift (see e.g. Chisari et al. 2017). More specifically, for redshift $z = 3$, the total number of galaxies with at least 100 star particles identified with a host (sub)halo with the same number of DM particles would rise from 47 656 to 67 301, out of 76 887 galaxies in total, i.e. an increase in matched fraction from 61 per cent to 87 per cent. Selecting amongst these galaxies those hosted in (sub)haloes containing more than 500 DM particles (as we do in this work), further reduces the numbers to 34 128 (given in Table 2) and 48 354 galaxies, respectively, out of 76 887. Whilst the number of galaxies excluded by the strict position and mass criteria we employ for twinning does represent a significant fraction of the sample, especially at high redshift, we have checked that relaxing them hardly alters the quantitative results presented in this work. This can be intuitively understood as the mass cuts chosen only eliminate low-mass galaxies from the sample. Low-mass galaxies are both the most numerous and virtually unaffected by AGN feedback at any redshift. This can be seen in Fig. 3, where the entire population (i.e. all 76 887 objects for H-AGN at $z = 3$), as opposed to only twinned galaxies, is plotted. Similarly, the use of a stricter position criterion mostly affects low-mass galaxies, as these have longer dynamical friction times and are more easily dislodged from the centre of their host (sub)haloes during gravitational interactions. In short, the galaxy sample as defined in this section is statistically robust enough, at all redshifts and galaxy masses, for us to draw conclusions about the impact of AGN feedback in the H-AGN simulation, whilst allowing us to get rid of virtually all mismatch errors in the galaxy twinning process.

3.3 The effect of AGN feedback on haloes

Fig. 3 shows that no matter the redshift, the DM (sub)halo mass functions (HMF) for H-AGN and H-noAGN are so similar that they are indistinguishable on the plot. Directly comparing the DM masses of (sub)halo twins (see Fig. 4) shows that (sub)haloes with masses below $M_{\text{vir}}^{\text{H-AGN}} < 10^{11} M_\odot$ have identical masses in both

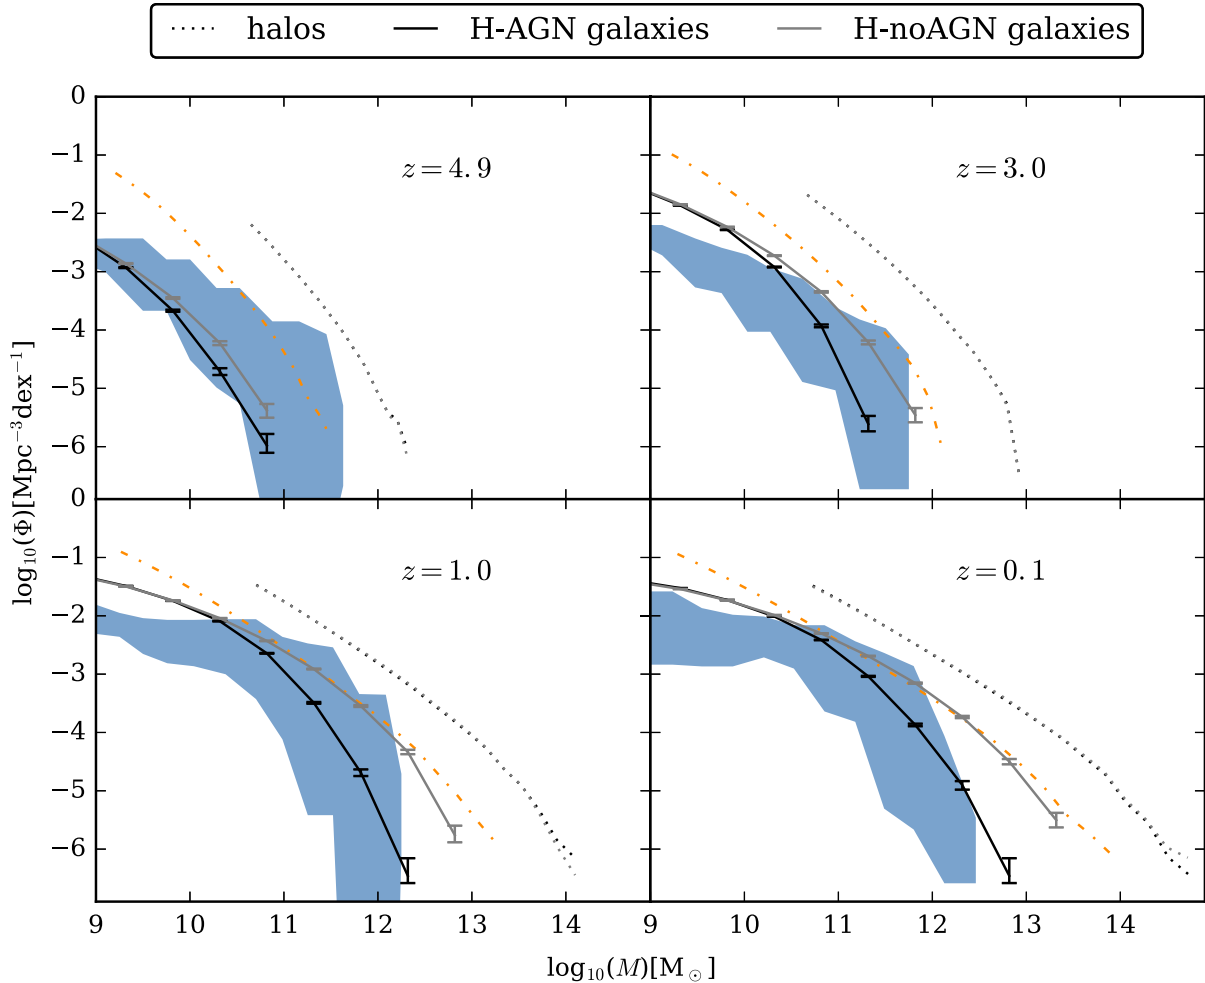


Figure 3. DM halo mass functions (HMFs; dotted) and corresponding galaxy stellar mass functions (GSMFs; solid) for H-AGN (black) and H-noAGN (light grey). We also plot combined observations from González et al. (2011), Bielby et al. (2014), Daddi et al. (2002), Bernardi et al. (2013), Moustakas et al. (2013), Tomczak et al. (2014), Song et al. (2016) and Davidzon et al. (2017) as the shaded area, and the GSMF expected if all the baryons within (sub)haloes (using a fraction of 16.5 per cent of the total mass) are converted into stars (dot-dashed line). A detailed comparison of individual observational data sets to the GSMF in H-AGN can be found in Kaviraj et al. (2017). Error bars on the simulated mass functions are Poisson, with those on the HMFs are not plotted as they are even smaller than that of the GSMF ones in the mass range where they overlap. HMFs for both H-AGN and H-noAGN are plotted but are so similar as to be indistinguishable. Note that this figure shows mass functions for the *full* sample of objects in H-AGN and H-noAGN, as opposed to the sample of twinned objects (see text for definition). At redshifts $z \leq 3$, the GSMFs for H-AGN lie within the region supported by observations for galaxies above stellar masses $M_* > 5 \times 10^{10} M_\odot$. H-AGN and H-noAGN overproduce the number of smaller galaxies by up to a factor of 3, partially due to a lack of magnitude cut applied to the GSMF from simulations at low stellar masses but predominantly because of the low efficiency of the stellar feedback model implemented (see the text for detail).

simulations, at all redshifts. The small spread in masses is mainly caused by variations in shape, which the structure finding algorithm translates into a small variation in virial mass. At redshifts of $z = 1$ and below, (sub)haloes with masses above $M_{\text{vir}}^{\text{H-AGN}} > 10^{11} M_\odot$ can have a DM mass up to 5 per cent lower in H-AGN. This is due to the fact that in the presence of AGN feedback, the baryon content of these massive (sub)haloes is strongly reduced (see Section 5), which translates into a reduced total (sub)halo mass, as the reduced gravitational pull slows down the cosmic inflow rate. As the (sub)haloes in H-AGN systematically exhibit lower masses, it makes sense to require, as we do, that 75 per cent of DM particles from H-AGN be present in the H-noAGN (sub)halo, and not the reverse. A more detailed analysis as to how AGN feedback affects the inner structure of DM (sub)haloes is carried out in Peirani et al. (2016).

4 AGN FEEDBACK AND STELLAR MASS

4.1 The galaxy stellar mass function

Comparing the galaxy population in H-AGN and H-noAGN at the various redshifts presented in this work shows that AGN feedback is instrumental in bringing the high-mass end of the galaxy stellar mass function (GSMF) in agreement with observations.¹ As Fig. 3 demonstrates, AGN feedback is able to suppress star formation in galaxies with masses $M_* \geq 5 \times 10^{10} M_\odot$ at $z = 3$, allowing the simulation to match the number of galaxies at and above the knee of the GSMF. It is important to note that H-AGN was *not* tuned

¹ A more detailed comparison of the GSMF in H-AGN to individual observational data sets can be found in Kaviraj et al. (2017).




Figure 4. Twinned (sub)halo DM mass ratios in H-AGN and H-noAGN at different redshifts, with quartile error regions around the median (shaded region). Note that the median twin (sub)halo with $M_{\text{vir}}^{\text{H-AGN}} < 10^{11} M_{\odot}$ has an identical DM mass with and without AGN feedback, whatever the redshift. Median intermediate mass (sub)haloes $2 \times 10^{11} M_{\odot} < M_{\text{vir}}^{\text{H-AGN}} < 2 \times 10^{13} M_{\odot}$ see their virial mass reduced by up to 5 per cent depending on mass and redshift.

to reproduce this result. As previously mentioned, the only tuning done on global galaxy properties in the simulation involves the radiative efficiency of the AGN feedback modes, which were set to reproduce the local $M_{\text{SMBH}}-M_*$ and $M_{\text{SMBH}}-\sigma_*$ relations.

In the absence of AGN feedback, the GSMF in H-noAGN agrees well with predictions that the uniform baryon mass fraction of $\Omega_b/\Omega_m = 0.165$ (dotted-dashed lines, Komatsu et al. 2011) is entirely converted into stars in the galaxy mass range of $5 \times 10^{10} M_{\odot} < M_* < 10^{12} M_{\odot}$ by $z = 1$. For galaxies with masses $M_* \geq 5 \times 10^{12} M_{\odot}$, a discrepancy between the GSMF in the absence of feedback (solid grey line) and expectation values from the CDM model (dotted line) starts to appear because gas cooling times in host (sub)haloes harbouring such massive galaxies become comparable to the Hubble time when the haloes assemble, so not all the baryons enclosed have yet been able to cool and form stars. Cooling is further hampered by the fact that in the absence of AGN feedback, heavy elements do not get distributed effectively throughout the halo but remain close to the central galaxy.

The simulations systematically overproduce the number of galaxies with masses below $M_* \leq 5 \times 10^{10} M_{\odot}$ and $z \leq 3$. This is partly caused by the fact that observed mass functions are derived from magnitude-limited data, whereas the GSMFs presented here for H-AGN and H-noAGN are raw stellar masses extracted from the simulation, with no completeness, surface brightness or luminosity cut applied. Indeed, comparing the GSMFs presented in Fig. 3 to those plotted in fig. 7 of Kaviraj et al. (2017), which are based on the *same* simulation, H-AGN, but include a magnitude cut to match observations, one realises that the effective number of galaxies with masses $M_* = 10^9 M_{\odot}$ is reduced by about 0.1–0.2 dex depending on redshift, whilst galaxies with masses $M_* \geq 5 \times 10^{10} M_{\odot}$ are completely unaffected. This does somewhat flatten the simulated GSMFs at the faint end, bringing them in better agreement with the data. However, the remaining discrepancy of about 0.3 dex with the data for galaxies with stellar masses $M_* \leq 10^{10} M_{\odot}$ can probably be attributed to the implementation of an insufficiently energetic stellar feedback model, coupled with numerical resolution effects

(Kaviraj et al. 2017),² although how efficient such a feedback can realistically be is still a matter of debate. Having said that, as AGN feedback, through the comparison of H-AGN and H-noAGN, is measured to have no effect on the low-mass end of the GSMFs (as clearly visible in Fig. 3), and the limitations discussed above are present in both simulations, they likely have a very limited impact on the work presented here. Still, the few absolute measurements presented here, such as outflow rates, have to be examined bearing in mind that stellar feedback is probably underestimated in the simulations, and thus that values derived for galaxies with stellar masses $M_* \leq 5 \times 10^{10} M_{\odot}$ are very likely too low.

Finally, both simulations systematically underproduce massive ($M_* \geq 5 \times 10^{10} M_{\odot}$) galaxies at redshift $z \geq 5$ (the top left-hand panel of Fig. 3). Given that the HMF multiplied by the universal baryon fraction (dot-dashed curve on the figure) seems to describe the observed data fairly well for galaxies in this mass range, this suggests that our inability to resolve the progenitors of haloes early enough leads to star formation being artificially postponed. Obviously, since the gas content of these haloes is still correctly estimated, galaxies will eventually catch up: their SFR will be slightly higher than expected, as long as more gas is present. However, at high redshift, galaxy star formation time-scales cannot be considered small in comparison to the time elapsed since their host halo formed, so their stellar masses can be significantly underestimated. Note that this resolution effect has completely vanished, at least for massive galaxies, by $z = 3$. Moreover, as this issue affects H-AGN and H-noAGN in the same way, it cancels out in the comparative analysis of the two simulations that we perform in this work.

4.2 Quenching

Instead of having to rely on statistical averages, such as those presented in the mass functions in Fig. 3, the twinning procedure described in Section 3 allows for a direct comparison of the stellar masses of each individual galaxy, with and without AGN feedback. Fig. 5 shows the results of this comparison for a range of redshifts. As expected from the local GSMF in Fig. 3, the most massive galaxies are the most strongly quenched at all redshifts (Fig. 5). However, the amount of quenching does not vary linearly with galaxy mass, with the function tailing off for both strongly quenched large galaxies and barely affected small galaxies.

We also measure a redshift dependence in the maximum amount of quenching observed, ranging from 40 per cent for galaxies with stellar masses in H-AGN of $M_*^{\text{H-AGN}} > 10^{10} M_{\odot}$ at redshift $z = 5$, up to a maximum of over 80 per cent for the largest galaxies at redshift $z = 0.1$, suggesting that galaxy quenching is a continuous process active throughout the merging history of galaxies. In general, the shape of the distribution is driven by small galaxies that show little influence of AGN feedback at any redshift, and a tailing off for massive galaxies. Large galaxies appear to converge to a constant quenching mass ratio $M_*^{\text{H-AGN}}/M_*^{\text{H-noAGN}} = 0.2$, as they grow from $M_* = 10^{11} M_{\odot}$ to $M_* = 10^{12} M_{\odot}$ both in H-AGN and H-noAGN. This is not due to any fundamental change in the impact of AGN feedback for galaxies with stellar masses above $M_* > 10^{11} M_{\odot}$ but rather reflects the fact that, even in the absence of AGN feedback, the GSMF in H-noAGN steepens due to long cooling times for massive objects (see Fig. 3 and Section 4.1) that lead to reduced SFRs. Therefore, the constant quenching mass ratio

² Note that a stronger stellar feedback is likely to affect black hole masses as well [Dubois et al. (2015); Habouzit, Volonteri & Dubois (2017)].




Figure 5. Quenching mass ratios, i.e. stellar mass ratios of twinned galaxies in H-AGN and H-noAGN. The twins are mass binned by stellar mass in H-AGN, M_*^{H-AGN} , with the median value (dashed lines) and quartile error ranges (shaded region) plotted for each bin. If AGN feedback had no effect, the quenching mass ratio would lie on the horizontal dotted line, which denotes twin galaxies having identical mass. Note the non-linear evolution of the quenching mass ratio with galaxy mass, with small galaxies ($M_*^{H-AGN} < 10^9 M_\odot$) barely affected, and large galaxies ($M_*^{H-AGN} > 10^{11} M_\odot$) being most strongly quenched. This plot also shows a redshift evolution, with smaller galaxies more affected at higher redshift.

for large galaxies is not driven by less effective AGN feedback, but rather by less effective cooling for galaxies in the absence of feedback. Note that this is not a selection effect either, as no mass ratio cut was applied to our galaxy sample.

Particularly noticeable for redshifts above $z > 3$, the minimum mass to experience quenching decreases with redshift: a typical $10^9 M_\odot$ galaxy at redshift $z = 5$ already has its stellar mass quenched by 10 per cent, whereas a $10^9 M_\odot$ galaxy at redshift $z = 1$ shows a median reduction in stellar mass of less than 1 per cent.

A transition mass between star forming and quenched galaxies is somewhat difficult to define in this context, as the quenching mass ratio is a cumulative measure, not an instantaneous one such as the SFR or the stellar mass growth time-scale often used to separate star forming and quiescent populations. However, based on Fig. 5, it seems natural to define the mass at which quenching due to AGN feedback becomes important to be the point where the quenching mass ratio equals 0.85 ± 0.05 , as it corresponds to the location of the sharp break in the quenching mass ratio versus galaxy stellar mass relation. As can be seen in Fig. 6, this quantity shows a redshift evolution, decreasing from $\log(M_*^{\text{quench}}/[M_\odot]) = 10.35_{-0.12}^{+0.16}$ at $z = 0.1$ to $\log(M_*^{\text{quench}}/[M_\odot]) = 9.17_{-0.11}^{+0.05}$ at $z = 5$. This result is in good agreement with the value of 10.3 found for a sample of SDSS galaxies by Baldry et al. (2004). At higher redshift, $z \approx 1$, we find good agreement with the transition mass quoted by Peng et al. (2010), who used a large sample of SDSS and zCOSMOS galaxies to study the stellar mass at which ‘mass quenching’, which includes the impact of AGN, becomes important. Note that our redshift evolution also suggests that the trend they observe can be extrapolated to $z = 2$ at least. However, mostly driven by results above $z > 4$, we find that our best-fitting power law for the redshift evolution,

$$M_*^{\text{quench}}(z) = 10^{10.49} (1+z)^{-1.50}, \quad (5)$$

has a steeper slope than that reported in Peng et al. (2010). Comparing to other simulations (Di Matteo et al. 2012; Keller et al. 2015;




Figure 6. Transition galaxy stellar mass at which AGN feedback becomes important, measured where the quenching mass ratio equals 0.85 ± 0.05 , as well as observational data from Peng et al. (2010) and Baldry et al. (2006) are shown in black, and results from a range of simulations (Di Matteo et al. 2012; Keller et al. 2015; Bower et al. 2017; Pontzen et al. 2017) are shown in grey. The values for Di Matteo et al. (2012) and Pontzen et al. (2017) are converted from the halo mass quoted in the paper, to stellar mass plotted here, using the $M_{\text{vir}}-M_*$ relation derived by Moster, Naab & White (2013).

Bower et al. 2017; Pontzen et al. 2017), we consistently find a lower value. This discrepancy is partially due to how the transition mass is defined in each work. The value in Pontzen et al. (2017), for example, is defined to be the value at which only AGN feedback quenches the galaxy. It is not surprising that this is higher than our definition of M_*^{quench} , which is the transition mass at which AGN feedback begins to be dominant, but stellar feedback might still play a role. We take the fact that we consistently measure lower transition masses to be evidence that quenching is a long-term process, whose effect accumulates over the evolution history of a galaxy. A galaxy whose SFR is reduced by 20 per cent would not sufficiently change colour to be counted amongst the ‘red and dead’ population, but could over time show a noticeably reduced quenching mass ratio. The transition mass based on the cumulative star formation history is therefore necessarily lower than that based on a more instantaneous measure, such as the stellar mass growth time-scale (Bower et al. 2017) or the SFR (Pontzen et al. 2017).

As the evolution in the quenching mass ratio (and the minimum mass of a galaxy affected) is in our case purely driven by AGN feedback, we expect the impact of AGN feedback to evolve with time. The evolution of the instantaneous AGN power plotted in Fig. 7 confirms this conjecture. Galaxies of the same stellar mass are subject to different amounts of feedback at different points in cosmic history. The median AGN power for a galaxy of a given stellar mass decreases strongly with redshift, with a galaxy with $M_*^{H-AGN} = 10^9 M_\odot$ at $z = 5$ subject to a feedback three orders of magnitude stronger than a galaxy of equivalent stellar mass at $z = 0$.

As the sound speed of the gas and its relative velocity with respect to the SMBH do not vary systematically with redshift, the feedback power of an AGN calculated according to equation (4) is mainly a function of local gas density in the vicinity of the BH, the Eddington ratio that determines the radiative efficiency of each mode of feedback and BH mass squared (the case of BHs accreting at the Eddington limit is rare and very short lived in the simulations, see e.g. Dubois et al. 2012 or Section 4.3). Therefore, the decreasing




Figure 7. Median AGN power \dot{E}_{AGN} (see equation 4) for each galaxy stellar mass bin in H-AGN, with the shaded region showing the samples quartile error ranges. At higher redshift, a galaxy of the same mass is subject to stronger feedback.




Figure 8. Median black hole mass (dashed lines) evolution as a function of galaxy stellar mass in H-AGN, as well as quartile ranges for the BH mass distribution (shaded region). The plot shows a clear redshift evolution, with smaller galaxies hosting larger black holes at higher redshift.⁴ Note that small galaxies with stellar masses below $M_{*}^{\text{H-AGN}} < 10^9 M_{\odot}$ are to be treated with some caution, as they are quite sensitive to resolution effects.

importance of AGN feedback in the evolution of small galaxies could be due to (i) the decreasing gas fractions associated with galaxy evolution, (ii) a shift in AGN feedback mode from quasar (high redshift) to radio dominated (low redshift), or (iii) an evolving black hole population (different M_{SMBH} versus M_{*} relation). We examine each of these three options, in turn, in the following section.

4.3 The coevolution of SMBHs and their hosts

The total AGN power is a function of the accretion rate on to the central SMBH, which in turn, in the Bondi regime, depends on the SMBH mass squared.³ Fig. 8 shows that the median mass

³ Or only of mass when accreting at the Eddington limit, but these episodes are rare and short lived in simulations with AGN feedback as we show later in this section; see equation (4), using equation (1) or equation (3).




Figure 9. Evolution of AGN feedback mode for the whole sample of SMBHs at different redshifts, with objects above the dotted line in quasar mode and objects below in radio mode. The solid lines denote the median Eddington ratio. There is a clear redshift evolution of the feedback mode, with 95.7 per cent of SMBHs in quasar mode at $z = 5$, in comparison to only 2.3 per cent at $z = 0.1$. The colour bar in each panel is normalized to the bin with the maximum number of objects at the given redshift.

of the central black hole undergoes a redshift evolution between $z = 0$ and $z = 5$, with galaxies of a given stellar mass hosting a more massive BH at higher redshift.⁴ For example, a galaxy with $M_{*}^{\text{H-AGN}} = 10^{10} M_{\odot}$ at redshift $z = 0.1$ typically hosts an SMBH with $M_{\text{SMBH}} = 1.1 \times 10^7 M_{\odot}$, whereas a galaxy with the same stellar mass at $z = 5$ hosts an SMBH with a median mass of $M_{\text{SMBH}} = 4.2 \times 10^7 M_{\odot}$. A similar evolutionary trend is reported observationally (Decarli et al. 2010; Merloni et al. 2010) and in other large-scale cosmological simulations (Khandai et al. 2015; Sijacki et al. 2015). A detailed discussion of the shape of the $M_{\text{SMBH}}-M_{*}$ relation can be found in Volonteri et al. (2016).

Looking at the evolution of AGN feedback mode with redshift, as determined by the Eddington ratio χ , Fig. 9 reveals that the SMBH population transitions from quasar mode to radio mode between redshifts $z = 3$ and $z = 1$, as χ falls below 0.01. At high redshift, the vast majority of AGN are found in quasar mode, namely 95.7 per cent of the sample at $z = 5$ and 85.5 per cent at $z = 3$. At lower redshift, the population is overwhelmingly in radio mode across all mass bins, with only 19.1 per cent and 2.3 per cent found in quasar mode at $z = 1$ and $z = 0.1$, respectively.

As the Eddington ratio is a measure of how efficiently a black hole of a given mass is accreting, the high Eddington ratios at redshift $z > 3$ explain the evolution in the $M_{\text{SMBH}}-M_{*}$ relation in Fig. 8. Indeed, whilst a black hole with $M_{\text{SMBH}} = 10^7 M_{\odot}$ at redshift $z = 5$ accretes with a mean Eddington ratio of $\chi = 6.27 \times 10^{-2}$, an SMBH with the same mass at redshift $z = 0.1$ accretes at only

⁴ The results presented here are consistent with fig. 10 of Volonteri et al. (2016), which reports no redshift evolution in the $M_{*}-M_{\text{SMBH}}$ relation. This is because the redshift evolution we see is mainly driven by low-mass black holes at low redshift, a sample excluded by these authors who apply a cut in host halo mass of $M_{\text{halo}} = 8 \times 10^{10} M_{\odot}$. By comparison, the sample analysed here includes all black holes identified within a host galaxy in haloes with $M_{\text{halo}} > 4 \times 10^{10} M_{\odot}$. The second difference between these two pieces of work concerns the statistical analysis chosen: whilst Volonteri et al. (2016) employ a linear fit as used in observational studies, we present median black hole masses and thus allow for a non-linear correlation between black hole and galaxy stellar masses.




Figure 10. Specific accretion rate on to the central SMBH in Eddington units versus average gas density of the host galaxy for various redshifts. The measured decrease in average specific accretion rate for SMBHs of any mass, as redshift decreases, is driven by a decreasing gas density in the host galaxy.

$\chi = 1.2 \times 10^{-4}$ Eddington. This means that the latter grows about two orders of magnitude more slowly than the former.

Fig. 9 also clearly shows that except for a few outlying objects, the bulk of the population is not accreting in an Eddington limited fashion at $z \leq 5$. This evolution in the median Eddington ratio reflects an underlying evolution in the gas density of galaxies, as can be seen in Fig. 10 where we have divided χ by the SMBH mass to calculate the specific accretion rate of these SMBHs in Eddington units, i.e. we have removed the dependence on BH mass to be able to intercompare BHs across the whole mass range. The story that emerges from this plot is that since high redshift galaxies are more gas rich, they fuel their central black holes more efficiently, regardless of their masses. As the gas supply is depleted, accretion on to the central black hole slows down, and AGN feedback transitions from a quasar to a radiative feedback mode around $z = 2$. Due to the different radiative efficiencies employed for the two feedback modes, $\epsilon_f = 0.15$ for the quasar mode at $\chi \geq 0.01$ and $\epsilon_f = 1.0$ for the radio mode at $\chi < 0.01$ (see Section 2.5), a larger percentage of accretion energy is converted into feedback at redshifts below $z < 2$, but nevertheless, the amount of energy available for feedback declines.

Summarizing the impact of all three effects, we conclude that the decreasing AGN power for a galaxy of a given stellar mass is driven by the cumulative effect of a proportionally larger central SMBH and the decreasing gas supply in the galaxy, for which the increasing efficiency of the feedback mode at $z < 2$ is unable to compensate. Not only do existing BHs of a given mass accrete less efficiently in the gas poor galaxies at $z < 2$, such that $\dot{M}_{\text{BH}}(M_{\text{SMBH}}, z > 2) > \dot{M}_{\text{BH}}(M_{\text{SMBH}}, z < 2)$ for all M_{SMBH} , but any galaxy of a given stellar mass also hosts a significantly smaller black hole than its equivalent counterpart at $z > 2$, i.e. $M_{\text{SMBH}}/M_*(z < 2) < M_{\text{SMBH}}/M_*(z > 2)$. The two effects combine together to produce the redshift evolution of AGN power seen in Fig. 7, despite the shift to a more efficient feedback mode around $z = 2$.

How efficiently an AGN of a given power is able to remove gas from a galaxy is dependent on the depth of the gravitational potential it has to overcome in the process. Repeating the analysis of the quenching mass ratios, but plotting it against $M_{\text{SMBH}}/M_{\text{vir}}^{\text{H-AGN}}$ instead of the galaxy stellar mass, shows that the redshift evolution




Figure 11. Stellar mass ratio between H-AGN and H-noAGN galaxy twins against ratio of central SMBH mass to virial mass of the host (sub)halo, instead of stellar mass as in Fig. 5. The redshift evolution is erased in that case because the feedback energy deposited is directly connected to the central BH mass (see the text for detail).

in the quenching mass ratio is truly driven by the evolution of the SMBH population and its feedback power (Fig. 11 in comparison to Fig. 5). Indeed, whilst the relation scatter augments, the different redshift curves now overlap within the quartile error bars: the redshift dependence in Fig. 5 has been erased. In other words, independently of redshift, BHs with a black hole to halo mass ratio of less than $M_{\text{SMBH}}/M_{\text{vir}}^{\text{H-AGN}} \leq 4 \times 10^{-5}$ quench their host galaxy by less than 20 per cent, whereas BHs with ratios only a factor of 2–3 larger than that suppress their host galaxy stellar masses by up to 50 per cent.

Although there exists a clear transition of the sample from one AGN feedback mode to the other, no significant difference was found when analysing the quenching mass ratio (such as in Fig. 11) by splitting the sample into quasar or radio mode galaxies. All results shown here can be reproduced by assuming that the entire sample can be found in quasar mode at redshifts $z > 2$, and in radio mode otherwise. The only notable discrepancy between the simulation data and this simplified model is that the scatter is somewhat reduced, which is expected as objects found in the opposite feedback mode to the majority of the population are statistical outliers. However, this analysis is based on an instantaneous measure of feedback mode at a specific redshift and does not capture the accretion history of a particular object. We defer a more careful analysis of the evolution of the AGN sample, together with an analysis of the time-scales on which quenching occurs in individual galaxies, to future work.

5 AGN FEEDBACK AND GAS FLOWS

Fig. 8 shows that SMBHs make up much less than 1 per cent of the mass of their host galaxy, so the mass of baryons accreted by the SMBHs is negligible compared to the reduction in stellar mass caused by quenching. The effect of AGN feedback on the cold gas supply of the galaxy must therefore be profound to suppress star formation by up to an order of magnitude over the evolution of the galaxy. There are three possible channels through which AGN feedback can affect the gas content of the galaxy: (i) it can drive powerful outflows, emptying the reservoir of gas available in the ISM of the galaxy; (ii) it can prevent cosmic inflows from




Figure 12. Ratio between the average baryon density $\bar{\rho}_{\text{bar}}^{\text{H-AGN}} / \bar{\rho}_{\text{bar}}^{\text{H-noAGN}}$ within the virial radius of the galaxy host halo twins: $\bar{\rho}_{\text{bar}} = (M_* + M_{\text{SMBH}} + M_{\text{gas}}) / (\frac{4\pi}{3} R_{\text{vir}}^3)$. The dotted line represents identical baryon mass in both host haloes, and the shaded regions show the quartile ranges of the sample. AGN feedback partially acts on star formation by reducing the total baryon mass in the halo, particularly for galaxies with stellar masses around $M_*^{\text{H-AGN}} = 10^{11} M_\odot$ in H-AGN.

replenishing the gas supply in the galaxy; or (iii) it can heat existing gas of the ISM and circum-galactic medium to prevent cooling flows and the associated star formation. In this section, we investigate the relative importance of these three feedback channels.

5.1 The evolution of the baryon content

Should AGN feedback primarily suppress star formation through heating the existing gas in the halo, one would expect twinned haloes to have the same total baryon mass, with that in H-AGN showing a much higher gas fraction as less gas is being turned into stars. Fig. 12 shows that AGN feedback directly lowers the baryon (gas + stars + BH) content embedded within the virial radius of DM haloes: the average baryon density ratio versus galaxy stellar mass relation follows a shape reminiscent of that of the quenching mass ratio previously discussed. Unfortunately all three major channels, through which AGN feedback is expected to affect star formation, lower the baryon density of the galaxy. Boosted outflows drive existing gas out of the galaxy, slowed down inflows prevent accretion in the first place and heating causes the gas to expand, lowering the average density. We compare the average baryon density, as opposed to the total mass within the halo, to correct for the small differences in halo mass shown in Fig. 4, which translate into a difference in virial radius, and therefore a difference in the volume over which the gas mass in the halo is measured.

Two features stand out in comparison to the quenching mass ratio. First, the average baryon density within the halo of small galaxies is unaffected by AGN feedback at all redshifts. These galaxies do however show a reduced stellar mass, particularly at redshift $z = 5$, where their stellar mass shows a median reduction of 10 per cent, as seen in Fig. 5. This suggests that feedback affects the star formation efficiency of these galaxies more than it alters their gas supply. Efficiency is reduced either by locally heating the gas, redistributing the gas within the halo, or by destroying dense, star-forming clumps, but not by driving outflows or preventing gas inflows through the halo virial sphere. Secondly, for large galaxies with stellar masses above

$M_* > 10^{11} M_\odot$, the baryon content in both simulations becomes increasingly comparable again with increasing mass, despite the fact that these galaxies see a reduction in their stellar mass of around 80 per cent for redshifts between $z = 1$ and $z = 0.1$. In this case, the deepening gravitational potential of the halo makes it difficult for the AGN to affect gas flows at the halo virial radius. However, a much smaller fraction of the existing gas is converted into stars, because of effects (gas heating/redistribution) similar to those which plague small galaxies. These translate into the significantly reduced galaxy masses presented earlier in Fig. 5. For medium galaxies at all redshifts, AGN feedback acts through depleting the gas reservoir at the halo scale, which reduces the supply of gas available for star formation. This means AGN feedback also directly influence the inflows and/or the outflows of the galaxy. We leave a detailed analysis of the interstellar and intergalactic medium under AGN feedback to future work.

5.2 The effect on inflows and outflows

There are two ways to decrease the total baryon mass of a galaxy: by reducing inflows or by boosting outflows. In this work, we measure flows at two different radii: halo scales, also called R95, which correspond to a radius of $R95 = 0.95 \times R_{\text{vir}}$, and galaxy scales, also called R20, which correspond to a radius of $R20 = 0.2 \times R_{\text{vir}}$. Flows are measured through spherical surfaces located at these radii, centred on the halo. Flow masses are calculated for all cells within a narrow shell, centred on the radius in question, where $\dot{M}_{\text{gas}} = \sum_i \rho \Delta x_i^3 \bar{v}_i \cdot \bar{r}_i / \omega$, where ρ is the gas density, Δx is the cell size, \bar{v}_i is the gas velocity, \bar{r}_i is the unit vector of the cell centre relative to the halo centre and $\omega = 2$ kpc is the width of the shell. \dot{M}_{outflow} includes all cells with $\bar{v}_i \cdot \bar{r}_i > 0$ and \dot{M}_{inflow} includes all cells for which $\bar{v}_i \cdot \bar{r}_i \leq 0$.

A first comparative look at the flow patterns for galaxies in H-AGN and H-noAGN (see Fig. 13) suggests that AGN feedback drives outflows in medium and large galaxies, particularly at halo scales. There are also some large-scale pseudo-flows present for large galaxies. These pseudo-flows appear in Fig. 13 because the algorithm used to extract the absolute flow values presented here assumes the halo can be accurately represented by a sphere. However, if the halo is non-spherical, pseudo-flows are created. When rotating a non-spherical object through a spherical surface across which absolute mass flows are measured, parts of the object passing out of the sphere will register as outflows, whilst parts passing in will register as inflows. However, these contributions are not mass flows in the common sense, and cancel out when calculating net mass flows.

Small galaxies undergo no outflows at halo scales, with or without AGN feedback, which matches the conclusion from Fig. 12 that the baryon mass of their haloes is identical in H-AGN and H-noAGN. AGN feedback reduces inflows for medium and large galaxies, at both halo and galaxy scales, but the effect is more pronounced at the latter. A more quantitative analysis of the outflows driven by AGN feedback is presented in Fig. 14, where residual flow values for the two simulations are plotted. Residual flows are defined as the mass flow rates in H-AGN relative to those of their twin galaxies in H-noAGN, i.e. $\dot{M}_{\text{gas}}^{\text{residual}} = \dot{M}_{\text{gas}}^{\text{H-AGN}} - \dot{M}_{\text{gas}}^{\text{H-noAGN}}$. This approach has the advantage of isolating the effect of AGN feedback and subtracting out any effects present in both simulations, such as the SN driven outflows for small and medium galaxies at halo scales, and the pseudo-flows seen for large objects at both radii.

The residual gas flows shown in Fig. 14 demonstrate that AGN feedback has an approximately equal and opposite effect




Figure 13. Gas inflows (dashed lines) and outflows (solid lines) at halo (R95, measured at $R = 0.95R_{\text{vir}}$) and galaxy (R20, measured at $R = 0.2R_{\text{vir}}$) scales for both H-AGN and H-noAGN for a range of redshifts. To simplify the comparison with other plots, flow values for both H-AGN and H-noAGN are plotted against H-AGN galaxy stellar masses. Note that both outflows and inflows represent the median value for a given mass bin and do not necessarily belong to the same object in H-AGN and H-noAGN. The overall inflows are reduced in the presence of AGN feedback, which also drives outflows at halo scales for medium-size galaxies.

on outflows and inflows. Galaxies with stellar masses $M_*^{H-AGN} \leq 10^{11} M_\odot$ see a similar amount of gas carried away by AGN driven outflows as that depleted in inflows. Apart for small galaxies at high redshift ($z = 5$), where AGN feedback seems able to heat up the gas in the vicinity of the galaxy, causing a gas pile up that triggers larger inflows in H-AGN, the differences between flows at halo and galaxy scales are rather modest. There is a weak trend for boosted outflows to be less dominant at galaxy scales (especially at $z \leq 1$) than at halo scales, and conversely for the suppression of inflows to be more relevant on small scales (especially for galaxies that are less massive than $M_*^{H-AGN} \leq 10^{10} M_\odot$), but overall material is neither being significantly swept up in the halo and kicked out nor preferentially deposited there. Rather, these two effects combine to reduce the net median inflow for medium galaxies into both the halo and the galaxy by up to 60 per cent for the most massive objects, as the direct comparison of net inflows for twins in Fig. 15 shows. This results in the reduction in baryon mass seen for the same galaxies in Fig. 12. In agreement with the same figure, small galaxies show no change in net flows at halo scales, which match their identi-

cal baryon mass in the presence and absence of AGN feedback. These results are in agreement with work by Dubois et al. (2013), who compare high resolution zoom simulations with and without AGN feedback of a single galaxy of $M_*(z = 6) = 6.2 \times 10^9 M_\odot$. The authors find strong evidence for the fact that the AGN significantly heat the gas at halo scales, driving a hot superwind and destroying cold flows. Pontzen et al. (2017) also report AGN boosted outflows in their hydrodynamical merger simulation and emphasize that long-term quenching requires the inflows to be suppressed.

For large galaxies, the situation is different, particularly at low redshift. At halo scales, the simulation with AGN feedback actually shows boosted inflows carrying an amount of mass similar to that in boosted outflows. This means the outflows are being recycled, as AGN feedback becomes unable to gravitationally unbind the gas from the halo. It is important to note that the values plotted here represent the median value for a given mass bin, so the outflow and inflow values do not necessarily belong to the same object. It is therefore not necessarily correct that the two curves cancel out to

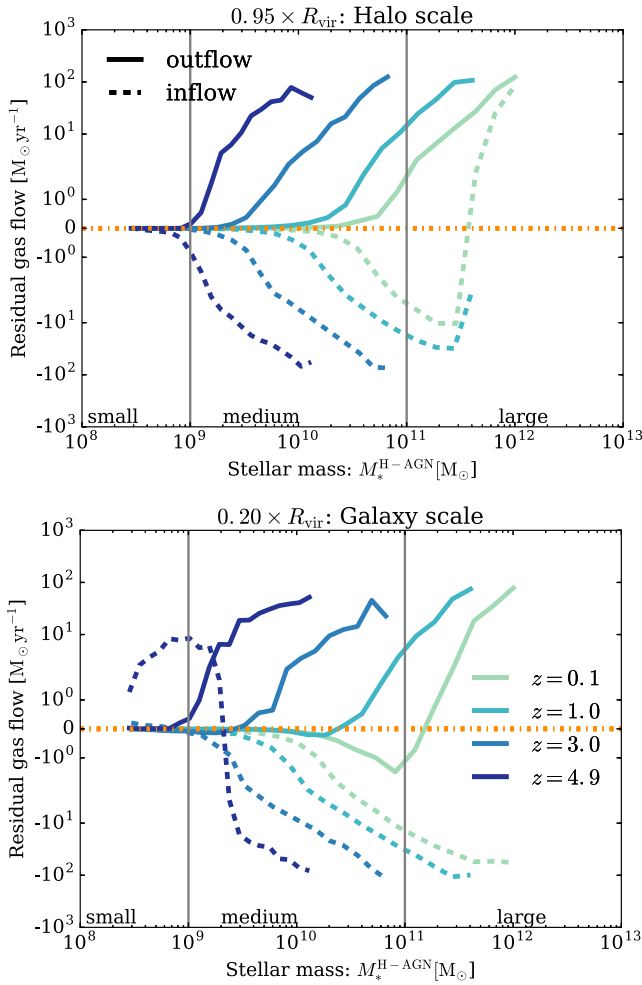


Figure 14. Residual gas flows $\dot{M}_{\text{gas}}^{\text{residual}} = \dot{M}_{\text{gas}}^{\text{H-AGN}} - \dot{M}_{\text{gas}}^{\text{H-noAGN}}$, plotted against the twin's galaxy stellar mass in H-AGN. Data are presented for flows at two different distances from galaxies, as indicated above each panel, and a range of redshifts. Higher mass flow rates in H-AGN appear above the dash-dotted line, which denotes identical flows in both simulations, and lower mass flow rates sit below this line. Note that both outflows and inflows represent the median value for that mass bin and do not necessarily belong to the same twin. In general, AGN feedback causes a boost in outflows and an approximately equal but opposite reduction in inflows.

produce no change in the net flow. Indeed, a comparative analysis of net inflows for each twin across H-AGN and H-noAGN (Fig. 15) shows that at halo scales, the overall inflow is boosted by up to 50 per cent for the most massive galaxies in the presence of AGN feedback.

At galaxy scales, the gas flow patterns for the most massive objects, $M_{*}^{\text{H-AGN}} > \times 10^{11} \text{ M}_{\odot}$ at $z = 0.1$, become harder to predict. In these galaxies, AGN fall into maintenance feedback mode (see Section 4.3) at redshifts below the peak of star formation, $z = 2$. This produces very bursty outflows, as SMBHs go through cycles of being fed, which triggers strong feedback episodes. The latter drive out the gas, starving the black hole, and the feedback abates until enough gas becomes available again for the SMBH to go through another accretion event. These cycles take place on time-scales much shorter than the time interval between the redshift outputs we are considering in this work. Furthermore, the number of such galaxies is quite limited (≈ 1000). We illustrate the im-

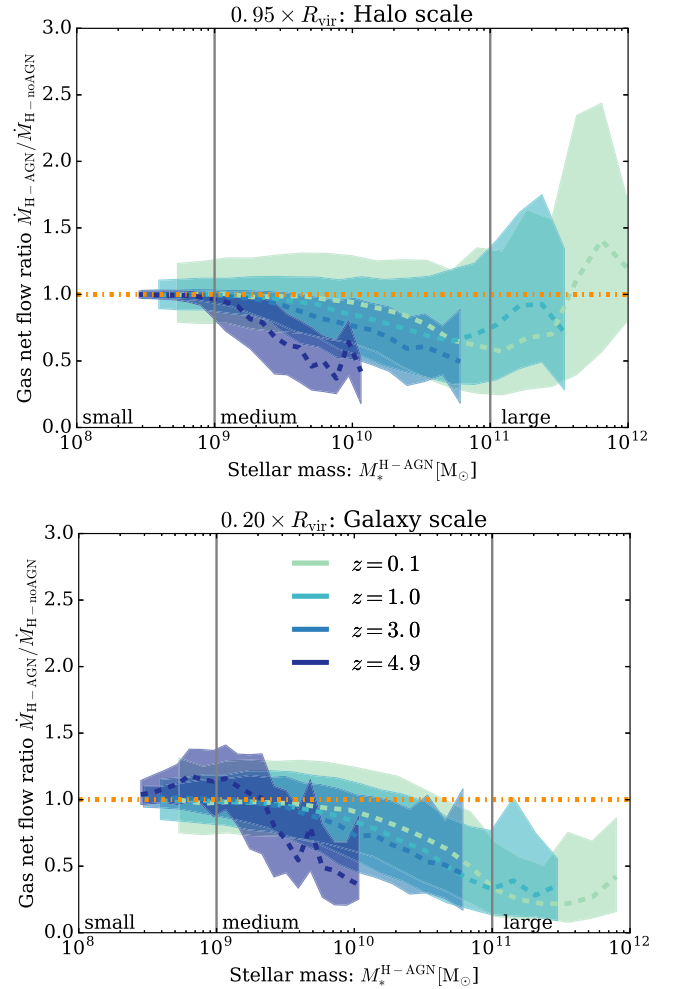


Figure 15. Ratio of net flow rates at halo and galaxy scales as a function of galaxy stellar masses in H-AGN, $M_{*}^{\text{H-AGN}}$. The dash-dotted horizontal line at a ratio of 1.0 corresponds to identical net flows in both H-AGN and H-noAGN. Data are presented for flows at two different distances from galaxies, as indicated above each panel, and at a range of redshifts. AGN feedback has no effect on flows in small galaxies, produces a reduction in median net flow rates of up to 60 per cent for medium-size galaxies and drives a bursty flow pattern leading to increased net flows at halo scales for large galaxies. Positive net flows are defined to be falling into the object.

pact this has on our results in Fig. 16, which shows that at redshift $z = 0.5$, the median outflow depends quite sensitively on the exact point in time at which the distribution of galaxies is sampled. As the SMBHs cycle rapidly through a wide variety of active and quiescent states, the distribution of residual outflows spans several orders of magnitude. In comparison, the outflows produced by the larger population of smaller galaxies are more steady on similar time-scales, and therefore their sampling is more robust (see the left-hand panel of Fig. 16). We would like to point out that the variation time-scales of the large-scale outflows studied here do not necessarily reflect the duty cycle of the SMBH, as each burst can be driven by a series of feedback events. We postpone a more detailed analysis of the SMBH duty cycles in the simulation to future work.

It is interesting to note that the lack of excess outflows for low-mass galaxies ($M_{*}^{\text{H-AGN}} < 5 \times 10^{10} \text{ M}_{\odot}$ at $z = 0.1$) in Fig. 14 does not mean that the outflows for any given matching galaxy pair in




Figure 16. Residual outflows due to AGN feedback at galaxy scales, $R_{20} = 0.2 \times R_{\text{vir}}$, plotted for 10 randomly selected galaxies within the small and large galaxy mass bins, between redshifts $0.64 > z > 0.42$. The solid line represents the median sample output, with the shaded region covering quartile error ranges. The vertical dotted line represents redshift $z = 0.52$, for comparison with Fig. 14. This shows that outflows from large galaxies (right-hand panel) with stellar masses $M_*^{\text{H-AGN}} > 2 \times 10^{11} M_\odot$ at redshifts around $z = 0.52$ (dotted vertical line) vary on very short time-scales, so the median outflow value (solid line, with quartile error ranges represented by the shaded regions) is very sensitive to the exact time at which it is sampled. By comparison, the outflows for small galaxies with stellar masses $M_*^{\text{H-AGN}} < 5 \times 10^{10} M_\odot$ (left-hand panel) vary much less on these time-scales, and the median value is less sensitive to variations in time.

H-AGN and H-noAGN are identical. As the panel on small galaxies in Fig. 16 demonstrates, individual objects show a variety of residual outflows (and inflows). The build-up of small differences in galaxy properties not necessarily induced by AGN feedback (e.g. stochastic star formation algorithm, seeding and growth of the central SMBH) can result in temporarily diverging residual outflow histories at low redshift. In other words, even though the precise amount of residual outflow from any specific twinned pair of galaxies does depend on the time-scale of the outflow and so is sensitive to the redshift at which it is measured, the lack of any marked systematic difference in the gas flow pattern due to AGN feedback registers as a median residual outflow of zero for the whole sample.

Another note of caution concerns the residual flows in Fig. 14, which likely underestimate the effect of AGN on inflows and outflows. Particularly for massive galaxies at low redshift, where the H-noAGN twin has a stellar mass $\sim 5 \times$ that of the H-AGN, the H-noAGN twin has stronger stellar flows that obscure some of the AGN-driven effects when calculating residual gas flows as $\dot{M}_{\text{gas}}^{\text{residual}} = \dot{M}_{\text{gas}}^{\text{H-AGN}} - \dot{M}_{\text{gas}}^{\text{H-noAGN}}$. Finally, we come back to the asymmetry between residual inflows and outflows with respect to the zero residual line, which is stronger at galaxy than halo scales. In light of the previous remarks about outflow time-scales, we can safely attribute this difference as meaning that AGN feedback does preferentially suppress inflows in the vicinity of all galaxies rather than eject gas from them, except perhaps at the very high mass end of the galaxy stellar mass function, $M_*^{\text{H-AGN}} > \times 10^{11} M_\odot$ at low redshifts ($z < 0.1$).

Overall, flow patterns due to AGN feedback, made up in roughly equal parts of boosted outflows and reduced inflows at halo scales, comfortably explain the non-linear distribution of baryon masses plotted in Fig. 12, which combines with a reduced star formation efficiency across all galaxy mass bins to produce the quenching mass ratios shown in Fig. 5.

6 DISCUSSION AND CONCLUSIONS

We have isolated the effect of AGN feedback on stellar quenching in massive galaxies by comparing two cosmological simulations, H-AGN and H-noAGN, which were run with and without AGN feedback, respectively. More specifically, by twinning individual DM haloes and galaxies across the two simulations, we have been able to quantify the effect of feedback on individual objects throughout cosmic time. In agreement with a large body of previous work (such as Springel et al. 2005; Bower et al. 2006; Cattaneo & Teyssier 2007; Ciotti & Ostriker 2007; Sijacki et al. 2007; Dubois et al. 2013; Gene et al. 2014; Schaye et al. 2014; Pontzen et al. 2017, amongst others), our results show that AGN feedback is instrumental in quenching the massive end of the GSMF. Whilst the stellar mass of galaxies without AGN feedback closely follows predictions based on the assumption that all baryons contained in DM haloes end up forming stars, galaxies subject to the influence of AGN feedback end up with masses distributed according to a GSMF that shows a characteristic exponential steepening at the high-mass end, in line with observations (Daddi et al. 2002; González et al. 2011; Moustakas et al. 2013; Bielby et al. 2014; Tomczak et al. 2014; Song et al. 2016).

The importance of AGN feedback has been emphasized in all recent large-scale simulations of galaxy evolution but the results differ in the details. Similar to results presented here, galaxies in the MassiveBlackII simulation exhibit signs of relatively strong quenching early on but then see a reduction in the impact of AGN feedback at lower redshifts (Khandai et al. 2015). A similar issue is reported by Illustris, who found that despite aggressive AGN feedback that produces unrealistically low gas fractions in DM haloes at low redshift, star formation is not suppressed strongly enough and the simulation overproduces massive galaxies (Gene et al. 2014; Vogelsberger et al. 2014a,b). As opposed to the dual AGN feedback model used in H-AGN and Illustris, the EAGLE simulation only

employs a single feedback mode, tuned to reproduce the GSMF at $z = 0$. They do also see good agreement with the GSMF at redshift $z > 2$, which supports our conclusion that the shift in feedback mode plays a subordinate role in the importance of AGN feedback (Schaye et al. 2014). Overall, regardless of the hydrodynamic scheme employed and the detail of the subgrid model implementation of AGN feedback, all four recent large-scale cosmological simulations of galaxy evolution agree that AGN feedback must play a crucial role in regulating the evolution of massive galaxies.

A closer comparison of the stellar mass for individual objects in H-AGN and H-noAGN reveals a non-linear dependence of quenching mass ratio on mass, with the most massive galaxies being the most strongly quenched, and the smallest galaxies mostly unaffected by AGN feedback. This leads to a characteristic shape for the mass ratio $M_*^{\text{H-AGN}}/M_*^{\text{H-noAGN}}$, which shows a linear dependence on $\log(M_*^{\text{H-AGN}})$ for medium-sized galaxies with $10^9 M_\odot \lesssim M_*^{\text{H-AGN}} \lesssim 10^{11} M_\odot$ (the exact values depend on redshift), but tails off at both low- and high-mass ends. The most massive galaxies, with $M_* > 10^{11} M_\odot$, are most strongly quenched and contain only 20 per cent of the stellar mass in the presence of AGN feedback at $z = 0.1$, in comparison to the case without feedback.

We also find a significant redshift evolution for the smallest galaxy mass to be affected by AGN feedback, with smaller galaxies being more quenched at higher redshift. This transition mass at which AGN feedback becomes important evolves from $2.24 \times 10^{10} M_\odot$ at $z = 0.1$ to $1.48 \times 10^9 M_\odot$ at $z = 4.9$. Such a redshift dependence seems in good agreement with observations (Baldry et al. 2004; Peng et al. 2010) but systematically leads to values at high redshift that are lower than those reported by a range of hydrodynamic simulations (Di Matteo et al. 2012; Keller et al. 2015; Bower et al. 2017; Pontzen et al. 2017). We argue that these discrepancies reflect the fact that quenching is very likely a cumulative process that builds up over the entire history of the galaxy, not just a one-off event that shuts down star formation forever when the galaxy reaches a particular stellar mass. If correct, the consequence is that differences induced by AGN feedback in SFRs can be small, particularly for objects near the transition mass, the effect of such feedback only becoming apparent for galaxies once its integral represents a significant fraction of their total stellar mass.

In our case, this evolution is caused by the median central black hole mass of galaxies of a given stellar mass, M_{SMBH}/M_* , increasing with redshift by a factor of a few between $z = 0$ and $z = 5$, a result for which there exists observational support (Decarli et al. 2010; Merloni et al. 2010). However, it has been suggested that the observation trend might be biased (Lauer et al. 2007; Volonteri & Stark 2011; Bongiorno et al. 2014). A similar trend has been reported in other large-scale simulations, such as MassiveBlackII (Khandai et al. 2015) and Illustris (Sijacki et al. 2015), but not in all of them (Rosas-Guevara et al. 2016). This evolution in black hole masses combines with higher accretion rates at high redshift, due to gas-rich galaxies, such that a galaxy of the same stellar mass is subject to AGN feedback up to three orders of magnitude stronger at redshift $z = 5$ than at $z = 0.1$. We also measure a shift in feedback mode, with at least 85.5 per cent of AGN at redshift $z = 3$ and above in quasar mode, compared with a maximum of 19.1 per cent at lower redshifts, but the increasing radiative efficiency associated with this shift is unable to offset the trend of lower M_{SMBH}/M_* and lower accretion rates.

A comparative analysis of the baryon content of haloes reveals that AGN feedback quenches star formation through a combination

of reducing the total gas supply within the halo by driving outflows, and preventing accretion of fresh gas by curbing inflows. Small galaxies with $M_*^{\text{H-AGN}} \leq 10^9 M_\odot$ show nearly identical baryon masses with and without AGN feedback, as the flows at halo-scales remain chiefly unaffected by feedback. Note that this is not true at galaxy scales and high redshift, where gas inflows can be somewhat enhanced by feedback for these small objects. On the other hand, medium-size galaxies $10^9 M_\odot \leq M_*^{\text{H-AGN}} \leq 10^{11} M_\odot$ experience a significant reduction in baryon mass, caused by an approximately equal contribution from AGN-driven gas outflows and a reduction of cosmic inflows. At galaxy scales, the reduction of inflows dominates. Finally, for large galaxies ($10^{11} M_\odot \leq M_*^{\text{H-AGN}}$) the baryon mass rises again, as inflows at halo scales are swelled by gas expelled by AGN feedback in the inner regions, which remains gravitationally bound and falls back into the halo. At galaxy scales, outflows for the most massive objects vary on very short time-scales, as the AGN enter a bursty maintenance mode. Thus, the gas mass has a tendency to increase on average (as compared to medium-size galaxies), even though the stellar mass does not, given the long characteristic time-scale of star formation.

The picture that emerges seems consistent with high-resolution work (20 pc instead of 1 kpc) by Dubois et al. (2013), who use zoom simulations of individual objects with and without AGN feedback and find that AGN drive hot super-winds that disrupt cold inflows on halo scales. The two effects combine to reduce the baryon content of galaxies by up to 30 per cent. Similar conclusions were reached by early semi-analytic work by, e.g. Benson et al. (2003), who argue that strong, hot winds are necessary to reproduce the observed luminosity function, and recent simulations by Pontzen et al. (2017) and Taylor & Kobayashi (2015) also report AGN-boosted outflows, concluding that long-term quenching requires gas inflows to be suppressed.

Looking at cosmic accretion specifically, Nelson et al. (2015) find that the simulation without feedback (neither AGN driven nor stellar) sees much higher levels of smooth accretion into the galaxy than the ones with (AGN and stellar) feedback. However, contrary to the results presented here, they find no evidence for recycling of gaseous material at the halo boundary. Whilst all current large-scale cosmological simulations include AGN feedback as in integral part of their galaxy evolution model, some authors contend that processes in massive galaxies that do not rely on AGN feedback can reproduce galaxy mass functions through e.g. cosmic quenching (Feldmann & Mayer 2014) and that stellar feedback super-bubble feedback can drive powerful outflows (Keller et al. 2015). Gabor & Bounaud (2014) and Roos et al. (2015), running idealized galaxy simulations, find AGN-driven outflows consistent with those we report here but lower impact on the SFR of galaxies.

Overall, we conclude that AGN feedback provides an effective mechanism to reproduce the distribution of galaxies at and above the knee of the GSMF, over a redshift range spanning 90 per cent of the age of the Universe. For local galaxies, AGN feedback plays an important role in stifling star formation in objects above a transition mass of $M_* \geq 2 \times 10^{10} M_\odot$. AGN feedback acts by reducing the stellar content of galaxies by up to 80 per cent (for the most massive objects) through a mixture of increased outflows and reduced inflows, combined with a decreased star formation efficiency of *in situ* gas. We predict that the influence of AGN feedback should already be noticeable by redshift $z = 5$, for galaxies with relatively modest stellar masses ($M_* \approx 2 \times 10^9 M_\odot$) by current epoch standards, as these objects are close to the top end of the GSMF at these redshifts. This is exciting news as the *James Webb Space Telescope* should be able to test this prediction in the near future.

ACKNOWLEDGEMENTS

We would like to thank the reviewer, Richard Bower, for his constructive comments that improved the quality of this paper. This work used the HPC resources of CINES (Jade supercomputer) under the allocation 2013047012 made by GENCI, and the horizon and Dirac clusters for post processing. This work is partially supported by the Spin(e) grants ANR-13-BS05-0002 of the French Agence Nationale de la Recherche and by the National Science Foundation under Grant No. NSF PHY11-25915, and it is part of the Horizon-UK project, which used the DiRAC Complexity system, operated by the University of Leicester IT Services, which forms part of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment is funded by BIS National E-Infrastructure capital grant ST/K000373/1 and STFC DiRAC Operations grant ST/K0003259/1. DiRAC is part of the National E-Infrastructure. The research of RSB is supported by STFC, and the research of AS, MLAR and JD at Oxford is supported by the Oxford Martin School and Adrian Beecroft. NC is supported by a Beecroft postdoctoral fellowship.

REFERENCES

- Arav N. et al., 2001, *ApJ*, 561, 118
- Aubert D., Pichon C., Colombi S., 2004, *MNRAS*, 352, 376
- Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Z., Lupton R. H., Nichol R. C., Szalay A. S., 2004, *ApJ*, 600, 681
- Baldry I. K., Balogh M. L., Bower R. G., Glazebrook K., Nichol R. C., Bamford S. P., Budavari T., 2006, *MNRAS*, 373, 469
- Barro G. et al., 2016, *ApJ*, 820, 120
- Benson A. J., Bower R. G., Frenk C. S., Lacey C. G., Baugh C. M., Cole S., 2003, *ApJ*, 599, 38
- Bernardi M., Meert A., Sheth R. K., Vikram V., Huertas-Company M., Mei S., Shankar F., 2013, *MNRAS*, 436, 697
- Bielby R. M. et al., 2014, *A&A*, 568, A24
- Bieri R., Dubois Y., Silk J., Mamon G. A., 2015, *ApJ*, 812, L36
- Binney J., Tabor G., 1995, *MNRAS*, 276, 663
- Bongiorno A. et al., 2014, *MNRAS*, 443, 2077
- 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
- Cattaneo A., Teyssier R., 2007, *MNRAS*, 376, 1547
- Cattaneo A. et al., 2009, *Nature*, 460, 213
- Chartas G., Brandt W. N., Gallagher S. C., Proga D., 2007, *AJ*, 133, 1849
- Chisari N. E. et al., 2017, preprint ([arXiv:1702.03913](https://arxiv.org/abs/1702.03913))
- Christensen C. R., Davé R., Governato F., Pontzen A., Brooks A., Munshi F., Quinn T., Wadsley J., 2016, *ApJ*, 824, 57
- Cicone C. et al., 2014, *A&A*, 562, A21
- Ciotti L., Ostriker J. P., 2007, *ApJ*, 665, 1038
- Cole S., Aragon-Salamanca A., Frenk C. S., Navarro J. F., Zepf S. E., 1994, *MNRAS*, 271, 781
- Costa T., Sijacki D., Haehnelt M. G., 2015, *MNRAS*, 448, L30
- Croton D. J. et al., 2006, *MNRAS*, 365, 11
- Daddi E. et al., 2002, *A&A*, 384, L1
- Davidzon I. et al., 2017, *A&A*, preprint ([arXiv:1701.02734](https://arxiv.org/abs/1701.02734))
- De Kool M., Arav N., Becker R. H., Gregg M. D., White R. L., Laurent-Muehleisen S. A., Price T., Korista K. T., 2001, *ApJ*, 548, 609
- Decarli R., Falomo R., Treves A., Labita M., Kotilainen J. K., Scarpa R., 2010, *MNRAS*, 402, 2453
- Dekel A., Silk J., 1986, *ApJ*, 303, 39
- Dekel A., Burkert A., 2014, *MNRAS*, 438, 1870
- Di Matteo T., Springel V., Hernquist L. E., 2005, *Nature*, 433, 604
- Di Matteo T., Colberg J., Springel V., Hernquist L., Sijacki D., 2008, *ApJ*, 676, 33
- Di Matteo T., Khandai N., DeGraf C., Feng Y., Croft R. A. C., Lopez J., Springel V., 2012, *ApJ*, 745, L29
- Dubois Y., Teyssier R., 2008, *A&A*, 477, 79
- Dubois Y., Devriendt J., Slyz A., Teyssier R., 2010, *MNRAS*, 409, 985
- Dubois Y., Devriendt J., Slyz A., Teyssier R., 2012, *MNRAS*, 420, 2662
- Dubois Y., Pichon C., Devriendt J., Silk J., Haehnelt M., Kimm T., Slyz A., 2013, *MNRAS*, 428, 2885
- Dubois Y. et al., 2014, *MNRAS*, 444, 1453
- Dubois Y., Volonteri M., Silk J., Devriendt J., Slyz A., Teyssier R., 2015, *MNRAS*, 452, 1502
- Dunn J. P. et al., 2010, *ApJ*, 709, 611
- Efstathiou G., 2000, *MNRAS*, 317, 697
- Fabian A., 2012, *ARA&A*, 50, 455
- Fabian A. C., Celotti A., Erlund M. C., 2006, *MNRAS*, 373, L16
- Farrah D. et al., 2012, *ApJ*, 745, 178
- Feldmann R., Mayer L., 2014, *MNRAS*, 446, 1939
- Ferrarese L., Merritt D., 2000, *ApJ*, 539, L9
- Foglizzo T., Galletti P., Ruffert M., 2005, *A&A*, 435, 397
- Franx M., van Dokkum P. G., Förster Schreiber N. M., Wuyts S., Labbé I., Toft S., 2008, *ApJ*, 688, 770
- Gabor J. M., Bournaud F., 2014, *MNRAS*, 441, 1615
- Gaibler V., Khochfar S., Krause M., Silk J., 2012, *MNRAS*, 425, 438
- Gebhardt K. et al., 2000, *ApJ*, 539, L13
- Geen S., Slyz A., Devriendt J., 2013, *MNRAS*, 429, 633
- Gene S. et al., 2014, *MNRAS*, 445, 175
- Girardi L., Bressan A., Bertelli G., Chiosi C., 2000, *A&AS*, 141, 371
- Gnedin N. Y., 2000, *ApJ*, 542, 535
- González V., Labbé I., Bouwens R. J., Illingworth G., Franx M., Kriek M., 2011, *ApJ*, 735, L34
- Governato F., Willman B., Mayer L., Brooks A., Stinson G., Valenzuela O., Wadsley J., Quinn T., 2007, *MNRAS*, 374, 1479
- Haardt F., Madau P., 1996, *ApJ*, 461, 20
- Habouzit M., Volonteri M., Dubois Y., 2017, *MNRAS*, 468, 3935
- Heckman T. M., Kauffmann G., 2006, *New Astron. Rev.*, 50, 677
- Heckman T. M., Lehnert M. D., Strickland D. K., Armus L., 2000, *ApJS*, 129, 493
- Hopkins P. F., Keres D., Onorbe J., Faucher-Giguere C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Kauffmann G., White S. D. M., Guiderdoni B., 1993, *MNRAS*, 264, 201
- Kaviraj S. et al., 2017, *MNRAS*, 467, 4739
- Keller B. W., Wadsley J., Couchman H. M. P., 2015, *MNRAS*, 453, 3500
- Kennicutt R. C., Jr, 1998, *ApJ*, 498, 541
- Khandai N., Di Matteo T., Croft R., Wilkins S., Feng Y., Tucker E., DeGraf C., Liu M. S., 2015, *MNRAS*, 450, 1349
- Kimm T., Cen R., Devriendt J., Dubois Y., Slyz A., 2015, *MNRAS*, 451, 2900
- Komatsu E. et al., 2011, *ApJS*, 192, 18
- Krumholz M. R., Tan J. C., 2007, *ApJ*, 654, 304
- Lauer T. R. et al., 2007, *ApJ*, 662, 808
- Leitherer C., Robert C., Drissen L., 1992, *ApJ*, 401, 596
- Leitherer C. et al., 1999, *ApJS*, 123, 3
- Leitherer C., Ortiz Otálvaro P. A., Bresolin F., Kudritzki R.-P., Lo Faro B., Pauldrach A. W. A., Pettini M., Rix S. A., 2010, *ApJS*, 189, 309
- Lutz D. et al., 2008, *ApJ*, 684, 853
- Mac Low M.-M., Ferrara A., 1999, *ApJ*, 513, 142
- Maiolino R. et al., 2012, *MNRAS*, 425, 66
- Mandelbaum R., Wang W., Zu Y., White S., Henriques B., More S., 2016, *MNRAS*, 457, 3200
- Matteucci F., Greggio L., 1986, *A&A*, 154, 279
- Merloni A. et al., 2010, *ApJ*, 708, 137
- Moe M., Arav N., Bautista M. A., Korista K. T., 2009, *ApJ*, 706, 525
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Moustakas J. et al., 2013, *ApJ*, 767, 50
- Murray N., Quataert E., Thompson T. A., 2005, *ApJ*, 618, 569
- Nardini E., Risaliti G., Salvati M., Sani E., Imanishi M., Marconi A., Maiolino R., 2008, *MNRAS*, 385
- Nelson D., Genel S., Vogelsberger M., Springel V., Sijacki D., Torrey P., Hernquist L., 2015, *MNRAS*, 448, 59

- Nomoto K., Saio H., Kato M., Hachisu I., 2007, *ApJ*, 663, 1269
- Okamoto T., Gao L., Theuns T., 2008, *MNRAS*, 390, 920
- Omma H., Binney J., Bryan G., Slyz A., 2004, *MNRAS*, 348, 1105
- Page M. J. et al., 2012, *Nature*, 485, 213
- Peirani S. et al., 2016, *MNRAS*, preprint ([arXiv:1611.09922](https://arxiv.org/abs/1611.09922))
- Peng Y.-j. et al., 2010, *ApJ*, 721, 193
- Pontzen A., Tremmel M., Roth N., Peiris H. V., Saintonge A., Volonteri M., Quinn T., Governato F., 2017, *MNRAS*, 465, 547
- Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, *MNRAS*, 338, 14
- Prunet S., Pichon C., Aubert D., Pogosyan D., Teyssier R., Gottloeber S., 2008, *ApJS*, 178, 179
- Quintero A. D. et al., 2004, *ApJ*, 602, 190
- Rafferty D. A., McNamara B. R., Nulsen P. E. J., Wise M. W., 2006, *ApJ*, 652, 216
- Rasera Y., Teyssier R., 2006, *A&A*, 445, 1
- Reeves J. N., O'Brien P. T., Ward M. J., 2003, *ApJ*, 593, L65
- Roos O., Juneau S., Bournaud F., Gabor J. M., 2015, *ApJ*, 800, 19
- Rosas-Guevara Y., Bower R. G., Schaye J., McAlpine S., Dalla Vecchia C., Frenk C. S., Schaller M., Theuns T., 2016, *MNRAS*, 462, 190
- Rosdahl J., Schaye J., Dubois Y., Kimm T., Teyssier R., 2017, *MNRAS*, 466, 11
- Saez C., Chartas G., Brandt W. N., 2009, *ApJ*, 697, 194
- Salpeter E. E., 1954, *ApJ*, 121, 161
- Santini P. et al., 2012, *A&A*, 540, A109
- Schawinski K. et al., 2014, *MNRAS*, 440, 889
- Schaye J. et al., 2014, *MNRAS*, 446, 521
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
- Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
- Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, 452, 575
- Silk J., 2005, *MNRAS*, 364, 1337
- Silk J., Rees M. J., 1998, *MNRAS*, 324, 128
- Somerville R. S., 2002, *ApJ*, 572, L23
- Song M. et al., 2016, *ApJ*, 825, 5
- Springel V., Di Matteo T., Hernquist L., 2005, *MNRAS*, 361, 776
- Sturm E. et al., 2011, *ApJ*, 733, L16
- Sutherland R. S., Dopita M. A., 1993, *ApJS*, 88, 253
- Taylor P., Kobayashi C., 2015, *MNRAS*, 452, L59
- Teyssier R., 2002, *A&A*, 385, 337
- Thomas D., Thomas D., Maraston C., Maraston C., Bender R., Bender R., de Oliveira C. M., de Oliveira C. M., 2005, *ApJ*, 621, 673
- Tomczak A. R. et al., 2014, *ApJ*, 783, 85
- Tomczak A. R. et al., 2016, *ApJ*, 817, 118
- Tremonti C. A., Moustakas J., Diamond-Stanic A. M., 2007, *ApJ*, 663, L77
- Tweed D., Devriendt J., Blaizot J., Colombi S., Slyz A., 2009, *A&A*, 506, 647
- Vassiliadis E., Wood P. R., 1993, *ApJ*, 413, 641
- Veilleux S., Cecil G., Bland-Hawthorn J., 2005, *ARA&A*, 43, 769
- Vogelsberger M. et al., 2014a, *MNRAS*, 444, 1518
- Vogelsberger M. et al., 2014b, *Nature*, 509, 177
- Volonteri M., Stark D. P., 2011, *MNRAS*, 417, 2085
- Volonteri M., Dubois Y., Pichon C., Devriendt J., 2016, *MNRAS*, 460, 2979
- Weiner B. J. et al., 2009, *ApJ*, 692, 187
- White S. D. M., Frenk C. S., 1991, *ApJ*, 379, 52
- Williams R. J., Maiolino R., Krongold Y., Carniani S., Cresci G., Mannucci F., Marconi A., 2017, *MNRAS*, 467, 3399
- Zolotov A. et al., 2015, *MNRAS*, 450, 2327

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