

A hybrid active galactic nucleus feedback model with spinning black holes, winds and jets

Filip Huško¹,¹★ Cedric G. Lacey², Joop Schaye³,¹ Matthieu Schaller⁴,^{1,3} Evgenii Chaikin⁵,¹ Sylvia Ploeckinger⁶,⁴ Alejandro Benítez Llambay⁷,⁵ Alexander J. Richings^{6,7} and James W. Trayford⁸

¹Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands

²Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK

³Lorentz Institute for Theoretical Physics, Leiden University, PO box 9506, NL-2300 RA Leiden, the Netherlands

⁴Department of Astrophysics, University of Vienna, Türkenschanzstrasse 17, A-1180 Vienna, Austria

⁵Dipartimento di Fisica G. Occhialini, Università Degli Studi di Milano Bicocca, Piazza della Scienza, 3 I-20126 Milano MI, Italy

⁶Centre for Data Science, Artificial Intelligence and Modelling, University of Hull, Cottingham Road, Hull HU6 7RX, UK

⁷E. A. Milne Centre for Astrophysics, University of Hull, Cottingham Road, Hull HU6 7RX, UK

⁸Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Burnaby Road, Portsmouth PO1 3FX, UK

Accepted 2026 February 10. Received 2026 January 18; in original form 2025 September 6

ABSTRACT

We present a hybrid active galactic nucleus (AGN) feedback model that features three accretion disc states (the thick, thin, and slim discs at low, moderate, and super-Eddington accretion rates, respectively), and two feedback modes: thermal isotropic and kinetic jets. The model includes black hole (BH) spin evolution due to gas accretion, BH mergers, jet spin-down, and Lense–Thirring torques. The BH spin determines the jet directions and affects the feedback efficiencies. The model is implemented in the SWIFT code and coupled with the COLIBRE galaxy formation model. We present the first results from hybrid AGN feedback simulations run as part of the COLIBRE suite, focusing on the impact of new parameters and calibration efforts. Using the new hybrid AGN feedback model, we find that AGN feedback affects not just massive galaxies, but all galaxies down to $M_* \approx 10^8 M_\odot$. BH spins are predicted to be near-maximal for intermediate-mass BHs ($M_{\text{BH}} \in [10^6, 10^8] M_\odot$), and lower for other BH masses, which is in good agreement with observations. The intergalactic medium is hotter and impacted on larger scales in the hybrid AGN feedback simulations compared to those using purely thermal feedback. In the hybrid AGN simulations, we predict that half of the cumulative injected AGN energy is in thermal and the other half in jet form, broadly independent of BH mass and redshift. Jet feedback is important at all redshifts and dominates over thermal feedback at $z < 0.5$ and $z > 1.5$, but only mildly.

Key words: galaxies: evolution – galaxies: formation – galaxies: jets – (galaxies:) quasars: supermassive black holes.

1 INTRODUCTION

Observations reveal a diversity in the star-formation (SF) activity of galaxies, with lower-mass galaxies typically having higher specific star formation rates (sSFRs), i.e. SFRs per unit stellar mass, than high-mass galaxies (e.g. A. E. Bauer et al. 2013; Y.-Y. Chang et al. 2015). High-mass ellipticals are typically gas-poor (e.g. T. Wiklund, F. Combes & C. Henkel 1995; L. M. Young et al. 2011; T. A. Davis et al. 2019) and have very low or undetectable levels of SF (e.g. S. Salim et al. 2007; K. E. Whitaker et al. 2012). Brightest-cluster galaxies (BCGs), the most massive elliptical galaxies in the Universe, are almost all devoid of SF (A. Fraser-McKelvie, M. J. I. Brown & K. A. Pimbblet 2014). The quiescent nature of these galaxies poses a problem. They are located in the centres

of galaxy clusters and are surrounded by hot, dense, pressure-supported gas, the intracluster medium (ICM), which emits X-rays (e.g. C. L. Sarazin 1986; R. Stanek et al. 2006). The typical cooling time of the ICM close to the BCG is shorter than a Hubble time. Without mitigation, we would expect massive cooling flows to develop, causing large molecular hydrogen reservoirs and starbursts within the BCG, with SFRs $\approx 1000 M_\odot \text{ yr}^{-1}$ (A. C. Fabian 1994; D. S. Hudson et al. 2010). With the exception of a few cases (e.g. A. C. Edge 2001; A. Fraser-McKelvie et al. 2014; M. McDonald et al. 2015), this is not observed.

The lack of SF in massive galaxies can be attributed to energy injection by supermassive black holes (hereafter BHs), i.e. active galactic nucleus (AGN) feedback (see review by A. C. Fabian 2012). X-ray observations of galaxy groups and clusters reveal cavities in the ICM (e.g. S. F. Gull & K. J. E. Northover 1973; H. Boehringer et al. 1993; L. Birzan et al. 2004; D. Eckert et al. 2021). These cavities are cospatial with lobes inflated by jets

* E-mail: husko@strw.leidenuniv.nl

of relativistic particles launched by the central BHs, which can mainly be observed at radio wavelengths (e.g. R. D. Blandford & A. Königl 1979; P. L. Biermann & P. A. Strittmatter 1987; C. M. Urry & P. Padovani 1995; S. Markoff, H. Falcke & R. Fender 2001). The jet powers, inferred from X-ray observations of cavities, are sufficient to offset the cooling luminosities of the ICM (e.g. D. A. Rafferty et al. 2006; J. Hlavacek-Larrondo et al. 2012; H. R. Russell et al. 2013), solving the aforementioned cooling flow problem. AGN feedback in this situation, as well as in other, lower-mass quenched galaxies, is known as radio jet, ‘maintenance mode’, or mechanical feedback (e.g. A. C. Fabian 2012; D. Eckert et al. 2021).

In addition to radio jets, AGN have been observed at other wavelengths (e.g. optical, X-ray, UV, mid-infrared, etc.; see compilation by X. Shen et al. 2020). The brightest cases are known as quasars (e.g. M. Schmidt 1968). These observations extend to the highest redshifts (e.g. X. Fan et al. 2001), with the peak in radiative AGN activity occurring around $z = 2$ (e.g. I. Delvecchio et al. 2014; J. Aird et al. 2015; J. C. J. D’Silva et al. 2023). AGN are observed to drive outflows (e.g. D. M. Crenshaw, S. B. Kraemer & I. M. George 2003; F. Tombesi et al. 2010; C. Feruglio et al. 2010). Their origin is unknown, with the radiation possibly driving a wind on the scales of the accretion disc surrounding the BH (e.g. N. Murray et al. 1995; M. Nomura et al. 2016; F. Fiore et al. 2017; A. Quera-Bofarull et al. 2023), or on larger scales (e.g. A. C. Fabian 1999; T. Costa et al. 2018). What fraction of the luminosity couples to the gas is also uncertain (e.g. C. M. Harrison et al. 2018). This mode of AGN feedback, often referred to as ‘quasar’ feedback (T. M. Heckman & P. N. Best 2014), should be more relevant at higher redshifts, when radiative AGN activity is stronger.

In the standard picture, radiative AGN eject gas from their galaxies/host dark matter (DM) haloes and thus quench them (e.g. T. Di Matteo, V. Springel & L. Hernquist 2005), while maintenance-mode (jet) feedback, which is more often observed in the local Universe, is thought to keep the galaxies quenched (e.g. R. G. Bower et al. 2006), with almost all massive quenched galaxies having some radio emission, even if at a low level (J. Sabater et al. 2019). In this picture, BHs occupy one of the two different feedback modes depending on their accretion rate expressed as a fraction of the Eddington rate (the Eddington ratio). They are thought to be in the radiatively efficient quasar mode if the Eddington ratio is above 1 per cent, and in the maintenance mode if it is below that value (T. M. Heckman & P. N. Best 2014).

The above picture of two-mode feedback, in separate Eddington ratio regimes, is very likely an oversimplification. Of order 10 per cent of AGN in the radiatively efficient, quasar regime, are observed to be radio-loud (T. M. Heckman & P. N. Best 2014), i.e. they have strong jets. In fact, some of the earliest observed quasars have these strong jets (e.g. J. L. Greenstein & M. Schmidt 1964; M. Kino & N. Kawakatu 2005). While these radio sources are relatively rare and constitute only a small fraction of all radio sources, they are very powerful. These radio sources are found to be hosted by active rather than quenched galaxies (D. J. Saikia 2022), unlike the majority of low-power radio sources (J. Sabater et al. 2019). The difference in host galaxy properties also implies that BHs in quasars have higher accretion rates, which in turn means that their radio jets occur in a fundamentally different accretion and feedback state to the one that causes ‘maintenance mode’ feedback in quenched galaxies. Conceivably, jets from quasars could be as important for galaxy evolution as the lower-power jets in the standard ‘maintenance mode’ (A. Merloni & S. Heinz 2008).

Conversely, low-power winds are observed not just in the high-Eddington ratio, quasar regime, but also at low Eddington ratios, in which case, these objects are referred to as ‘red geysers’ (E. Cheung et al. 2016; F. Shi et al. 2025).

Furthermore, the properties of AGN may change at super-Eddington rates, with both jets and winds possibly becoming relevant, and the radiative efficiency dropping (e.g. M. A. Abramowicz et al. 1988; A. Sądowski et al. 2014; also see review by P. Marziani et al. 2025). The super-Eddington regime is very rare, especially in the local Universe, but the accretion rates can be so high that this regime may be important for BH and galaxy evolution, despite its short-lived and infrequent nature (F. Huško et al. 2025). Powerful jets, possibly launched in the super-Eddington state, have been found not only at redshifts up to $z = 3$ (V. Smolčić et al. 2017; B. Šlaus et al. 2024), but also up to $z = 6$ (G. Ghisellini et al. 2014; T. Sbarrato et al. 2021; T. Sbarrato et al. 2022; N. Roy et al. 2024).

Accretion theory has had great success in explaining many aspects of the AGN observations described above. Standard accretion disc theory (I. D. Novikov & K. S. Thorne 1973; N. I. Shakura & R. A. Sunyaev 1973) has been used to explain observed AGN in the moderate Eddington-ratio regime, including the details of their spectra (e.g. W.-H. Sun & M. A. Malkan 1989). In this accretion model, which we will refer to as the ‘thin disc’ in this study, the accretion disc is relatively cool, geometrically thin, optically thick and radiatively efficient. Of order 10 per cent (the precise value depends on BH spin and is between ≈ 4 and ≈ 40 per cent; I. D. Novikov & K. S. Thorne 1973) of the mass-energy accreting through the disc is released by means of viscous heating and then radiated away.

The theory of the advection-dominated accretion flow (ADAF; R. Narayan & I. Yi 1994), which we will refer to as the ‘thick disc’, describes a hot, geometrically thick, optically thin and radiatively inefficient accretion flow. This model corresponds to the low-Eddington ratio regime, in which the AGN are not bright in the optical or X-rays, but instead produce energetically dominant jets. In all accretion discs, the magnetic dynamo effect can lead to an increase of the magnetic field strength. However, in the ADAF, the field is easily advected inwards (F. Yuan & R. Narayan 2014). This results in the launching of a jet through the R. D. Blandford & R. L. Znajek (1977) (BZ) mechanism. The BZ jet efficiency depends strongly on BH spin, as it is the rotational energy of the BH that powers jet launching. The ADAF model predicts that the transition between the thick and thin discs should occur at ≈ 1 per cent Eddington (R. Narayan & I. Yi 1995), which is consistent with observations of X-ray binaries and AGN (T. J. Maccarone, E. Gallo & R. Fender 2003; C. Done, M. Gierliński & A. Kubota 2007; H. R. Russell et al. 2013; H. Noda & C. Done 2018).

In addition to the thin and thick discs, other models have also been developed. The advection-dominated, inflow–outflow solution (ADIOS; R. D. Blandford & M. C. Begelman 1999) predicts a similar accretion flow as the thick disc (ADAF), but it features strong outflows on account of thermal pressure of the gas. This can explain observations of red geysers (E. Cheung et al. 2016). One of the main features of this model is that the winds take away most of the accreting mass, so that the horizon-scale mass accretion rate at the inner edge of the accretion disc and into the BH is much smaller than that at large radii (on to the disc). In other words, the accretion efficiency is much less than unity. The ADIOS model is preferred over the ADAF one in observations of individual objects such as M87* and SgrA* (F. Yuan 2003; C. Y. Kuo et al. 2014), cavity power observations in galaxy groups and

clusters (R. S. Nemmen & A. Tchekhovskoy 2015), and numerical simulations (F. Yuan, D. Bu & M. Wu 2012).

The slim disc model (e.g. M. A. Abramowicz et al. 1988; J.-M. Wang & Y.-Y. Zhou 1999; A. Sądowski 2009) has been used to describe the super-Eddington regime. In the slim disc, the pressure is dominated by radiation rather than gas, but the disc is otherwise quite similar to the thick disc. Most notably, it is highly advective, which should result (in combination with the dynamo effect) in similar jet efficiencies as in the thick disc.

Recent developments in numerical simulation methods have allowed direct tests of accretion disc models. The exception to this is of very thin discs corresponding to the radiatively efficient thin disc model (N. I. Shakura & R. A. Sunyaev 1973), which cannot yet be simulated due to their very thin nature. Recent simulations of accretion disc formation from cosmological initial conditions find a hyper-magnetized disc state at moderate Eddington ratios, which may supplant the original N. I. Shakura & R. A. Sunyaev (1973) solution (P. F. Hopkins et al. 2024; P. F. Hopkins 2025). General-relativistic magneto-hydrodynamical (GRMHD) simulations of thick discs have confirmed most of the properties predicted by the ADAF/ADIOS models (e.g. R. Narayan, I. V. Igumenshchev & M. A. Abramowicz 2003; A. Tchekhovskoy, R. Narayan & J. C. McKinney 2010; J. C. McKinney, A. Tchekhovskoy & R. D. Blandford 2012; A. Sądowski et al. 2013). These simulations have also confirmed the BZ mechanism of jet launching, and they have converged on a fitting formula for the jet efficiency that is more accurate than the BZ prediction (R. Narayan et al. 2022). The jet efficiency in this formula depends strongly on BH spin and on the magnetization of the disc. It can reach values of up to several hundred per cent. This may appear to violate energy conservation. However, the efficiencies are defined as the power output relative to the mass-energy input into the BH. BHs can tap their own rotational energy through jet launching and lose more energy in the process than they gain mass-energy through accretion. In addition to confirming the operation of the BZ process, GRMHD simulations show that very mass-loaded winds are operating in the thick disc (F. Yuan et al. 2012; M. Guo et al. 2023), resulting in accretion efficiencies of a per cent or less.

GRMHD simulations reveal two classes of objects that differ based on the strength of the magnetic field near the event horizon. The first are so-called ‘standard and normal evolution’ (SANE) discs (e.g. R. Narayan et al. 2012), while the second are magnetically arrested (MAD) discs (e.g. R. Narayan et al. 2003). In the SANE state, the event horizon magnetization is a free parameter, while in the MAD state, it is saturated and is determined by the accretion rate and BH spin. We expect most thick discs to be in the MAD state on theoretical grounds, due to strong advection (F. Yuan & R. Narayan 2014), and this is supported by observations (Event Horizon Telescope Collaboration et al. 2021; K. Akiyama et al. 2022; F. Yuan, H. Wang & H. Yang 2022). GRMHD simulations of super-Eddington, slim discs show that jet efficiencies can be similarly high as in the thick disc. We also expect them to be MAD, as they are advection-dominated. GRMHD simulations of thinner, sub-Eddington accretion discs (M. J. Avara, J. C. McKinney & C. S. Reynolds 2016; A. Ricarte, R. Narayan & B. Curd 2023) have found that the MAD state is possible in the thin disc, with jet efficiencies given by the same BZ formula as for the thick and slim discs. However, the magnetization, and therefore the jet efficiency, is lower, and heavily dependent on the Eddington ratio. Furthermore, it is unknown if the MAD assumption is correct for thin discs, as they are not advection-dominated.

Models of galaxy formation, starting with semi-analytical models (R. G. Bower et al. 2006; D. J. Croton et al. 2006; C. D. P. Lagos, S. A. Cora & N. D. Padilla 2008) and early hydrodynamical simulations (T. Di Matteo et al. 2005; C. M. Booth & J. Schaye 2009), began incorporating AGN feedback after X-ray observations showed that it is important in offsetting cooling flows in galaxy clusters (e.g. L. Birzan et al. 2004; B. R. McNamara et al. 2005; M. W. Wise et al. 2007). Many cosmological hydrodynamical simulation models have included AGN feedback using a single, typically thermal AGN feedback mode, corresponding to the high-Eddington ratio, ‘quasar’ mode (e.g. OWLS: J. Schaye et al. 2010; Magneticum: M. Hirschmann et al. 2014; EAGLE: J. Schaye et al. 2015; MassiveBlack-II: N. Khandai et al. 2015; Romulus: M. Tremmel et al. 2017; Astrid: S. Bird et al. 2022; Y. Ni et al. 2022). Other simulation models have often used two-mode AGN feedback where at low-Eddington ratios, the ‘quasar’ mode is replaced by either winds representing the ADIOS outflow (IllustrisTNG; R. Weinberger et al. 2017; D. Nelson et al. 2019), manual bubble injection representing lobes inflated by jets (e.g. D. Sijacki et al. 2007; Illustris; M. Vogelsberger et al. 2014), or direct kinetic jet feedback (HorizonAGN: S. Kaviraj et al. 2017; NewHorizon: Y. Dubois et al. 2021; SIMBA: R. Davé et al. 2019). Some simulation models have included additional feedback modes (e.g. X-ray heating in SIMBA). Sophisticated sub-grid models have been developed to more realistically represent AGN feedback, including accretion disc modelling and BH spin evolution (e.g. M. Volonteri et al. 2005; C. D. P. Lagos, N. D. Padilla & S. A. Cora 2009; N. Fanidakis et al. 2011; Y. Dubois et al. 2012; L. K. Steinborn et al. 2015; D. Fiacconi, D. Sijacki & J. E. Pringle 2018; S. Bustamante & V. Springel 2019; A. J. Griffin et al. 2019; F. Huško et al. 2022; S. Koudmani et al. 2024; L. Sala et al. 2024), as well as inclusion of super-Eddington accretion and its feedback effects (e.g. J. S. Bennett et al. 2024; D. Rennehan et al. 2024; F. Huško et al. 2025).

Motivated by observational evidence of the complexity of AGN feedback, as well as advances in theoretical modelling, we have developed a hybrid AGN feedback model for use in cosmological, hydrodynamical simulations of galaxy formation. The model synthesizes different aspects of AGN activity that have been studied with accretion disc theory, including recent developments from GRMHD simulations. We include three accretion disc states: the thick, thin and slim disc, at low, moderate and super-Eddington ratios, respectively. Two feedback modes are included: (1) thermal isotropic, representing the effects of radiation and winds, and (2) kinetic jets, representing the effects of relativistic jets. We also model BH spin to more realistically implement AGN feedback, especially for jets. In addition to its usefulness as an AGN feedback model, galaxy formation simulations run with the model can be used for detailed studies of BH and accretion physics in a cosmological context.

Aspects of the model have already been introduced in previous studies, where we coupled it with the SWIFT-EAGLE model (J. Schaye et al. 2015; M. Schaller et al. 2024). In F. Huško et al. (2022), we studied BH spin and jet feedback from the thick disc using idealized galaxy clusters. We found that spin-driven jets can quench BCGs, even in very massive galaxy clusters. In F. Huško et al. (2024), we compared this mode with the thermal isotropic one from the thin disc, as well as studied a simple hybrid model that used kinetic jets at low Eddington ratios and thermal isotropic feedback at large Eddington ratios. The hybrid model performed best, yielding moderate SFRs in the BCGs (but not

permanent quenching), and the lowest central entropies of the ICM. In F. Huško et al. (2025) we introduced the super-Eddington aspect of the model (the slim disc), with both jet and thermal feedback, and also considered jets in the thin disc. In F. Huško & C. G. Lacey (2023a), we performed the first idealized jet simulations using smoothed particle hydrodynamics (SPH) to validate the numerical behaviour of jets and the lobes they inflate. We found that most aspects of their behaviour are captured well with as few particles as a few hundred.

Here, we combine the different aspects of the hybrid AGN feedback model discussed above into a three-state model. We also add winds to the thick disc at low Eddington ratios, and introduce accretion efficiencies, which take account of the fact that the vast majority of the mass accreting through a thick or slim disc is blown away by these winds. The model builds upon the AGN feedback model used for the OWLs and EAGLE simulations (C. M. Booth & J. Schaye 2009; C. M. Booth & J. Schaye 2010; J. Schaye et al. 2015). It has been integrated into the SWIFT simulation code and included in the COLIBRE,¹ galaxy formation simulation project (J. Schaye et al. 2025), as an alternative AGN feedback model alongside its primary thermal-only model. In Fig. 1, we show the widespread and large-scale impact of jet activity at $z = 0.2$ in one of our hybrid AGN feedback simulations run as part of the COLIBRE project. We describe COLIBRE in Section 2. The main purpose of this study is to provide a comprehensive and self-contained description of the hybrid AGN feedback model (Section 3), as well as to describe its calibration as part of the COLIBRE project (Section 4). We also provide predictions for BH spin (Section 5) and a first analysis of the operation of AGN feedback in the hybrid simulations (Section 6). We compare with the simulations using purely thermal AGN feedback, focusing on the energetics of feedback. In Section 7, we summarize and conclude.

2 THE COLIBRE MODEL AND SIMULATIONS

Here we will describe the simulation methods and the COLIBRE galaxy formation model. We provide a summary, and refer the reader to J. Schaye et al. (2025), as well as individual methods papers, for further details. The simulations are performed using a version of the SWIFT² code (M. Schaller et al. 2024) in which the COLIBRE model has been included. In Section 2.1, we first describe the broad aspects of the simulation methods. In Section 2.2, we present the different resolutions and simulation runs that are used in this study. In Section 2.3, we summarize the COLIBRE subgrid model. BHs and AGN feedback are also included in the fiducial thermal model, but we describe these separately in Section 2.4, and we do so in more detail than for other aspects of COLIBRE, since BH physics is particularly relevant for this study.

2.1 Simulation methods

2.1.1 Initial conditions

The simulations include four particle types: DM, gas, stars, and BHs. The first two are included in the initial conditions (ICs), which are created at redshift $z = 63$, while the second two are created during the simulation. Four times as many DM particles

as gas particles are included in the ICs, which results in the particle masses of the two types being nearly equal. The large number of DM particles suppresses spurious energy transfer from DM to baryons in galaxies (A. D. Ludlow et al. 2019, 2023), which can affect galaxy properties negatively. The ICs are created using MONOFONIC (O. Hahn et al. 2020; M. Michaux et al. 2021).

2.1.2 Gravity and hydrodynamics

Gravitational forces are computed using a fourth-order fast multipole method (L. Greengard & V. Rokhlin 1987; W. Dehnen 2014), and softened using a softening length ϵ . The effects of massive neutrinos, which to first order influence the expansion rate of the universe, are included using the linear response method (Y. Ali-Haïmoud & S. Bird 2013). The equations of hydrodynamics are solved using the SPHENIX SPH scheme (J. Borrow et al. 2022). SPHENIX is a density-energy SPH scheme that includes artificial viscosity and conduction, as well as respective limiters of the two.

2.1.3 Structure finding

As the simulation progresses, a friends-of-friends (FoF) algorithm is run to identify groups, but also to seed BHs (Section 2.4.1). The FoF algorithm is run on the DM particles starting from $a = 0.05$, after every $\Delta a = 0.00751a$, using a linking length of 0.2 times the mean distance between them. Other particles are ‘attached’ to FoF groups based on their nearest DM particles within the linking length. The HBT-HERONS (V. J. Forouhar Moreno et al. 2025) subhalo finder, an updated version of HBT+ (J. Han et al. 2018), is then run using the catalogues and the snapshots. FoF groups with fewer than 32 particles of all types are discarded when constructing halo catalogues. HBT-HERONS, unlike most other halo finders, uses temporal information to assign hierarchical (sub)halo membership, alongside dynamical boundness criteria. Finally, we run the SOAP post-processing tool (R. McGibbon et al. 2025) to calculate all relevant galaxy and halo properties in spherical and projected apertures of fixed proper size, inclusive or exclusive of unbound particles and self-bound substructures. The fiducial choice for galaxy properties in this study is spherical exclusive apertures of 50 proper kpc (hereafter pkpc) in radius. These are centred on the most bound particle of each subhalo.

2.2 Runs

The COLIBRE simulations assume the ‘ $3 \times 2pt+$ All Ext.’, spatially flat Λ CDM cosmology from the Dark Energy Survey (DES) year three results (T. M. C. Abbott et al. 2022). The parameter values are $h = 0.681$, $\Omega_m = 0.306$, $\Omega_{\text{CDM}} = 0.256011$, $\Omega_b = 0.0486$, $\Omega_\Lambda = 0.693922$, $\Omega_\gamma = 0.000053$. For the definitions of these parameters, as well as other cosmological parameters and their values, see either J. Schaye et al. (2025) or T. M. C. Abbott et al. (2022).

In this study, we mostly use simulations with a (50 Mpc)³ comoving volume (hereafter L050) at $m7$ resolution (hereafter L050m7). Here, ‘ $m7$ ’ refers to the initial mean particle mass (i.e. the mass resolution), which is $m_g = 1.47 \times 10^7 M_\odot$ for gas particles, while DM particles have $m_{\text{DM}} = 1.94 \times 10^7 M_\odot$. The gravitational softening length for all particles is $\epsilon = 1.4$ kpc in proper coordinates at $z < 1.57$, and $\epsilon = 3.6$ kpc in comoving coordinates at higher redshifts. We also run simulations at $m6$ and $m5$ resolutions, for which the masses of all particles are

¹<https://colibre-simulations.org>

²<https://swiftsim.com>




Figure 1. Visualization of jet activity in the L100m6h simulation at $z = 0.2$ in proper coordinates. The background image shows the cosmic web using the gas surface density, while particles kicked into jets by BHs are displayed using both information on the elapsed time since when they were kicked into jets (colour) and their surface density (opacity). Side panels zoom in on individual jets of interest, with further zoom-ins on to their host galaxies (except in the case of a cluster and its BCG) displayed using luminosities in the SDSS i , r , and g bands assigned to RGB colours.

decreased by successive factors of 8; the gas particle masses are $m_g = 1.84 \times 10^6 M_\odot$ and $m_g = 2.3 \times 10^5 M_\odot$, respectively. The gravitational softening lengths are reduced by successive factors of two relative to m7.

In addition to L050, we used L025 and L012.5 simulations at m6 and m5 resolution during the calibration process. In this paper we also present results from L100 (at m6 and m7) and L200 (at m7) volumes run with the hybrid AGN feedback model, which we refer to as L100m6h, L100m7h and L200m7h. Here, ‘h’ refers to the fact that these were run with the hybrid model, unlike the fiducial simulations that use purely thermal AGN feedback.

2.3 Subgrid physics

2.3.1 Radiative cooling and heating

Unlike in many previous large-volume cosmological simulations, gas is allowed to cool down to $T \approx 10$ K, i.e. no pressure floor is applied to prevent cooling below warm interstellar medium (ISM) temperatures ($\approx 10^4$ K). The radiative cooling and heating rates are calculated with HYBRID-CHIMES (S. Ploekinger et al. 2025), which uses the CHIMES reaction network (A. J. Richings, J. Schaye & B. D. Oppenheimer 2014a, b) to calculate the species abundances of electrons and nine hydrogen and helium species

in non-equilibrium. The contribution to cooling and heating due to an additional 147 metal species is accounted for using tables that assume steady-state chemistry and ionization equilibrium. In HYBRID-CHIMES, the pre-tabulated metal cooling rates are corrected by the non-equilibrium electron density throughout the simulation, therefore also including important non-equilibrium effects in the metal cooling rates. The parametrization of the interstellar radiation field, the cosmic ray rate, and the shielding column density is linked to the local Jeans length, as introduced in S. Ploeckinger & J. Schaye (2020), with the updates described in S. Ploeckinger et al. (2025).

2.3.2 Chemistry, dust, and diffusion

Chemical enrichment of gas by heavy elements (described in detail in Correa et al. 2025, in preparation) is assumed to occur due to mass-loss of Asymptotic Giant Branch (AGB) stars, massive stars, core-collapse supernovae (CCSN), SNIa, neutron star mergers, common-envelope SN jets and collapsars. 14 elements are tracked, 11 of which contribute to radiative cooling and heating: H, He, C, N, O, Ne, Mg, S, Si, Ca, Fe (S and Ca are not explicitly tracked, but solar abundances relative to Si are assumed for these two elements). In addition, Ba, Sr and Eu are also tracked. The chemical yields of the processes that generate these elements depend on the metallicity of star particles (inherited from the gas metallicities of particles from which stars are born in COLIBRE). Gas particles gain mass in the simulation as they are enriched. Gas particles that reach a mass of $4m_g$, where m_g is the mean initial baryonic particle mass, are split into two.

The dust grain model, which is described in detail in J. W. Trayford et al. (2025), includes three grain compositions (graphite, forsterite, and fayalite) and 2 grain sizes (0.01 and 0.1 μm), and is coupled to the radiative cooling module. Gas particles track their fractions of mass contained in the relevant dust grains types. Grains are assumed to be seeded by AGB stars and CCSN using metallicity-dependent dust yields. The growth and destruction of dust grains in the ISM are modelled on the fly. The model includes accretion on to grains, as well as sputtering (destruction by collisions with hot gas particles), grain collisions (shattering and coagulation), and astration. Dust is also directly destroyed in gas particles that are affected by thermal SN and AGN feedback.

The turbulent diffusion model (Correa et al. 2025, in preparation) accounts for mixing below the resolution scale using the diffusion equation, with a diffusion coefficient that depends on the local turbulent velocity dispersion. It is applied to elemental mass fractions, mass fractions of dust grains of different compositions and sizes, and diagnostic tracers that follow material emerging through different stellar mass-loss channels. A time-step condition is applied to ensure that diffusion is sampled well.

2.3.3 Star formation and mass-loss

The star formation model is detailed in F. S. J. Nobels et al. (2024). Star formation follows the Schmidt law: $\dot{m}_* = \varepsilon_{\text{ff}} m_g / t_{\text{ff}}$, where t_{ff} the gravitational free-fall time-scale and $\varepsilon_{\text{ff}} = 0.01$ the star formation efficiency per free-fall time. Only gas that is gravitationally unstable on the scale of the SPH kernel is eligible to form stars, with both the local velocity dispersion and thermal pressure contributing to the stability of the gas against collapse. Stellar

particles represent simple stellar populations (SSPs) of subgrid stars whose masses are distributed according to the G. Chabrier (2003) initial mass function (IMF) in the range M_\odot . Once formed, star particles lose mass due to winds and SN. Most of the mass loss occurs in the first 40 Myr once a stellar particle has formed. To sample this mass loss phase well, star particles are limited to time-steps no longer than 1 Myr for the first 40 Myr of their evolution.

2.3.4 Stellar feedback

The main channel of stellar feedback in COLIBRE is SN (SNIa and CCSN, the latter being dominant). The energy associated with SNe is released stochastically, mostly in thermal form and using relatively large temperature increments so that numerical overcooling is suppressed (C. Dalla Vecchia & J. Schaye 2012). The gas neighbours targeted for feedback are chosen isotropically (see E. Chaikin et al. 2022). The heating temperatures used are given by

$$\Delta T_{\text{SN}} = 10^{6.5} \text{ K} \left(\frac{n_{\text{H}}}{n_{\text{H,pivot}}} \right)^{2/3}, \quad (1)$$

where n_{H} is the SPH-smoothed gas number density at the location of the star particle, and $n_{\text{H,pivot}}$ is a free parameter (different values are used for the thermal and hybrid AGN feedback models; see Section 4.5). The index of 2/3 is motivated by the derivation by C. Dalla Vecchia & J. Schaye (2012) for a minimum heating temperature required to prevent numerical overcooling. Minimum and maximum heating temperatures are also used, and they depend on resolution (see E. Chaikin et al. 2025; J. Schaye et al. 2025).

In addition to the variable heating temperature, the model includes variable SN energies (for CCSN feedback, but not for SNIa). The energies injected per unit time by CCSN are not equal to those expected from the SSP represented by the particle (given its age and metallicity), but rather f_E times that value. f_E for CCSN can be both smaller and larger than unity, and varies with the stellar birth pressure, with minimum and maximum values between 0.1 and 4 (but dependent on resolution). This variation is required because the stochastic thermal feedback may result in an over- or underestimate of the true radiative losses, which cannot be predicted robustly at our numerical resolution (E. Chaikin et al. 2025). Therefore, calibration of f_E is necessary if one wishes to reproduce both galaxy masses and sizes simultaneously. For CCSN feedback, a fraction $f_{\text{kin}} = 0.1$ of the energy is injected kinetically in the form of low-velocity kicks, using a target velocity of 50 km s^{-1} . The kicks are done in pairs, in random but opposite directions in a way that conserves energy, momentum and angular momentum (E. Chaikin et al. 2023). The role of kinetic feedback is to generate turbulence in the ISM, helping to stabilize the gas against star formation.

COLIBRE also includes early stellar feedback alongside CCSN (despite calling these ‘early’, these processes continue to operate after the SN phase). This includes: (1) stellar winds, (2) radiation pressure, and (3) H II regions. Early stellar feedback is detailed in A. Benítez-Llambay et al. (2025). Stellar winds and radiation pressure are both modelled as momentum injections, which depend on the metallicity and age of the SSP. These momentum injections are modelled as kinetic feedback with velocity increments of 50 km s^{-1} .

2.4 Black holes

Here, we summarize how supermassive BHs are treated in the fiducial COLIBRE model, which uses thermal AGN feedback. These prescriptions are described in more detail in J. Schaye et al. (2025) (see also F. S. J. Nobels et al. 2022; Y. M. Bahé et al. 2022). In Section 3, where we introduce the hybrid model of AGN feedback, we explain the differences of that model relative to the one presented in this section.

2.4.1 Seeding

BH seeding follows the approach introduced by T. Di Matteo et al. (2003): they are seeded in haloes identified by the FoF algorithm (Section 2.1) whose mass exceeds $M_{\text{FoF,seed}}$ and that do not yet have a BH. BHs are seeded with a mass of $M_{\text{BH,seed}}$ at the location of the densest gas particle in the FoF group, which is effectively converted into the collisionless BH particle. BHs carry a subgrid mass, M_{BH} , and a particle mass, $M_{\text{BH,part}}$ (initially equal to the mass of the gas particle that created the BH). The subgrid mass is the physical, true mass of the BH. In the early evolution of the BH, its subgrid mass is generally much lower than the mass of the other particles in the simulation. The particle mass is thus used for all N-body gravitational interactions with other particles within the code, since large differences between masses used in the gravity calculations lead to undesirable effects. As the BH grows through accretion, it eventually reaches $M_{\text{BH}} \approx M_{\text{BH,part}}$.

For m7 resolution, $M_{\text{FoF,seed}} = 5 \times 10^{10} M_{\odot}$, while for m6 and m5, $M_{\text{FoF,seed}} = 10^{10} M_{\odot}$. The BH seed mass, $M_{\text{BH,seed}}$, is a free parameter with values of order $\approx 10^4$ - $10^5 M_{\odot}$, corresponding to the direct collapse seed formation scenario. Its value affects the onset of BH growth and AGN feedback. It is used in the calibration of the model at each resolution (E. Chaikin et al. 2025), and different values are used for the thermal and hybrid AGN feedback simulations (Section 4.5).

In the hybrid simulations, we seed BHs with an initial spin value of $|a| = 0.01$, directed randomly (see Section 3, and more specifically Section 3.6, for details on BH spin). This effectively assumes that BHs are not spinning at the time they form, but we avoid using $|a| = 0$, since this can lead to numerically unstable expressions in our model. We tested seeding BHs with higher spins, up to $|a| = 1$, but found that this had no appreciable effect on either BH or galaxy evolution.

2.4.2 Gas accretion

BHs are assumed to only accrete from gas in COLIBRE. In the hybrid AGN feedback model, there are two relevant mass accretion rates: $\dot{M}_{\text{acc,d}}$, the accretion rate on to a subgrid accretion disc (hereafter the disc-scale accretion rate), and $\dot{M}_{\text{acc,H}}$, the accretion rate on to the BH itself (i.e. its event horizon, hereafter the horizon-scale accretion rate). The first of these accretion rates is generally agnostic to what occurs on subgrid (accretion-disc) scales, while the second can be heavily affected by assumptions about subgrid accretion discs and feedback. These two accretion rates can be very different in the hybrid model (see Section 3.3), since we include subgrid processes occurring on accretion disc scales in more detail in the hybrid model.

In the thermal AGN feedback model, mass is assumed to accrete unimpeded from disc scales all the way into the event horizon. In other words, $\dot{M}_{\text{acc,d}} = \dot{M}_{\text{acc,H}}$. The BH accretion rate in the

thermal AGN feedback model (and the disc-scale accretion rate in the hybrid model) is given by

$$\dot{M}_{\text{acc,d}} = \dot{M}_{\text{BHL}} \frac{f_{\text{turb}} f_{\text{ang}}}{(f_{\text{turb}}^2 + f_{\text{ang}}^2)^{1/2}}, \quad (2)$$

where \dot{M}_{BHL} is the Bondi–Hoyle–Lyttleton (BHL) accretion rate (H. Bondi 1952, F. Hoyle & R. A. Lyttleton 1939), given by

$$\dot{M}_{\text{BHL}} = 4\pi G^2 \frac{\rho_{\text{g}}}{c_{\text{s}}^3} M_{\text{BH}}^2, \quad (3)$$

with ρ_{g} and c_{s} the kernel-weighted gas density and sound speed at the location of the BH, and f_{turb} and f_{ang} are correction terms that account for the suppression of the accretion rate due to turbulence and vorticity, respectively (M. R. Krumholz, C. F. McKee & R. I. Klein 2006). These corrections are due to turbulence and vorticity on resolved scales, well outside a subgrid accretion disc. The modified BHL formula may work better with a realistic multiphase ISM that has non-zero angular momentum with respect to the BH, and perform better than the traditional formula in the context of galaxy formation (P. F. Hopkins & E. Quataert 2011; A. Hobbs et al. 2012; M. Gaspari, M. Ruszkowski & S. P. Oh 2013; D. Anglés-Alcázar et al. 2021).

The turbulence limiter is given by

$$f_{\text{turb}} = \left[\frac{\lambda^2 + \mathcal{M}^2}{(1 + \mathcal{M}^2)^4} \right]^{1/2}, \quad (4)$$

where $\lambda = 1.1$ and $\mathcal{M} = \sqrt{(v^2 + \sigma_{\text{g}}^2)}/c_{\text{s}}$ is the Mach number, with v the kernel-interpolated bulk velocity of gas relative to the BH, and σ_{g} the non-thermal velocity dispersion of the gas in the BH frame. For low Mach numbers, f_{turb} asymptotes to λ , while for high Mach numbers, it asymptotes to $1/\mathcal{M}^3$. In the latter case, the accretion rate scales as $\dot{M}_{\text{BHL}} \propto 1/(v^2 + \sigma_{\text{g}}^2)^{3/2}$, with no dependence on the sound speed. In the case that $v \gg \sigma_{\text{g}}$, the accretion rate corresponds to the classical limit $\dot{M}_{\text{BHL}} \propto 1/v^3$ for supersonic bulk flow (F. Hoyle & R. A. Lyttleton 1941).

The vorticity limiter is given by (M. R. Krumholz et al. 2006)

$$f_{\text{ang}} = \frac{0.34}{1 + \omega_{\star}^{0.9}}, \quad (5)$$

where $\omega_{\star} = \omega R_{\text{B}}/c_{\text{s}}$, $R_{\text{B}} = GM_{\text{BH}}/c_{\text{s}}^2$ is the Bondi radius and $\omega = |\nabla \times \mathbf{v}|$ is the kernel-weighted vorticity of the ambient gas. For a negligible amount of vorticity, f_{ang} asymptotes to 0.34, and the accretion rate is suppressed by that factor relative to BHL rate with no vorticity. This suppression is a result of accretion proceeding through a torus (rather than spherically), which is always expected to be present on some scales, even if the angular momentum is very small. For a large vorticity, the accretion rate is suppressed even more.

The accretion rate is limited to 100 times the Eddington rate³, which is given by

$$\dot{M}_{\text{Edd}} = \frac{4\pi G m_{\text{p}} M_{\text{BH}}}{\epsilon_{\text{rad}} \sigma_{\text{T}} c}, \quad (6)$$

where σ_{T} is the Thomson cross-section and ϵ_{rad} the accretion efficiency of a thin accretion disc. We assume $\epsilon_{\text{rad}} = 0.1$ in the thermal AGN feedback model and use the BH spin-dependent

³This limit is a compromise between allowing effectively uncapped BH accretion, and not allowing extremely high accretion rates that may lead to numerical issues.

formula in the hybrid AGN feedback model (I. D. Novikov & K. S. Thorne 1973, equation 23).

The (subgrid) BH mass evolves according to

$$\dot{M}_{\text{BH}} = (1 - \epsilon_{\text{rad}})\dot{M}_{\text{acc,d}}, \quad (7)$$

in the thermal AGN feedback model, since a fraction ϵ_{rad} of the accreting mass is radiated away as energy. In the hybrid AGN feedback model, we replace this equation by a similar equation including additional efficiency terms corresponding to the accretion disc winds and jets (see equation 15). Furthermore, in that version, the accretion rate on to the disc, $\dot{M}_{\text{acc,d}}$, is replaced with the accretion rate through the inner boundary of the accretion disc, into the BH event horizon, which is suppressed relative to $\dot{M}_{\text{acc,d}}$ by an accretion efficiency factor, $\epsilon_{\text{acc}} \leq 1$ (accounting for mass-loss due to accretion disc winds, as mass flows through an accretion disc from larger scales).

BHs grow in the simulation by ‘nibbling’ rather than swallowing entire gas particles (Y. M. Bahé et al. 2022). BHs are seeded on a gas particle that is converted into a BH particle, with the BH subgrid mass, M_{BH} , much smaller than the initial particle mass, $M_{\text{BH,part}}$. The subgrid BHs therefore initially ‘accrete’ from the particle that they are attached to since being seeded, as long as $M_{\text{BH}} < M_{\text{BH,part}}$. In this phase, the particle mass does not grow, while the subgrid mass does. Once the subgrid mass M_{BH} exceeds the particle mass $M_{\text{BH,part}}$, BHs ‘nibble’ on surrounding particles in order to conserve mass and energy. The nibbling is done in a kernel-weighted fashion, but only on gas particles whose mass is larger than half their initial baryonic mass. In the case that the BH grows from its particle mass reservoir rather than through nibbling on gas neighbours, in the thermal AGN feedback model we reduce the particle mass at the rate $\dot{M}_{\text{BH,part}} = -\epsilon_{\text{rad}}\dot{M}_{\text{acc,d}}$, which is consistent with equation (7), and which accounts for energy losses. In the hybrid AGN feedback model, we replace the radiative efficiency in this equation with the sum of all energy release efficiencies, and the disc-scale accretion rate with the accretion rate into the BH horizon.

2.4.3 Mergers

In order to merge, BHs must be located within three softening lengths of each other, the less massive BH needs to be within the SPH kernel of the more massive BH, and the relative velocity of the BHs must not be too large: $\Delta v < \sqrt{2G(M_{\text{BH,part,1}} + M_{\text{BH,part,2}})/r}$, where $M_{\text{BH,part},i}$ is the particle mass and r the separation between the two BHs. This condition ensures that the relative velocity does not exceed the escape velocity of the system. When merging the BHs, we also conserve linear momentum. During a merger, a fraction of the total mass-energy of the two BHs is lost to gravitational waves. For the thermal AGN feedback model, this fraction is equal to

$$\frac{E_{\text{GW}}}{(M_{\text{BH,1}} + M_{\text{BH,2}})c^2} = 0.057191\nu + 0.543763\nu^2, \quad (8)$$

which is accurate to order ν^2 , where $\nu = M_{\text{BH,1}}M_{\text{BH,2}}/(M_{\text{BH,1}} + M_{\text{BH,2}})^2$ is the symmetric mass ratio. This formula is based on numerical relativity simulations (E. Barausse, V. Morozova & L. Rezzolla 2012), is valid for non-spinning BHs, and accounts for the losses during the initial BH inspiral, the final plunge, the merger, and the ringdown. In the hybrid AGN feedback model, we use a BH spin-dependent formula (Section 3.7, equation 35), taken from the same work, of which the above is a special case. In the

non-spinning case, the mass-loss can reach ≈ 5 per cent for equal-mass mergers. In the maximally spinning, counter-aligned case, this can reach ≈ 12 per cent. We do not implement gravitational-wave recoils.

2.4.4 Repositioning

Instead of modelling dynamical friction acting on to BHs, whose effects we cannot reliably resolve in our simulations, we reposition them manually to the local minimum of the potential well, following the prescription detailed in Y. M. Bahé et al. (2022). At every time-step, the BH is instantaneously moved to the position of the neighbouring gas particle with the lowest gravitational potential that satisfies the following criteria. The particle must: (1) be within the SPH kernel of the BH, (2) be within three gravitational softening lengths of the BH, and (3) have a lower potential than that at the BH’s location (excluding the BH’s own contribution). If no such particle is found, the BH does not reposition. When repositioning, we do not alter the BH’s velocity.

2.4.5 Thermal AGN feedback

The thermal AGN feedback model follows the prescription of C. M. Booth & J. Schaye (2009). Instead of continuous energy injection at every time-step, energy from accretion is stored in a reservoir until it reaches a value large enough to heat a single gas neighbour by a large amount, to prevent numerical overcooling and drive powerful outflows (for similar reasons as for SN feedback, Section 2.3.4). The rate at which the energy reservoir is incremented at every time-step is the thermal feedback power

$$P_{\text{th}} = \epsilon_{\text{f}}L_{\text{bol}} = \epsilon_{\text{f}}\epsilon_{\text{rad}}\dot{M}_{\text{acc,d}}c^2, \quad (9)$$

where L_{bol} is the (emitted) bolometric luminosity of the AGN, $\epsilon_{\text{rad}} = 0.1$ is the assumed thin disc radiative efficiency, and ϵ_{f} the coupling efficiency, a parameter that determines what fraction of the emitted radiation couples to gas (either on accretion disc or larger scales, and either by directly heating it or driving a wind through radiation pressure that then shocks on subgrid scales). The coupling efficiency is calibrated on the BH mass–stellar mass relation (Section 4.5), with different values (between 0.02 and 0.1) adopted in the thermal and hybrid AGN feedback models, and at different resolutions. In the hybrid AGN feedback model, a version of the above equation is applied only at moderate Eddington ratios (corresponding to the thin disc state), while otherwise (in the thick and slim disc), the factor $\epsilon_{\text{f}}\epsilon_{\text{rad}}$ is replaced by ϵ_{wind} , which corresponds to an accretion disc wind that is shocked on subgrid scales and does energy-driven feedback on larger scales (Section 2.4.5). In the hybrid AGN feedback model, the accretion rate $\dot{M}_{\text{acc,d}}$ is also suppressed by the accretion efficiency factor, ϵ_{acc} .

The energy required for a single thermal AGN feedback event is

$$\Delta E_{\text{AGN}} = \frac{m_{\text{g}}k_{\text{B}}\Delta T_{\text{AGN}}}{(\gamma - 1)\mu m_{\text{p}}}, \quad (10)$$

where $\gamma = 5/3$ is the adiabatic index for monatomic gas, and $\mu = 0.6$, valid for ionized gas of primordial composition. The heating temperature for a single feedback event scales with BH mass:

$$\Delta T_{\text{AGN}} = 10^9 \text{ K} \left(\frac{M_{\text{BH}}}{10^8 M_{\odot}} \right), \quad (11)$$

with a minimum of $\Delta T_{\text{AGN},\text{min}} = 10^{6.5}$ K and a maximum of $\Delta T_{\text{AGN},\text{max}} = 10^{9.5}$ K at m7 resolution and $\Delta T_{\text{AGN},\text{max}} = 10^{10}$ K at m6 and m5 resolutions. While the feedback efficiencies associated with thermal feedback change in the hybrid AGN feedback model, we use the same heating temperatures as for the thermal model.

The variation of the heating temperature with BH mass, as opposed to the use of a large, constant value such as in EAGLE ($\Delta T_{\text{AGN}} = 10^9$ K, J. Schaye et al. 2015), has a few benefits, all of which are connected. First, it allows better sampling of AGN feedback in low-mass galaxies. In other words, a constant value biases the simulations towards better resolving AGN feedback in more massive systems. While it is typically assumed that that regime is the only one where AGN feedback is important, variable feedback temperatures (and jet kick velocities; Section 3.10) is what allows the hybrid AGN feedback model to feature significant AGN effects in galaxies with masses of order $M_* = 10^9 M_\odot$ or even lower, rather than only at $M_* \gtrsim 10^{10.5} M_\odot$ (as we show in Section 4.3.1). Secondly, using the specific formulation $\Delta T_{\text{AGN}} \propto M_{\text{BH}}$ ensures that, as BHs grow, they do the same number of heating events per unit of growth on a logarithmic scale. This again means that feedback from higher-mass BHs is not preferred over feedback from lower-mass ones, in terms of how well resolved it is. Finally, using variable heating temperatures allows one to implement feedback more gently in the early stage of BH growth.

In addition to gravitational time-steps, BHs use a time-step condition $\Delta t_{\text{AGN}} \leq \Delta E_{\text{AGN}}/P_{\text{th}}$, which ensures that BHs are woken up by the time they need to heat one gas neighbour. In the hybrid AGN feedback model, we apply an equivalent time-step limiter for jet feedback, and two additional time-step conditions related to BH spin evolution (Section 3.8). A further minimum of 100 yr is applied to BH time-steps (both in the thermal-only and hybrid AGN feedback simulations). This is very short, but it is sometimes reached in the m5 simulations, since extremely high densities are reached at that resolution.

2.5 Parameters and calibration

In Section 4, we discuss the strategy adopted for calibrating the parameter values of the hybrid AGN feedback model (Section 4). The strategy is different to the one used for calibrating the thermal AGN feedback simulations, which is discussed in detail in E. Chaikin et al. (2025). Briefly, the final calibration involved simulations with a comoving volume of $(50 \text{ cMpc})^3$ at m7 resolution, in a Latin hypercube configuration of parameter values. Gaussian process emulation was used to infer the best-fitting parameter values. These parameters (see below for the list) are present also in the hybrid AGN feedback model. The best-fitting parameter values found by E. Chaikin et al. (2025) for all three resolutions were used as initial guesses for the hybrid AGN feedback model. They are varied by hand, alongside new parameters specific to the hybrid AGN feedback model, and calibrated as discussed in Section 4.

The Gaussian process emulation employed by E. Chaikin et al. (2025) aimed to find the parameter values that simultaneously yield the best agreement with the $z = 0$ galaxy stellar mass function (GSMF) data from S. P. Driver et al. (2022) (GAMA survey year 4 release) and $z = 0$ galaxy size data from J. A. Hardwick et al. (2022). Equal weights were given to both datasets, and the fitting was done only over the range $10^9 < M_* < 10^{11.3} M_\odot$,

due to resolution constraints at the low-mass end and box-size limitations at the high-mass end.

Most of the calibration process involved varying four subgrid parameters: (1) $M_{\text{BH},\text{seed}}$, the BH seed mass, (2) f_{kin} , the fraction of CCSN energy injected in kinetic form, (3) $n_{\text{H},\text{pivot}}$, the pivot density for SN heating temperatures, and (4) $P_{\text{E},\text{pivot}}$, the pivot stellar birth pressure at which the SN energy fraction (f_{E}), which depends on stellar birth pressure, is halfway between its minimum and maximum value.

For the m6 and m5 resolutions, instead of employing emulation, slight adjustments to parameters were chosen by hand relative to their m7 values, to achieve similar agreement with the observational data. This process was done using L025 and L012p5 boxes for m6 and m5, respectively. In short, SN heating temperatures and pivot pressures increase, in order to suppress numerical overcooling, since higher densities are reached in higher-resolution simulations. For the same reason, BHs also grow more in higher-resolution simulations. To adjust for this effect, BH seed masses are reduced at higher resolution. For all three resolutions, the thermal AGN feedback coupling efficiency, ϵ_{t} , is adjusted independently of other parameters, to yield BH masses similar to $z = 0$ observations.

3 HYBRID AGN ACCRETION AND FEEDBACK MODEL

Here, we introduce our hybrid AGN model for BH accretion, spin evolution and AGN feedback. We provide a full and self-consistent description. In Section 3.1, we provide a brief description of the model, while other subsections present the model in detail. For readers interested only in the outline of the model, we suggest to read Section 3.1 and skip the rest of this section.

In Section 3.11, we compare our model to similar previously published ones. Different aspects of our model were already introduced and discussed in F. Huško et al. (2022, 2024, 2025). In Appendix A, we discuss how these different aspects are combined and what is different about our final model compared to these previously published, simpler versions.

3.1 Overview of the model

The BH accretion and feedback subgrid model that we present here has three accretion disc states, representing qualitatively different properties of accretion discs. These different states, namely the thick, thin and slim disc, occur at different Eddington ratios (low, moderate, and super-Eddington, respectively). These accretion disc states are illustrated schematically in Fig. 2. We discuss them in detail in Section 3.2. It is important to emphasize that there is no subgrid accretion disc that we evolve (i.e. no subgrid mass or angular momentum reservoir); we instead assume that accretion occurs from large scales on to the BH instantaneously. However, assumptions about accretion disc properties can significantly affect mass accretion, BH spin evolution and feedback even in this simplified scenario.

The model includes two AGN feedback modes: thermal isotropic and kinetic jets. The thermal isotropic mode represents the effects of winds that are adiabatically shocked on subgrid scales, as well as effects of radiation (which may couple to gas on larger than accretion disc scales). Both feedback modes are active in all three accretion disc states, resulting in a total of six effective feedback channels. In practice, we find that only three of the six channels are actually relevant for the galaxy and BH population

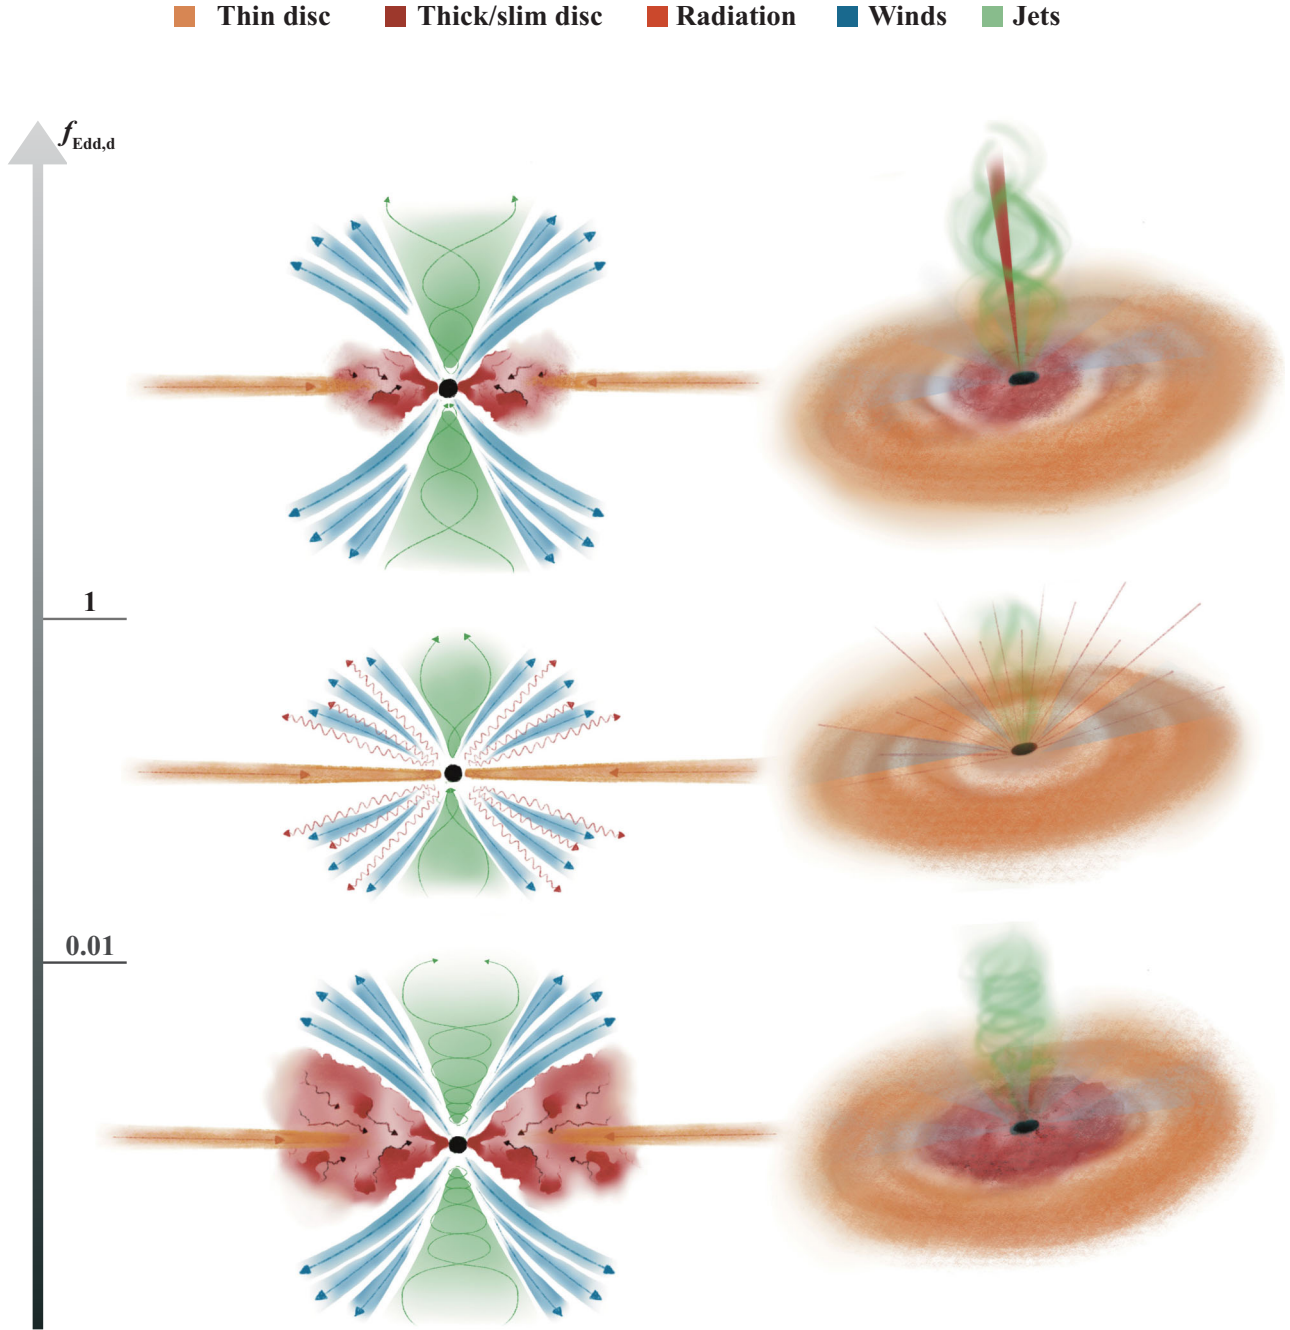


Figure 2. The three different accretion disc states assumed in our model, corresponding to different Eddington ratio regimes, from bottom to top displaying the thick, thin and slim discs, with the thin disc component present at large radii. The accretion disc state changes depending on the disc-scale Eddington ratio, $f_{\text{Edd},d}$, defined in Section 3.2, at two critical values of 0.01 and 1. The left and right illustrations show edge-on and inclined views of the accretion disc, respectively. Radiation and winds are assumed to originate on accretion disc scales, and the effects of both are implemented using thermal isotropic feedback in the simulations. Jets, which consist purely of Poynting flux on these accretion disc scales, are implemented as bipolar velocity kicks.

as a whole (1) jets at low Eddington ratios (at low redshifts), (2) jets at moderate Eddington ratios (at high redshifts), and (3) thermal feedback at moderate Eddington ratios (at all redshifts).

In our model, BH growth can be significantly impeded by the direct effects of feedback, for two separate reasons. First, the feedback efficiencies can be very high and result in net mass loss by BHs. Secondly, our model assumes that subgrid accretion

disc winds strip away most of the accreting material at low and super-Eddington accretion rates, which we represent using an accretion efficiency parameter that quantifies this mass-loss. This parameter relates the accretion rate on to a subgrid accretion disc to the accretion rate through the BH event horizon. The BH mass evolution and various accretion rates related to the above are discussed in Section 3.3.

In order to realistically employ feedback (especially concerning jets), we model the evolution of the angular momentum of a BH, \mathbf{J}_{BH} (see Sections 3.6 and 3.7). The dimensionless BH spin parameter (hereafter the BH spin) has a magnitude defined as

$$a = \frac{J_{\text{BH}} c}{M_{\text{BH}}^2 G}, \quad (12)$$

where $J_{\text{BH}} = |\mathbf{J}_{\text{BH}}|$ is the angular momentum magnitude, and a sign which indicates whether accretion is prograde or retrograde (either an aligned or counter-aligned inner accretion disc). The BH spin direction is used as the direction for jet launching. Since the BH spin vector is effectively a time-integrated quantity, it is typically much more stable in time than the instantaneous direction of the angular momentum of gas near the BH. This facilitates stable jet launching and results in large-scale jet activity, as shown in Fig. 1. In addition to being used for jet directions, the BH spin also affects the strength of feedback (i.e. the efficiencies).

Some of the feedback efficiencies in our model (the jet efficiency especially) depend on the magnetization state of the accretion disc. We assume the largest possible (steady-state) magnetization. In other words, we adopt the magnetically arrested disc (MAD) assumption (R. Narayan et al. 2003), as discussed in Section 3.4. We stress that the feedback efficiencies in our model are largely not free parameters; they instead depend on BH spin and Eddington ratio, as discussed in Section 3.5. The exception to this is the coupling efficiency of radiation from the thin accretion disc, but this same parameter exists in the purely thermal AGN feedback simulations. On the other hand, our model includes two new important parameters that are related to feedback, which set the values of accretion efficiencies in the thick disc state and jet velocities.

Due to the inclusion of jet feedback and BH spin, additional time-step conditions need to be adopted (see Section 3.8). Our jet launching scheme is discussed in Section 3.9, while the implementation of AGN feedback more generally is discussed in Section 3.10.

3.2 Accretion disc states

We assume different accretion disc states depending on the Eddington ratio defined using the accretion rate on to the subgrid accretion disc (the disc-scale accretion rate), before mass loss due to winds. This Eddington ratio is defined as

$$f_{\text{Edd,d}} = \frac{\dot{M}_{\text{acc,d}}}{\dot{M}_{\text{Edd}}}, \quad (13)$$

where $\dot{M}_{\text{acc,d}}$ is given by equation (2). For most calculations within our model, we use the Eddington ratio⁴ defined as

$$f_{\text{Edd,H}} = \frac{\dot{M}_{\text{acc,H}}}{\dot{M}_{\text{Edd}}}, \quad (14)$$

which is the same as equation (13), but with the disc-scale accretion rate, $\dot{M}_{\text{acc,d}}$, replaced by the accretion rate into the BH (event horizon), $\dot{M}_{\text{acc,H}}$ (see Section 3.3 for subtleties associated

with this difference). This second Eddington ratio includes the effects of mass loss from the accretion disc due to winds and corresponds more closely to observed Eddington ratios. When deciding the nature of an accretion disc, one could use either of the two Eddington ratios. We discuss the merits of both options in Appendix B1, and based on the discussion therein, we choose to use the disc-scale Eddington ratio (equation 13) to decide the accretion disc state.

The three accretion disc types we include in the model are the following:

(i) **Thick disc** for $f_{\text{Edd,d}} < 0.01$. This accretion flow is also known as the ADAF (advection-dominated accretion flow), ADIOS (advection-dominated inflow–outflow solution if winds are included), hot accretion flow, RIAF (radiatively inefficient accretion flow), the hard state (in terms of X-ray spectra) and the low state (in terms of accretion rate). The disc is geometrically thick (with an aspect ratio $H/R \approx 0.5$, H being the scale height of the disc at a cylindrical radius R) and optically thin. The gas in this disc is very hot and relatively diffuse, with low radiative efficiencies and strong advection that traps most of the thermal energy. Gas orbits are not fully circular, instead having a significant radial component. The gas flow is continuous all the way down to the event horizon, with no abrupt change in properties at the radius of the innermost stable circular orbit (R_{ISCO}). The poloidal magnetic flux at the event horizon of the BH is very large, leading to strong jets if the BH is spinning. Winds driven by thermal pressure and MHD effects are also present, taking away most of the accreting mass. We take the solution for this type of disc from R. Narayan & I. Yi (1995), and we also use results from numerical GRMHD simulations (A. Tchekhovskoy et al. 2010; R. Narayan et al. 2022). BH spin evolution and AGN jet feedback in this accretion disc state was studied in detail in F. Huško et al. (2022).

(ii) **Thin disc** for $0.01 < f_{\text{Edd,d}} < 1$. This accretion disc state is also known as the radiatively efficient, standard or N. I. Shakura & R. A. Sunyaev (1973) disc (SD or SSD), as well as the soft or high state. The disc is geometrically thin (usually $H/R \in [0.001, 0.01]$) and optically thick, with a high radiative efficiency of ~ 10 per cent. Gas orbits are almost fully circular and extend down to R_{ISCO} . Within this inner radius, they quickly become unstable. Jets are likely present, but their strength depends sensitively on the accretion rate and magnetic state of the disc. This state is effectively the one that is assumed at all accretion rates in the thermal AGN feedback model. We take the solution for this disc from N. I. Shakura & R. A. Sunyaev (1973). We studied feedback in this accretion mode in detail in F. Huško et al. (2024), alongside the thick disc that was used at low Eddington ratios.

(iii) **Slim disc** for $f_{\text{Edd,d}} > 1$. This accretion flow appears in the super-Eddington regime of accretion. The disc is geometrically thick (H/R of order 0.1–0.5), but less so than the thick disc. It shares some features with both the thick and thin discs: it is both advection-dominated and radiatively efficient. It launches strong jets and is very luminous (the efficiencies for jets and radiation are lower than in the thick and thin disc, respectively, but the accretion rate is higher, resulting in potentially higher jet powers and luminosities). Radiation pressure and MHD effects drive a wind that may take away most of the accreting mass. We take the solution for this type of disc from J.-M. Wang & Y.-Y. Zhou (1999), and also use numerical general-relativistic radiation magnetohydrodynamical (GRRMHD) simulation results from A. Ricarte et al. (2023). We introduced this disc state and the feedback from it in our model in F. Huško et al. (2025).

⁴To define the Eddington ratios in the hybrid AGN feedback model, following recent conventions, we use the BH spin-dependent Eddington accretion rate \dot{M}_{Edd} (equation 6), in which the constant radiative efficiency $\epsilon_{\text{rad}} = 0.1$ is replaced with a BH spin-dependent one that is appropriate for a thin, relativistic disc (equation 23). This can lead to differences in Eddington ratios of up to a factor of 3.




Figure 3. The fraction of BHs occupying different accretion disc states in the L200m7h simulation: *left*: thick disc, *middle*: thin disc, *right*: slim disc, multiplied by a factor of 10 for visual comparison purposes. Different colours indicate different redshifts. The shaded regions show uncertainties computed assuming binomial proportion confidence intervals.

The choices for the critical Eddington ratios ($f_{\text{Edd},d} = 0.01$ and $f_{\text{Edd},d} = 1$) that separate the above states are discussed in Appendix B2. The modelling of accretion discs is highly dependent on assumptions about the kinematic viscosity ν in the discs, whose properties are largely unknown. All of the above accretion disc models are of so-called α -discs, where α is a numerical parameter that is related to the kinematic viscosity ν through $\nu = \alpha c_s H$. Here, c_s and H are the sound speed and height of the disc at a given radius, respectively. In the α -disc models, α does not vary with radius, but is instead given by a constant value. We set a value of $\alpha = 0.1$, as motivated by the discussion in Appendix B2.

In order to illustrate the importance of the three accretion disc states, in Fig. 3 we show the fraction of all BHs that are in a given accretion disc state, as a function of BH mass and for different redshifts. Here, we use results from one of our own simulations (specifically, L200m7h) that use the hybrid AGN feedback model in combination with the COLIBRE galaxy formation model. Most BHs are in the thick disc state at $z = 0$, with the remainder (5-15 per cent) in the thin disc state. The fraction of low-redshift BHs in the thin disc state peaks at $M_{\text{BH}} \approx 10^7 M_{\odot}$, corresponding to host galaxies with stellar masses $M_* = 10^{10}-10^{10.5} M_{\odot}$, a mass regime where BHs grow most rapidly. The BHs in the thin disc state are observable as classical optical and X-ray AGN. The $z = 0$ thin disc occupation fraction that we find here matches AGN fractions found observationally in the local Universe (e.g. Z. Zhang et al. 2021; J. L. Steffen et al. 2023), where AGN are identified using the BPT diagram. In these observations, the AGN fraction is non-negligible for galaxies with stellar masses between 10^9 and $10^{11} M_{\odot}$, with a peak value of 15-20 per cent at $M_* \approx 10^{10.5} M_{\odot}$. Although Fig. 3 has the BH mass on the x -axis, this overall shape (see black line in the middle panel) and the peak value of 15 per cent found in the simulations, reached at $M_{\text{BH}} \approx 10^7 M_{\odot}$, are in good agreement with observations given that galaxies with $M_* \approx 10^{10.5} M_{\odot}$ host BHs with median masses $M_{\text{BH}} \approx 10^7 M_{\odot}$ (Section 4).

As redshift increases, the fraction of BHs occupying the thin disc state increases, at all BH masses, but most prominently around that same mass scale, $M_{\text{BH}} \approx 10^7 M_{\odot}$. By $z = 1, 2, 4$, only $\approx 25, \approx 5$, and ≈ 2 per cent of such BHs are in the thick disc state, with most of the remainder growing efficiently in the thin

disc state. Above this mass scale, the fraction of BHs in the thin disc state drops, since more massive BHs are more likely to have quenched their host galaxy and stifled their own growth. The fractions of BHs occupying different accretion states reflects the evolution of the median Eddington ratio at fixed BH mass, which increases with redshift (Huško et al. in preparation).

The slim disc state is very rare at all redshifts, and exceedingly so at lower redshifts and/or high BH masses. Low-mass BHs ($M_{\text{BH}} \in [10^6, 10^8] M_{\odot}$) are the most likely to be in the slim disc state. The BH mass range where this is likely at a given redshift is even narrower, and moves from higher to lower-mass BHs with increasing redshift. At $z = 3, z = 4, z = 6$, and $z = 8$, the peak in the fraction of BHs occupying the slim disc state is at $M_{\text{BH}} \approx 10^8, \approx 10^8, \approx 10^7$, and $\approx 10^6 M_{\odot}$, respectively, with the corresponding peak occupation fractions equal to $\approx 1, \approx 2.5, \approx 6$, and ≈ 6 per cent.

3.3 Black hole mass evolution

The BH mass evolves according to

$$\dot{M}_{\text{BH}} = (1 - \epsilon_{\text{tot}}) \dot{M}_{\text{acc,H}} = (1 - \epsilon_{\text{rad}} - \epsilon_{\text{jet}} - \epsilon_{\text{wind}}) \epsilon_{\text{acc}} \dot{M}_{\text{acc,d}}, \quad (15)$$

where ϵ_{tot} is the sum of all energy release efficiencies (radiative efficiency ϵ_{rad} , jet efficiency ϵ_{jet} , and wind efficiency ϵ_{wind}), and the individual efficiency terms are defined using the relation $P_i = \epsilon_i \dot{M}_{\text{acc,H}} c^2$, where P_i is the feedback power into a given channel i . Equation (15) is analogous to that used for the thermal AGN feedback model (equation 7), but it has additional feedback terms alongside radiation. The radiation and wind terms are both associated with the thermal isotropic feedback channel used in our model, with the radiation term used in the thin disc state (and $\epsilon_{\text{wind}} = 0$), and the wind term in the thick and slim disc states (with $\epsilon_{\text{rad}} = 0$ for the purpose of equation 15). The sum of the three efficiencies may exceed unity, in which case the BH can reduce its mass as it accretes (due to losing more mass-energy associated with its rotational energy, than it gains through accretion).

In addition to these feedback terms, equation (15) also features the accretion efficiency, $\epsilon_{\text{acc}} < 1$, which relates the disc-scale accretion rate, $\dot{M}_{\text{acc,d}}$, and the accretion rate into the event




Figure 4. The assumed accretion efficiencies as a function of the disc-scale Eddington ratio. The accretion efficiency connects the accretion rate on to a subgrid accretion disc (equation 15) with the accretion rate into the BH itself, and it represents the effects of disc winds that take away most of the mass. In the thin disc we assume a 100 per cent accretion efficiency, while in the slim disc we assume 1 per cent. In the thick disc, we combine results of GRMHD simulations and analytical calculations (equations 16 and 17). For the free parameter $r_{tr,0}$, we take the value 10^4 , corresponding to the value we found through calibration on the bolometric AGN luminosity function (see Section 4). The accretion efficiency in the thick disc is limited by the ratio of the Bondi radius and the gravitational radius of the BH, which depends only on the gas sound speed. Using different line styles, we show the accretion efficiency for several values of the sound speed, ranging from values appropriate for the hot ICM in galaxy clusters to those appropriate for cold, dense gas in the ISM.

horizon, $\dot{M}_{acc,H}$, through $\dot{M}_{acc,H} = \epsilon_{acc}\dot{M}_{acc,d}$. This relation encapsulates mass-loss due to accretion disc winds that we assume to be present in the thick and slim discs. In the next subsection, we discuss our assumptions for the accretion efficiency in different accretion disc states.

3.3.1 Accretion efficiency

The accretion efficiency, ϵ_{acc} , represents the fraction of mass accreting on to a subgrid accretion disc that actually reaches the event horizon of a BH. We assume that it is less than unity only in the thick and slim disc, where we assume strong winds to be operating. We show our assumed accretion efficiencies in Fig. 4 and in the first row of Table 1.

We assume the accretion efficiency to be given by

$$\epsilon_{acc} = \begin{cases} \left(\frac{10}{r_{th}}\right)^s, & f_{Edd,d} < 0.01 \text{ (thick disc)} \\ 1, & 0.01 < f_{Edd,d} < 1 \text{ (thin disc)} \\ 0.01, & f_{Edd,d} > 1 \text{ (slim disc)} \end{cases} \quad (16)$$

For the thin disc, we assume a 100 per cent accretion efficiency, while we assume 1 per cent for the slim disc, as motivated below. For the thick accretion disc state, the accretion efficiency depends on the (dimensionless) radius of the thick component of the accretion disc $r_{th} = R_{th}/R_G$, where R_{th} is the physical radius of the thick disc, and $R_G = M_{BH}G/c^2$ is the gravitational radius of the BH. At radii larger than R_{th} , the disc may be present, but it is in the

Table 1. The accretion efficiency, dimensionless magnetic flux, and feedback efficiencies assumed for the three accretion disc states. The values (which depend on the BH spin, a) of the (dimensionless) event horizon angular velocities Ω_H and radius of the innermost stable circular orbit, ISCO, are given shortly after equation (24) in the text and in Appendix C, respectively. The thick disc size, r_{tr} , is given by equation (17). ϵ_f is the coupling efficiency of thin disc radiative feedback, while $r_{tr,0}$ is the dimensionless size of the thick accretion disc at a critical Eddington ratio $f_{Edd,d} = 0.01$. Both ϵ_f and $r_{tr,0}$ are free parameters.

	Thick disc ($f_{Edd,d} < 0.01$)	Thin disc ($0.01 < f_{Edd,d} < 1$)	Slim disc ($f_{Edd,d} > 1$)
ϵ_{acc} (accretion efficiency)	$(10/r_{th})^{0.5}$	1	0.01
ϕ (magnetic flux)	$r_{th} = \min\left[r_{tr,0}\left(\frac{0.01}{f_{Edd,d}}\right)^2, r_{tr}\right]$		
ϵ_{th} (thermal efficiency)	$\phi_0(a) = -20.2a^3 + 34a + 52.6$	$\epsilon_f \left(1 - \sqrt{1 - \frac{2}{3r_{ISCO}}}\right)$	$0.065 \left[1 + \left(\frac{\phi}{50}\right)^2\right] \max(0, 1 + \Omega_H - 8\Omega_H^2)$
ϵ_{jet} (jet efficiency)	$0.0005 \left[1 + \left(\frac{\Omega_H}{0.2}\right)\right]$	$(\kappa/4\pi)\phi^2\Omega_H^2(1 + 1.38\Omega_H^2 - 9.25\Omega_H^4)$	

thin state and is thus assumed in our model to not be contributing to the launching of a wind. The scaling of the accretion efficiency with the size of the thick disc, $\epsilon_{\text{acc}} = (10/r_{\text{th}})^s$, is well motivated by theory (F. Yuan & R. Narayan 2014), where larger discs have smaller accretion efficiencies, implying that the exponent $s > 0$. Its value is also limited to $s < 1$ on theoretical grounds (R. D. Blandford & M. C. Begelman 1999), but it cannot be predicted analytically. Observations suggest $0.3 < s < 1$ (R. S. Nemmen & A. Tchekhovskoy 2015), while recent MHD and GRMHD simulations have converged on $s \approx 0.5$ (M. Guo et al. 2023, H. Cho et al. 2024), so we assume that value. The normalization constant equal to 10 (dimensionless radii) is present in the scaling since accretion disc winds are not thought to be launched from within ≈ 10 gravitational radii (F. Yuan & R. Narayan 2014).

For the size of the thick disc, we take the scaling from R. Narayan & I. Yi (1995), which assumes that the transition from the thin disc at large radii to the thick disc at small radii occurs at a radius where half of the energy is radiated away, and half advected inwards⁵ We also assume that the disc can be no larger than the Bondi radius. For the dimensionless size of a thick disc we thus assume

$$r_{\text{th}} = \min \left[r_{\text{tr},0} \left(\frac{0.01}{f_{\text{Edd},d}} \right)^2, r_{\text{B}} \right]. \quad (17)$$

where $r_{\text{tr},0} = R_{\text{tr},0}/R_{\text{G}}$ is the dimensionless size of the thick accretion disc that is accreting at the critical Eddington ratio of 0.01, and $r_{\text{B}} = R_{\text{B}}/R_{\text{G}} = (c/c_s)^2$ is the dimensionless Bondi radius. The value of $r_{\text{tr},0}$, which serves as the normalization of the dependence of r_{th} on $f_{\text{Edd},d}$, is in principle predicted by theory (R. Narayan & I. Yi 1995). However, it is highly dependent on the accretion disc viscosity α , more steeply than $r_{\text{tr},0} \propto \alpha^4$. The value of α could be in the range 0.1–0.4, so the value of $r_{\text{tr},0}$ itself is highly uncertain. We hence leave it as a free parameter, and calibrate its value (see Section 4.4).

Given the assumptions discussed above, the accretion efficiency for the thick disc, i.e. at $f_{\text{Edd},d} < 0.01$, can be written as

$$\epsilon_{\text{acc,th}} \approx \begin{cases} 3.16 \times 10^{-2} \left(\frac{r_{\text{tr},0}}{10^4} \right)^{-0.5} \left(\frac{f_{\text{Edd},d}}{0.01} \right), & r_{\text{B}} \geq r_{\text{tr},0} \left(\frac{f_{\text{Edd},d}}{0.01} \right)^{-2} \\ 10^{-3} \left(\frac{c_s}{100 \text{ km s}^{-1}} \right), & r_{\text{B}} < r_{\text{tr},0} \left(\frac{f_{\text{Edd},d}}{0.01} \right)^{-2}, \end{cases} \quad (18)$$

for our fiducial values. The first of these cases is valid if the transition radius (between an outer thin and an inner thick disc) is smaller than the Bondi radius, i.e. if an outer thin disc is present. In this case, the accretion efficiency is linear with $f_{\text{Edd},d}$ and reaches values no larger than 3.16×10^{-2} . In the second case, the entire disc (out to the Bondi radius) is thick. The dimensionless Bondi radius depends only on the sound speed of the gas, so the accretion efficiency scales as $\epsilon_{\text{acc,th}} \propto c_s$. For accretion from hot gas (e.g. $c_s = 1000 \text{ km s}^{-1}$), the accretion efficiency reaches values of 10^{-2} , but it will generally be lower.

For the slim disc, a similar picture as above likely holds. However, slim discs become relatively more advective and more efficient at launching winds at higher accretion rates (A. Ricarte et al. 2023), rather than at lower ones like for the thick disc. This

⁵Other calculations, such as from the disc evaporation model (e.g. R. E. Taam et al. 2012 and references therein), give quantitatively different expressions for the transition radius. However, the qualitative behaviour of a larger thick disc component at smaller Eddington ratios also appears in such models.

implies that the accretion efficiency for the slim disc should have the form $\epsilon_{\text{acc}} \propto 1/f_{\text{Edd},d}^a$, with $a > 0$. In other words, it should decline with increasing Eddington ratio, which is opposite behaviour to that in the thick disc state. Since no simulations of slim discs have, to our knowledge, yet studied wind mass loading as a function of the Eddington ratio, we instead assume a constant slim disc accretion efficiency parameter, $\epsilon_{\text{acc,slim}}$. We do not calibrate it on any particular quantity or relation, but instead assume a value of $\epsilon_{\text{acc,slim}} = 0.01$, which is similar to the values of the accretion efficiency in the thick disc (although lower than in some GRRMHD simulations, see e.g. Y.-F. Jiang, J. M. Stone & S. W. Davis 2019).

3.4 Accretion disc magnetisation

The magnetic flux threading the BH event horizon affects the efficiency of thermal feedback in the thick and slim disc states, as well as jet efficiencies in all accretion states. In general, accretion discs may occupy two qualitatively distinct states in terms of how strong their magnetic fields are: the ‘standard and normal evolution’ (SANE) state (R. Narayan et al. 2012), or the magnetically arrested disc (MAD) state (R. Narayan et al. 2003). In the SANE state, magnetic fields are unconstrained and dynamically unimportant, while in the MAD state, they are dynamically important and have reached a saturated value, that depends exclusively on accretion rate and BH spin. This state is reached as a result of a strong magnetic dynamo effect, in combination with advection of the field from larger radii.

We assume that all of our accretion discs are in the MAD state, and summarize our choice for the dimensionless magnetic flux, $\phi = \Phi_{\text{H}}/\sqrt{\dot{M}_{\text{acc,H}}R_{\text{G}}^2c}$, Φ_{H} being the magnetic flux threading the event horizon, in the second row of Table 1. The dimensionless magnetic flux in the MAD state is given by

$$\phi(f_{\text{Edd,H}}, a) = \phi_0(a) \times \begin{cases} 1, & f_{\text{Edd},d} < 0.01 \\ \frac{(f_{\text{Edd,H}}/1.88)^{1.29}}{1+(f_{\text{Edd,H}}/1.88)^{1.29}}, & f_{\text{Edd},d} > 0.01 \end{cases} \quad (19)$$

Here, $\phi_0(a)$ is the dimensionless magnetic flux of a thick or sufficiently super-Eddington slim disc, which depends only on BH spin and is given by (R. Narayan et al. 2022)

$$\phi_0(a) = -20.2a^3 - 14.9a^2 + 34a + 52.6. \quad (20)$$

For thin discs and mildly super-Eddington slim discs, the value of ϕ is lower than $\phi_0(a)$. The functional form given by equation (19) for those discs gives a strong dependence on the Eddington ratio, $\phi \propto f_{\text{Edd,H}}^{1.29}$ if $f_{\text{Edd,H}} \ll 1$, and saturation to $\phi_0(a)$ if $f_{\text{Edd,H}} \gg 1$.

The MAD assumption is supported by observations for thick discs (Event Horizon Telescope Collaboration et al. 2021; K. Akiyama et al. 2022). It is also strongly supported on theoretical grounds for advection-dominated discs (F. Yuan & R. Narayan 2014). We thus expect it to be valid at least in the thick disc state and for very super-Eddington slim discs. However, for thin discs and accretion near the Eddington rate, our assumed magnetic flux may be too high.

3.5 Feedback efficiencies

We include two modes of AGN feedback: thermal isotropic and kinetic jets, both of which are used at the same time, and in each accretion disc state. The first of these modes encapsulates the effects of radiation and accretion disc winds, while the second

one represents the effects of jets driven by the R. D. Blandford & R. L. Znajek (1977) mechanism. The feedback powers are given by

$$P_{\text{th/jet}} = \epsilon_{\text{th/jet}} \dot{M}_{\text{acc,H}} c^2 = \epsilon_{\text{th/jet}} \epsilon_{\text{acc}} \dot{M}_{\text{acc,d}} c^2, \quad (21)$$

where $\epsilon_{\text{th/jet}}$ is either the thermal or jet efficiency. The feedback efficiencies are shown in Fig. 5, and are introduced in the rest of this section. In the left panel we show only the feedback efficiencies, while in the right panel, we also multiply them by the accretion efficiencies, to make the net energy release efficiencies more comparable between different accretion disc states. The feedback efficiencies that we assume are also summarized in the 3rd and 4th rows of Table 1, for thermal and jet feedback, respectively.

3.5.1 Thermal feedback

We assume that AGN emit radiation at the rate

$$L_{\text{AGN}} = \epsilon_{\text{rad}} \dot{M}_{\text{acc,H}} c^2, \quad (22)$$

where ϵ_{rad} is the radiative efficiency and $\dot{M}_{\text{acc,H}}$ is the accretion rate into the event horizon. For the thin disc, we assume that a fraction ϵ_f , the coupling efficiency, of that radiation couples to gas (either on accretion disc or larger scales), causing feedback,⁶ and the observable AGN luminosity is thus equal to $(1 - \epsilon_f)L_{\text{AGN}}$. The coupling efficiency, ϵ_f , is a free parameter. Its value is calibrated by ensuring our simulated BH masses match observations at $z = 0$ (Section 4).

For the radiative efficiency, for the thin disc, we assume the expression for a relativistic thin accretion disc (I. D. Novikov & K. S. Thorne 1973):

$$\epsilon_{\text{rad,thin}}(a) = 1 - e_{\text{ISCO}}(a) = 1 - \sqrt{1 - \frac{2}{3r_{\text{ISCO}}(a)}}. \quad (23)$$

Here, e_{ISCO} is the specific binding energy (in units of c^2) at the innermost stable circular orbit (ISCO). $r_{\text{ISCO}} = R_{\text{ISCO}}/R_G$, is the dimensionless radius of the ISCO. We give the radius of the ISCO as a function of BH spin in Appendix C. The above formulation gives a radiative efficiency for the thin disc that varies between $\epsilon_{\text{rad}} \approx 4$ per cent at $a = -0.998$ and $\epsilon_{\text{rad}} \approx 30$ per cent at $a = 0.998$, which is to be compared with the constant value of 10 per cent in the thermal AGN feedback model.

In the thick and slim discs, we assume different, but very small radiative efficiencies (see Appendix D), since both of those disc types are radiatively inefficient. We also assume that $\epsilon_f = 0$ in those discs, since the slim disc winds may be partly driven by radiation on the accretion disc scales. The radiative efficiencies for those discs are thus used exclusively for calculating AGN luminosities.

For the thick and slim discs, we instead assume the presence of accretion disc winds that are shocked on subgrid scales (T. Costa, D. Sijacki & M. G. Haehnelt 2014), so we implement the effects of these winds using the thermal feedback channel (in other words, we assume $\epsilon_{\text{th}} = \epsilon_{\text{wind}}$ for these disc states). The thermal feedback

efficiencies are thus given by

$$\epsilon_{\text{th}}(f_{\text{Edd,H}}, a) = \begin{cases} 0.005 \left[1 + 0.3 \left(\frac{\phi}{50} \right) \left(\frac{\Omega_{\text{H}}}{0.2} \right) \right], & f_{\text{Edd,d}} < 0.01 \\ \epsilon_f \left(1 - \sqrt{1 - \frac{2}{3r_{\text{ISCO}}}} \right), & 0.01 < f_{\text{Edd,d}} < 1 \\ 0.065 \left[1 + \left(\frac{\phi}{50} \right)^2 \right] M(\Omega_{\text{H}}), & f_{\text{Edd,d}} > 1 \end{cases} \quad (24)$$

where $M(\Omega_{\text{H}}) \equiv \max(0, 1 + \Omega_{\text{H}} - 8\Omega_{\text{H}}^2)$. The second expression corresponds to $\epsilon_f \epsilon_{\text{rad,thin}}$ for the thin disc. The first and third expressions correspond to wind efficiencies for the thick and slim discs obtained from A. Sadowski et al. (2013) and A. Ricarte et al. (2023),⁷ respectively. Here, $\Omega_{\text{H}}(a)$ is the dimensionless angular velocity of the BH event horizon, given by $\Omega_{\text{H}}(a) = a/2r_{\text{H}}(a)$, where $r_{\text{H}}(a) = 1 + \sqrt{1 - a^2}$ is the dimensionless radius of the event horizon.

The thermal feedback efficiency for the thin disc, due to radiation (assuming $\epsilon_f = 0.025$, in the mid-range of values we find by calibration for the hybrid AGN feedback model, see Section 4), is shown in Fig. 5 using dashed horizontal lines at moderate Eddington ratios. It varies from ≈ 0.2 per cent at low BH spin ($a = 0.2$) to ≈ 1 per cent at maximal BH spin ($a = 0.998$). The wind efficiency for the thick disc is shown with horizontal lines at low Eddington ratios in Fig. 5. It is generally similar to or larger than the thermal efficiency in the thin disc. However, one needs to keep in mind that these winds are applied at lower accretion rates than in the thin disc, and that we apply an accretion efficiency in the thick disc (resulting in much lower net accretion rates than in the thin disc, see the right panel of Fig. 5). As a result, thick disc winds have a very small effect in practice, at least at the level of calibration of galaxy formation simulations. This is also related to the fact that jets are the dominant feedback mechanism in this regime, which is also visible in Fig. 5 (we introduce the jet efficiencies in the next subsection). The wind efficiency in the slim disc is shown using horizontal dashed lines at super-Eddington accretion rates in Fig. 5. The wind efficiencies have values of order 5-20 per cent, and are thus generally the largest among the thermal efficiencies used in our model. However, as in the thick disc, the winds from the slim disc are generally subdominant to its jets (unless the BH spin is quite low, $|a| < 0.4$).

3.5.2 Jet efficiencies

We assume that all accretion discs launch jets on account of energy (and angular momentum) extraction through the R. D. Blandford & R. L. Znajek (1977) (BZ) process. There have been many studies on jet launching from MAD thick discs, and recently it has been shown that the results from these studies may also be applied to slim and thin discs, once differences in the dimensionless magnetic flux, ϕ , between these disc types are accounted for (A. Ricarte et al. 2023). We assume the jet efficiency formula from A. Tchekhovskoy et al. (2010):

$$\epsilon_{\text{jet}} = \frac{\kappa}{4\pi} \phi^2 \Omega_{\text{H}}^2 (1 + 1.38\Omega_{\text{H}}^2 - 9.2\Omega_{\text{H}}^4), \quad (25)$$

⁶Our model is agnostic as to the exact mechanism by which radiation causes feedback. Radiation may launch a wind on accretion disc scales, or on larger scales. It may also directly heat gas on large scales. In either case, we assume that the thermal isotropic feedback employed in the simulations is representative the effects of whichever of these effects that is dominant.

⁷The thick disc wind efficiency is taken directly from A. Sadowski et al. (2013), while the slim disc wind efficiency is a fitting function to the results of A. Ricarte et al. (2023) that we derived and presented in F. Huško et al. (2025, see appendix A therein).

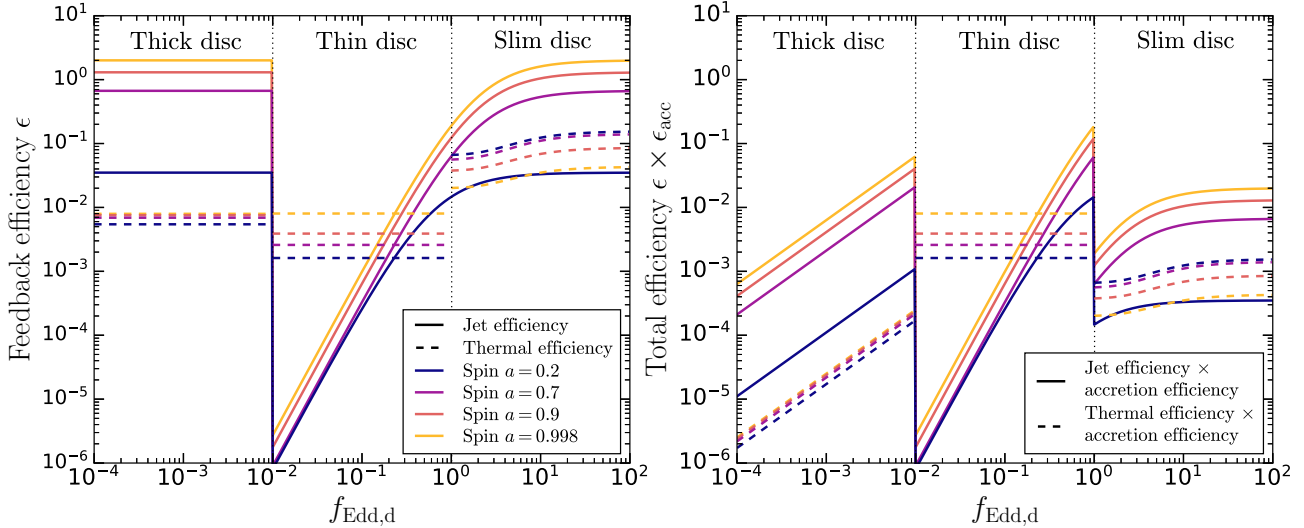


Figure 5. The dependence of feedback efficiencies used in the hybrid AGN feedback model on the accretion disc-scale Eddington ratio $f_{\text{Edd},d}$ for the three accretion and feedback states. Solid lines show the jet efficiencies (see Section 3.5.2), and dashed the thermal efficiencies (see Section 3.5.1), assuming a coupling efficiency $\epsilon_f = 0.025$. We show these for several values of BH spin using different colours (see legend); note that the jet efficiency is $\epsilon_{\text{jet}} = 0$ for $a = 0$. In the left panel we show the pure feedback efficiencies, while in the right panel, we further multiply these by accretion efficiencies (see Section 3.3.1) to emphasize the net feedback efficiencies and to facilitate comparison between the different modes. For this figure, we used $r_{\text{tr},0} = 10^4$ as found through calibration (Section 4.5).

where κ is a numerical factor that depends on the initial geometry of the magnetic field (e.g. 0.054 for split-monopole versus 0.044 for parabolic, we assume $\kappa = 0.05$). For a non-spinning BH ($a = 0$), the jet efficiency given by the above formula is $\epsilon_{\text{jet}} = 0$. All three of our accretion disc states use equation (25). They differ only in their values of ϕ , which introduces large differences in the jet efficiencies. For strongly super-Eddington rates, the jet efficiency of the slim disc is the same as that of the thick disc, since they have the same values of ϕ . For the thin disc, on the other hand, ϕ is generally much lower, except at near-Eddington rates (see equation 19).

The jet efficiency given by equation (25) is shown in Fig. 5 using solid lines. The jet efficiencies can be of order unity (or exceed it) for highly spinning BHs. For moderate and large values of BH spin, jets dominate over thermal feedback at all Eddington ratios except in the thin disc for $f_{\text{Edd},H} \in [0.01, 0.3]$ (jets dominate even in the thin disc for $f_{\text{Edd},H} \gtrsim 0.3$, since the value of ϕ is large enough at those near-Eddington rates to yield jet efficiencies larger than ≈ 1 per cent).

3.6 Black hole spin evolution

The spin of every BH, given by equation (12), evolves due to the following set of physical mechanisms: (1) gas accretion, (2) BH–BH mergers), (3) jet-induced spin-down, (4) radiation torques, and (5) Lense–Thirring torques. These mechanisms affect both the magnitude and direction of the BH spin. However, in practice, the inner region of the accretion disc is usually either perfectly aligned (spin positive) or counter-aligned (spin negative) with respect to the BH spin vector (J. M. Bardeen & J. A. Petterson 1975), or precessing around that same vector (L. Stella & M. Vietri 1998). Thus, almost all quantities that depend on BH spin and that appear in our model can be assumed to depend exclusively on its magnitude and sign.

Generally, the BH spin could occupy the full range $a \in [-1, 1]$, with magnitudes larger than 1 not allowed in general relativity

(R. P. Kerr 1963). In reality, however, the BH spin is thought to have an effective upper limit of $a_{\text{max}} \approx 0.998$ due to torques from photon emission during gas accretion (K. S. Thorne 1974). We thus impose that upper limit (which does have a non-negligible impact on the maximum values of some of the efficiencies that we use).

In Section 3.6.1, we discuss how accretion and jet-induced spind-down impact the evolution of the BH spin magnitude. The choice of the sign of BH spin and the evolution of its direction as a vector are discussed in Sections 3.6.2 and 3.6.3, respectively. Lense–Thirring torques have a great impact on both.

3.6.1 Evolving the magnitude of black hole spin

The BH spin magnitude evolves according to the spin-up/spin-down function s , given by (J. M. Bardeen 1970, R. Moderski & M. Sikora 1996, N. Fanidakis et al. 2011)

$$s \equiv \frac{da}{dM_{\text{acc},H}/M_{\text{BH}}} = \ell_{\text{in}} - 2a(1 - \epsilon_{\text{tot}}) + s_{\text{EM}}, \quad (26)$$

where $\ell_{\text{in}} = L_{\text{in}}/(R_{\text{G}}c)$ is the dimensionless specific angular momentum at the inner boundary of the accretion disc (L_{in} being the specific angular momentum in physical units). $dM_{\text{acc},H}$ is the rest mass of the material accreted through the BH event horizon, and ϵ_{tot} is the total efficiency of all energy loss processes as defined in equation (15). The first term in equation (26) is due to gas accretion on to the BH, the second originates from the definition of the BH spin a through the presence of the BH mass, while the last term encapsulates angular momentum loss due to electromagnetic effects (the launching of winds and mostly jets).

For the thick disc, instead of using equation (26) directly, we use results from GRMHD simulations, to be consistent with our assumptions about feedback efficiencies. We assume the fitting function of R. Narayan et al. (2022) for discs in the MAD




Figure 6. The BH spin-up/spin-down function. This is given by equation (27) for the thick disc at $f_{\text{Edd},d} < 0.01$ (blue line) and equation (28) for the thin and slim discs ($f_{\text{Edd},d} > 0$, red and orange lines, respectively). For the latter, we show the spin-up/spin-down function for several values of $f_{\text{Edd},H}$. Positive or negative values of the spin a correspond to the inner region of the accretion disc co-rotating or counter rotating relative to the BH spin. The bottom left and top right quadrants correspond to the magnitude of the spin increasing, and the top left and bottom right quadrants to the magnitude of the spin decreasing.

state:

$$s_{\text{thick}} = \left(\frac{da}{dM_{\text{acc},H}/M_{\text{BH}}} \right)_{\text{thick}} = 0.45 - 12.53a - 7.8a^2 + 9.44a^3 + 5.71a^4 - 4.03a^5. \quad (27)$$

We show this spin-up/spin-down rate in Fig. 6 using the blue line. In this figure, spin-up (an increase in BH spin magnitude) corresponds to the bottom left and top right quadrants, while spin-down (a decrease in BH spin magnitude) corresponds to the top left and bottom right quadrants. BHs in the thick accretion disc state are almost always spinning down, except in a very narrow region $a \in [0, 0.05]$. This is a signature of strong jet spin-down, and it means that jets will effectively fully spin down a BH given enough time. In practice, however, as the BH spin becomes low enough, feedback efficiencies also drop to low values, making it more likely that the BH will increase its accretion rate due to self-regulation (e.g. C. M. Booth & J. Schaye 2010) and transition to the thin disc state, where spin-up is more efficient (see below). Furthermore, BH–BH mergers (see Section 3.7) will typically leave remnant BHs with non-zero spin, so we generally do not expect to find many BHs with spin values that are very low (e.g. $|a| < 0.3$).

For the slim and thin discs, we use results from A. Ricarte et al. (2023), who provide a fitting formula that smoothly interpolates between the thin disc state without significant jet feedback (for $f_{\text{Edd},H} \ll 1$), and a highly super-Eddington slim disc where jet feedback essentially matches the thick disc (and so jet spin-down should also be similar). Their formula, which is analogous to equation (26), is

$$s_{\text{thin/slim}} = \left(\frac{da}{dM_{\text{acc},H}/M_{\text{BH}}} \right)_{\text{thin/slim}} = s_{\text{HD}} + s_{\text{EM}}, \quad (28)$$

where the first term is a pure hydrodynamical term and the second is an electromagnetic term. s_{HD} is given in their approach by

$$s_{\text{HD}} = \frac{s_{\text{thin}} + s_{\text{min}}\xi}{1 + \xi}, \quad (29)$$

where $\xi = 0.017f_{\text{Edd},H}$, $s_{\text{min}} = 0.86 - 1.94a$ and $s_{\text{thin}} = \ell_{\text{ISCO}} - 2ae_{\text{ISCO}}$ is the spin-up/spin-down function of the thin disc (without outflows and outside the MAD regime), in which ℓ_{ISCO} and e_{ISCO} are the (dimensionless) specific angular momentum and binding energy, respectively, at the ISCO (see Appendix C). The EM term, s_{EM} , is given by⁸

$$s_{\text{EM}} = -\text{sgn}(a)\epsilon_{\text{EM}} \left(\frac{1}{k\Omega_H} - 2a \right), \quad (30)$$

where sgn is the sign function, $\epsilon_{\text{EM}} = \epsilon_{\text{jet}} + \epsilon_{\text{wind}}$ is the total (jet+wind) EM efficiency of the slim disc, and k is given by The slim disc spin-up/spin-down rate is shown in Fig. 6 for several values of $f_{\text{Edd},H}$. For small values, $0.01 < f_{\text{Edd},H} \ll 1$, the spin-up/spin-down function shows monotonic spin-up to $a \approx 1$. As $f_{\text{Edd},H}$ approaches and exceeds unity (super-Eddington accretion), the impact of jets and winds on spin-down becomes stronger, and the spin-up/spin-down function becomes similar to that of the thick disc (albeit there is a small residual difference at $a < 0$, even for strongly super-Eddington slim discs).

$$k = \begin{cases} 0.23, & a \leq 0 \\ \min(0.35, 0.1 + 0.5a), & a > 0 \end{cases} \quad (31)$$

3.6.2 Choosing the sign of BH spin

We assume that the effects of LT torques (J. Lense & H. Thirring 1918) are effectively instantaneous in the inner regions of the accretion disc (compared to the time-scale of accretion). In this picture, the inner accretion disc is defined as the region where the BH’s angular momentum dominates over that of the angular momentum of the accretion disc. As a result, the accretion disc becomes either aligned or counter-aligned with the BH’s spin through LT torques, or it precesses around the BH spin. In the case of counteralignment, we consider accretion to be retrograde and the BH spin negative; otherwise BH spin is positive. The thin disc thus develops a warp due to LT torques and is assumed to be perfectly aligned or counteraligned with the BH within a warp radius R_{warp} (see Appendix E for the expressions that we use), which is the radius out to which the ‘communication’ of the BH and the disc is effective in terms of torques (J. M. Bardeen & J. A. Petterson 1975). Outside this radius, the accretion disc is undisturbed and aligned with the larger-scale accretion flow. For the thick and slim discs, the assumption of exact (counter-) alignment is invalid. Instead, the disc precesses about the BH spin vector. This precession occurs on very short time-scales, much shorter than those we are simulating. For this reason we also assume (counter-) alignment of the thick disc, in a time-averaged sense.

The sign of the BH spin (i.e. whether the disc aligns or counteraligns) is decided based on the criterion developed by A. R. King et al. (2005). In this prescription, the BH and the inner accretion disc are assumed to come into (counter-) alignment in a way that the magnitude of the BH spin does not change, and the

⁸We removed a superfluous minus sign in the original equation provided by A. Ricarte et al. (2023).

total angular momentum (of the BH + inner accretion disc out to R_{warp}) is conserved. In this process, the BH spin aligns with the total angular momentum. The condition for counteralignment (and for BH spin to be negative) in this approach can be stated as follows:

$$\cos \theta < -\frac{J_{\text{warp}}}{2J_{\text{BH}}}, \quad (32)$$

where $\cos \theta = \hat{\mathbf{J}}_{\text{BH}} \cdot \hat{\mathbf{J}}_{\text{d}}$ is the misalignment between the BH spin and the larger-scale angular momentum of the disc (i.e. outside R_{warp}), whose direction is $\hat{\mathbf{J}}_{\text{d}}$. J_{warp} is the total angular momentum of the inner accretion disc out to R_{warp} (before alignment occurs). In our cosmological application of the model, $\hat{\mathbf{J}}_{\text{d}}$ is the kernel-weighted direction of the gas angular momentum, computed from the same gas neighbours that are used to compute the disc-scale accretion rate. This assumption of alignment of the gas angular momentum between our simulated scales and the outer regions of a subgrid accretion disc is quite strong, and discussed in detail in F. Huško et al. (2022).

3.6.3 Evolving the direction of the black hole spin

From a numerical standpoint, the direction of the BH spin is evolved in the following way. For each increment of mass M_{warp} (the mass within R_{warp}) consumed by the BH, the BH-inner accretion disc system is assumed to come into equilibrium (with the inner accretion disc aligned or counter-aligned with the BH), so that the direction of the angular momentum of both the BH and the inner accretion disc is parallel or anti-parallel to the direction of the total angular momentum, $\mathbf{J}_{\text{tot}} = \mathbf{J}_{\text{BH}} + \mathbf{J}_{\text{warp}}$. Here, $\mathbf{J}_{\text{warp}} = J_{\text{warp}} \hat{\mathbf{J}}_{\text{d}}$, is the angular momentum of a single warp increment, which is assumed to be aligned with the angular momentum of the outer accretion disc (i.e. the large-scale accretion flow, which we calculate directly from the simulation). In practice we do not evolve the BH one warp increment at a time, but instead a number of increments, $N_{\text{warp}} = \Delta M / M_{\text{warp}}$, are used over each time-step, where $\Delta M = \dot{M}_{\text{acc},H} \Delta t$ is the total mass to be consumed over a single time-step of length Δt . This approach is valid as long as none of the relevant quantities (BH mass, BH spin magnitude and direction) change appreciably during a single time-step. We ensure this by implementing additional time-step limiters related to BH spin (Section 3.8).

To calculate M_{warp} and J_{warp} , we require knowledge of (1) the warp radius, R_{warp} , (2) the surface density, $\Sigma(R)$, and (3) the specific angular momentum, $L(R)$, of the accretion disc. The warp mass, M_{warp} , is calculated as the integral over the surface density out to R_{warp} , while the warp angular momentum, J_{warp} , includes the specific angular momentum, $L(R)$, in the same integral. For the thin disc we assume Keplerian rotation, $L(R) = \sqrt{M_{\text{BH}} GR}$, which is appropriate if pressure (both radiative and thermal) is negligible compared to gravity. For the thick and slim discs, the specific angular momentum is smaller by a numerical factor of order 0.5, and this is due to thermal and radiation pressure, respectively, being important in those two accretion disc states. We discuss our choices for R_{warp} and Σ in Appendix E. These choices then specify M_{warp} and J_{warp} . We also take into account the possible truncation of the accretion discs at their self-gravity radius, as discussed in Appendix E.

3.7 Final BH spin and mass loss to gravitational waves after BH mergers

The BH spin of a remnant after a BH–BH merger depends on the mass ratio of the BHs, $q = M_{\text{BH},2}/M_{\text{BH},1} < 1$ and three vectors: the two BH spin vectors of the BHs, \mathbf{a}_1 and \mathbf{a}_2 , as well as the orbital angular momentum of the system around their centre of mass, \mathbf{L} . We use a fitting formula from E. Barausse & L. Rezzolla (2009), derived from numerical relativity simulations:

$$\mathbf{a}_{\text{fin}} = \frac{1}{(1+q)^2} (\mathbf{a}_1 + \mathbf{a}_2 q^2 + \mathbf{l}q), \quad (33)$$

where \mathbf{l} is a vector whose direction is the same as that of the orbital angular momentum \mathbf{L} (in the centre-of-mass frame), while its magnitude is given by

$$\begin{aligned} |\mathbf{l}| = & \frac{s_4}{(1+q^2)^2} (|\mathbf{a}_1|^2 + |\mathbf{a}_1|^2 q^4 + 2|\mathbf{a}_1||\mathbf{a}_2| q^2 \cos \phi) \\ & + \left(\frac{s_5 \mu + t_0 + 2}{1+q^2} \right) (|\mathbf{a}_1| \cos \theta + |\mathbf{a}_2| q^2 \cos \xi) \\ & + 2\sqrt{3} + t_2 \mu + t_3 \mu^2. \end{aligned} \quad (34)$$

Here, $\mu = q/(1+q)$ is the symmetric mass ratio, while $s_4 = -0.1229$, $s_5 = 0.4537$, $t_0 = -2.8904$, $t_2 = -3.5171$, $t_3 = 2.5763$ are fitting parameters. The cosines that appear in equation (34) are given by $\cos \phi = \hat{\mathbf{a}}_1 \cdot \hat{\mathbf{a}}_2$, $\cos \theta = \hat{\mathbf{a}}_1 \cdot \hat{\mathbf{l}}$ and $\cos \xi = \hat{\mathbf{a}}_2 \cdot \hat{\mathbf{l}}$.

Given the information available within the model, we could in principle calculate the recoil velocity of the remnant, as well as the total mass fraction lost to gravitational waves. We do not implement the former at this stage, since we cannot reliably track the movement of BHs in their host galaxies due to our use of BH repositioning. However, we do implement the latter. We use results from numerical relativity simulations (E. Barausse et al. 2012) and calculate the final mass of the remnant BH as

$$\begin{aligned} M_{\text{BH,fin}} = & (M_{\text{BH},1} + M_{\text{BH},2}) \left\{ 1 - [1 - e_{\text{ISCO}}(\bar{a})] \mu \right. \\ & \left. - 4\mu^2 [4p_0 + 16p_1 \bar{a}(\bar{a} + 1) + e_{\text{ISCO}}(\bar{a}) - 1] \right\}, \end{aligned} \quad (35)$$

where $p_0 = 0.04827$, $p_1 = 0.01707$ and $e_{\text{ISCO}}(\bar{a})$ is the dimensionless specific binding energy (the expression can be read off from equation (23)) at the ISCO calculated using an effective BH spin variable defined as

$$\bar{a} = \frac{|\mathbf{a}_1| \cos \theta + |\mathbf{a}_2| \cos \xi}{(1+q)^2}. \quad (36)$$

3.8 Additional time-step conditions

Given the changes to the BH model in the form of AGN jets and BH spin evolution, a few additional time-step conditions need to be implemented. The minimum of these time-steps is taken to actually evolve the BH, alongside other time-step conditions already used for the BH in the code. We introduce a jet-related time-step condition that is given by

$$\Delta t_{\text{jet}} = \frac{\Delta E_{\text{jet}}}{P_{\text{jet}}}, \quad (37)$$

where P_{jet} is the current, instantaneous jet power, while ΔE_{jet} is the energy to be given to a pair of particles (see Section 3.10). This time-step limiter ensures that the BH is woken up by the time it needs to ‘hand out’ a pair of kicks.

We also introduce two time-step conditions related to the angular momentum of the BH. The first of these ensures that the

magnitude of the BH spin does not change much over a single time-step, and it is given by

$$\Delta t_a = 0.1 \frac{|a|M_{\text{BH}}}{sM_{\text{acc,H}}}, \quad (38)$$

where $s = da/(dM_{\text{acc,H}}/M_{\text{BH}})$ is the spin-up/spin-down function discussed in Section 3.6.1, given by equations (27) and (28). The numerical factor 0.1 quantifies how finely we want to evolve BH spin; it ensures that the value of BH spin changes by no more than 10 per cent (relative to the current value) over the next time-step.

The second angular momentum time-step condition is related to the redirection of the BH spin vector. Since the BH spin vector may be redirected very quickly relative to its magnitude (due to LT torques), this condition is separate to that mentioned above. This time-step condition is given by

$$\Delta t_a = 0.1 \frac{M_{\text{warp}}J_{\text{BH}}}{M_{\text{acc,H}}J_{\text{warp}} \sin \theta}, \quad (39)$$

where θ is the angle between the current BH spin vector and the angular momentum of gas in the accretion disc on large scales. The numerical prefactor is again present to ensure a sufficiently fine evolution of the BH spin direction. In particular, in the case that the BH spin vector and the gas angular momentum are perpendicular ($\sin \theta = 1$), this criterion will lead to a change of no more than $\approx 5^\circ$ in the BH spin vector direction per time-step.

3.9 Jet launching scheme

For each BH, we parametrize the choice of when and how to kick gas particles in the form of a target jet velocity, v_{jet} . We express the energy received by the particles through the target jet velocity as $\Delta E = 2 \times \bar{m}_{\text{ngb}} v_{\text{jet}}^2 / 2$, where \bar{m}_{ngb} is the mean gas particle mass of neighbours in the SPH kernel of the BH, and the multiplication by two is present since we always kick in pairs. For every BH, jet energy is funneled to a reservoir until it surpasses the value ΔE , at which point a pair of particles are kicked.

We do not kick particles perfectly along the jet direction (which is parallel or anti-parallel with the BH spin vector for each BH), but instead implement a finite half-opening angle $\theta_{\text{jet}} = 7.5^\circ$ (we found little difference when varying this from 0° to 30°). This is accomplished by assigning a new kick direction every time a kick event occurs; this direction is given by a unit vector \mathbf{n}_{jet} that is drawn randomly and uniformly in solid angle within a cone with half-opening angle θ_{jet} directed along the BH spin axis. Since we always kick in pairs, the above procedure is done for one particle in the ‘positive’ direction (along one side of the jet) and for another particle in the ‘negative’ direction (the other side of the jet). We draw only one direction for every pair of kicks, so that the imparted momentum is perfectly opposite in direction for the two particles.

We kick particles by increasing their velocity (in the frame of each BH) by $\Delta \mathbf{v} = \Delta v \mathbf{n}_{\text{jet}}$, where \mathbf{n}_{jet} is the kick direction for a given particle. The magnitude of the velocity increase, Δv , is chosen in such a way that the kinetic energy of each particle increases exactly by $\Delta E_{\text{jet}}/2$. Δv thus depends on the initial particle velocity (in the frame of the BH) and varies for every particle. Conservation of kinetic energy gives

$$\frac{1}{2} m_g (\mathbf{v}_i + \Delta \mathbf{v})^2 - \frac{1}{2} m_g \mathbf{v}_i^2 = \frac{\Delta E_{\text{jet}}}{2}, \quad (40)$$

where \mathbf{v}_i is the initial velocity in the BH frame. This equation can be solved for the magnitude of the velocity increase Δv , yielding

$$\Delta v = \sqrt{v_{i,\text{jet}}^2 + v_{\text{jet}}^2} - v_{i,\text{jet}}, \quad (41)$$

where $v_{i,\text{jet}} = \mathbf{v}_i \cdot \mathbf{n}_{\text{jet}}$ is the initial velocity projected on to the kick direction.

Ideally, one would want to apply radial kicks to particles with respect to the BH. In our scheme, the relative position vector pointing from the BH to each particle may be significantly misaligned with respect to the (new) relative velocity vector. For example, a particle that is near the equatorial plane (defined by the BH spin vector) would be kicked with a velocity that is near-perpendicular to the relative position vector that connects the BH and the particle. While this situation is not desirable, it is our favoured approach for the following reasons. In order to ensure that jets are being launched rather than winds, small opening angles have to be used. For a spherical kernel with ≈ 60 neighbours, a cone with an opening half-angle of 7.5° occupies only ≈ 13 per cent of the kernel volume. Thus, if we were kicking particles radially, on average, we would expect to find only a few neighbours in this region to pick from, assuming an isotropic density distribution. Jet feedback would then become more prone to noise, as we would not find any neighbours to kick in many cases. With non-isotropic particle distributions, e.g. one planar with respect to the BH spin vector, this issue would be exacerbated.

3.10 The numerical implementation of AGN feedback

The thermal mode of AGN feedback is implemented in the same way as in the purely thermal AGN feedback model (see Section 2.4.5), with the only difference being that different feedback efficiencies are used. We implement jets as described in the previous subsection. Both feedback modes use their own energy reservoirs that are independent of each other.

The reader may wonder whether it is reasonable to describe AGN jet feedback using a prescription as simple as pure kinetic kicks, especially given the complexity of jet and lobe physics (e.g. M. A. Bourne & H.-Y. K. Yang 2023), which includes instabilities, relativistic effects, magnetic fields and cosmic rays. These processes are either not resolved or not included in our model. Nonetheless, we argue that the effects of kinetic jets indeed represent the first-order effects of jet feedback as it is thought to be occurring in the real Universe: the driving of anisotropic bow shocks that heat the ambient medium. This issue is discussed further in Appendix F.

The velocity in our model represents both a physical and a numerical parameter. On the physical side, increasing it leads to stronger shocks, and hotter and less dense lobes inflated by jets. Feedback thus becomes more effective. However, increasing it should not affect the dynamical evolution of jets and lobes for very supersonic jets that inflate self-similar lobes (S. A. E. G. Falle 1991; C. R. Kaiser & P. Alexander 1997). In idealized simulations, we found the opposite to be the case: varying jet velocities can have a major impact (F. Huško & C. G. Lacey 2023a, 2023b; F. Huško et al. 2022). This is largely caused by numerical sampling effects. Consider the relation between the total kinetic energy of a jet event and the velocity: $E_{\text{kin}} = 0.5 N_{\text{jet}} m_g v_{\text{jet}}^2$, where N_{jet} is the number of particles kicked into the jet. For a jet event of a fixed energy, increasing the jet velocity by some factor X causes the number of particles by which the jet is sampled to drop as $N_{\text{jet}} \propto$

$1/X^2$. Therefore, the choice of jet velocities heavily impacts the sampling of AGN jet feedback, and it needs to be made carefully.

We assume the following relation between the target jet velocity and BH mass:

$$v_{\text{jet}}(M_{\text{BH}}) = v_{\text{jet},0} \times \min\left(1, \sqrt{\frac{M_{\text{BH}}}{10^9 M_{\odot}}}\right), \quad (42)$$

where $v_{\text{jet},0}$ is the normalization of the scaling (and also the maximum achieved value, for BHs with $M_{\text{BH}} \geq 10^9 M_{\odot}$). We also apply a minimum to the above relation, which is $v_{\text{jet},\text{min}} = 10^{2.5} \approx 316 \text{ km s}^{-1}$. The choice of the normalization, $v_{\text{jet},0}$, is discussed in Section 4. The form of equation (42) is chosen to satisfy the same argument as for thermal feedback (see Section 2.4.5). In particular, given the above form, the energy per feedback event (individual particle kick) scales linearly with BH mass, which means that for every unit of BH growth on a logarithmic scale, BHs are expected to experience a similar number of feedback events.⁹ This in turn means that there is no BH mass scale below which AGN feedback is not sampled sufficiently well. With variable velocities, we sample AGN feedback better than we would be able to if a constant velocity were used.

3.11 Comparison with other models

Accretion disc and BH spin modelling, in the context of galaxy formation and evolution, has previously been included in numerical calculations that use Monte Carlo-generated hierarchical merger trees (e.g. M. Volonteri et al. 2005; S. L. Shapiro 2005), semi-analytical models (e.g. C. D. P. Lagos et al. 2009; N. Fanidakis et al. 2011; E. Barausse 2012; A. J. Griffin et al. 2019) and hydrodynamical simulations (e.g. Y. Dubois, M. Volonteri & J. Silk 2014; D. Fiacconi et al. 2018; S. Bustamante & V. Springel 2019; R. Y. Talbot, M. A. Bourne & D. Sijacki 2020; Y. Dubois et al. 2021; S. Koudmani et al. 2024; L. Sala et al. 2024). Most such models have included only the thin accretion disc, while the thick accretion disc, with its strong jet spin-down effects, is typically neglected. Recent models have begun including it, however (A. J. Griffin et al. 2019; R. Y. Talbot et al. 2020; Y. Dubois et al. 2021; S. Koudmani et al. 2024). The slim disc, which is relevant for super-Eddington accretion, is also usually neglected.

Our subgrid accretion disc and BH spin evolution model inherits many aspects of the A. J. Griffin et al. (2019) semi-analytical model, adapted for use in hydrodynamical simulations by assuming that the angular momentum of a subgrid accretion disc is aligned with that of the gas in the BH kernel. We improve upon that model by explicitly including the thick disc and its associated jet-induced spin-down. We also model super-Eddington accretion using the slim disc, where jets are also very important. Our model is similar to the one employed in the NewHorizon simulations (Y. Dubois et al. 2021), with two main differences: (1) it includes super-Eddington accretion and the slim disc, and (2) it includes an accretion efficiency for the thick disc (also for the slim disc), which is much smaller than unity. In practice, this vastly reduces accretion, feedback and jet-induced spin-down from the thick

disc in our model compared to NewHorizon. The strong jets from the thick disc in NewHorizon are probably the reason why the BH masses and spins are generally too low in that model.

Our model is, in some aspects, quite similar to the unified accretion disc model presented by S. Koudmani et al. (2024), although they do not model the slim disc. Their model includes an accretion efficiency for the thick disc, which varies with the unsuppressed Eddington ratio in the same way as in our model ($\epsilon_{\text{acc}} \propto f_{\text{Edd}}$), although with a much different normalization (so that accretion in the thick disc is less efficient in our model). Their model assumes that the inflow rates of mass and angular momentum on to the accretion disc are well-resolved, which they achieve using super-Lagrangian refinement (e.g. M. A. Bourne et al. 2024; E.-J. Shin et al. 2025). This means they can explicitly track the total mass and angular momentum in the subgrid accretion disc–BH system. Paired with assumptions on the disc structure, their model yields an accretion rate through the disc and on to the BH that is generally much different from the accretion rate on to the disc itself. Due to viscous time delays the accretion rate on to the BH can be very different in magnitude and much less variable.

Including a subgrid accretion disc (as a separate reservoir of mass and angular momentum that modulates the accretion rate) as described above is hard to do in cosmological galaxy formation simulations such as the ones presented in this paper, since we are generally very far from resolving the inflow rates on to the subgrid accretion disc–BH system. For this reason, we assume that mass flows through the disc unimpeded (modulo accretion disc winds that strip away some of the mass) and with no time delay with respect to the inflow on to the disc. In practice, some assumptions on the size of a subgrid accretion disc (see e.g. E. Cenci et al. 2021) could be sufficient, when paired with the assumed mass inflow rate (e.g. equation 2), to model the mass and angular momentum of the disc. This would then yield an accretion rate through the disc and on to the BH, as in the S. Koudmani et al. (2024) model, which might represent a much more realistic accretion model than the one where the BH immediately accretes the material that is falling on to a subgrid disc.

Many previous models have used an accretion efficiency or a concept similar to it (e.g. D. Anglés-Alcázar et al. 2017). Mass loading factors for jets and winds (e.g. E. Choi et al. 2012; D. Rennehan et al. 2024) may produce a similar effect in that only a small fraction of the mass accreting from large scales makes it down to the BH event horizon, with the rest being expelled into feedback by a wind or jet. In these models, it is assumed that the mass being affected by feedback on resolved scales originates entirely from a subgrid wind or jet. This means there is an explicit link between the wind or jet velocity and the mass loading (accretion efficiency). We do not make such an assumption: the mass being heated by thermal feedback or kicked by jets in our model on resolved scales could be the mass in the ambient medium that is swept up by a subgrid wind or jet. Therefore, our accretion efficiencies are independent of the feedback prescription.

Our model is also significantly different compared to others in relation to AGN feedback. Many galaxy formation models use multiple feedback modes, with jets often being used at low Eddington ratios and thermal feedback at high Eddington ratios (e.g. S. Kaviraj et al. 2017; Y. Dubois et al. 2021). Compared to such models, ours differs most in also including jets at high Eddington ratios alongside thermal feedback (in the thin disc), which bears similarity to SIMBA’s low-velocity directed winds at high Eddington ratios (R. Davé et al. 2019). It also differs in

⁹This is not strictly true in the hybrid AGN feedback model, for neither the thermal nor jet feedback modes. The reason is that feedback efficiencies depend on BH spin and Eddington ratio, which introduces an indirect dependence on BH mass. This implies that the amount of energy injected by BHs does not scale linearly with BH mass, as it does in the purely thermal AGN feedback model. However, these effects are relatively small.

including thermal feedback at low Eddington ratios alongside jets (in the thick disc, albeit here the jets are completely dominant). Furthermore, our model includes the slim disc, with both thermal and jet feedback being important (see the similar model by W. Massonneau et al. 2023). Unlike most other hydrodynamical galaxy formation models (except that of NewHorizon; Y. Dubois et al. 2021), all of our AGN feedback efficiencies depend on BH spin. Our AGN feedback is in some respects similar to that of D. Rennehan et al. (2024), who also use three accretion states, both types of spin-dependent feedback across the different states, and feedback mass loadings that effectively behave like our accretion efficiencies (although they are explicitly linked with the feedback prescription).

4 CALIBRATION AND INITIAL COSMOLOGICAL RESULTS

Here we give a brief overview of our calibration efforts, before describing the procedure and its relevant ingredients in more detail. Readers not interested in details of the calibration procedure can skip to Section 4.5, where we discuss the results of our calibration.

In principle, the hybrid AGN feedback model could be calibrated in a similar way as the thermal AGN feedback model used for the flagship COLIBRE simulations (E. Chaikin et al. 2025). This would involve Gaussian-process emulation and a hypercube spanning many parameters, including (1) stellar feedback parameters in the COLIBRE model, (2) thermal AGN feedback parameters present in both the thermal and hybrid AGN feedback models, and (3) new hybrid AGN feedback-related parameters. This process was beyond the scope of our current efforts, but will be done in the future. We opted to instead perform a manual calibration varying a set of parameters as small as possible.

For calibrating the hybrid AGN feedback model, we chose to start from the fiducial set of parameter values found by E. Chaikin et al. (2025) for the thermal AGN feedback model at each resolution. Out of all parameters that vary with resolution, we chose a minimal set that are present in the thermal AGN model and varied these independently at every resolution, alongside a minimal set of parameters present only in the hybrid AGN model. The parameter variation and calibration process, which is described in more detail in Section 4.4, was used to first calibrate the hybrid AGN feedback model at the m7 resolution. That set of parameter values was then used as the best guess for the calibration at the m6 resolution. We found that less than half of our calibration parameters needed changing at higher resolution, none of which were the new parameters introduced in the hybrid AGN feedback model. Once a satisfactory model (when compared with calibration observables) was found for the m6 resolution by tuning these few parameters, this model was used as a best guess for m5, and they were then tuned again until an acceptable model was found.

In Section 4.1, we provide an overview of the parameters that we varied, and a brief motivation for why we chose to vary them. In Section 4.2, we discuss which galaxy and BH relations and properties were chosen for calibration, as well as details of observational data used for comparison during calibration. Section 4.3 includes the variations of several of the important calibration parameters, and the impact of these variations on the calibration relations. Informed by these variations, in Section 4.4, we discuss how we carried out the actual calibration. Finally, in Section 4.5 we discuss the calibrated parameter values that result from our

calibration procedure. We also discuss the results of the largest hybrid AGN simulations in terms of the calibration observables, and the convergence between resolutions.

4.1 The choice of calibration parameters

The five parameters that were varied are:

- (i) $v_{jet,0}$: the normalization (which is also the maximum) of the jet velocity–BH mass scaling (equation 42),
- (ii) $r_{tr,0}$: the dimensionless size of the thick accretion disc at a critical Eddington ratio $f_{Edd,d} = 0.01$, which determines the accretion efficiency at $f_{Edd,d} < 0.01$ (see Section 3.3), such that $\epsilon_{acc,th} \propto f_{Edd,d} / \sqrt{r_{tr,0}}$,
- (iii) ϵ_f , the coupling efficiency for thermal AGN feedback in the thin accretion disc state at $0.01 < f_{Edd,d} < 1$ (Section 3.5.1),
- (iv) $M_{BH,seed}$, the BH seed mass (Section 2.4.1),
- (v) $n_{H,pivot}$, the pivot density for SNIa and CCSN feedback heating temperatures (see equation 1).

The first two parameters are present only in the hybrid AGN feedback model, and relate to jet velocities and accretion efficiencies, while the remaining three parameters are present both in the hybrid and in the fiducial, thermal-only AGN feedback model. In principle, we could have varied other hybrid AGN parameters. However, these parameters (whose values were set in Section 3) were not varied for two reasons: (1) the effects of varying them are generally smaller than for the above two hybrid AGN parameters (as confirmed by many variations not shown here), especially for our target calibration observables, and (2) they are better constrained by either theory or observations. In most cases, both of these reasons were true at the same time. The two new AGN parameters that we vary have both unconstrained values and have a strong effect on our target calibration observables across the ranges we considered (see Sections 4.3.1 and 4.3.2).

We found that the modifications introduced in the hybrid AGN feedback model have a large effect on how feedback operates and on how BHs self-regulate their growth. For this reason, varying the coupling efficiency, ϵ_f , is necessary to reproduce the observed $M_{BH}-M_*$ relation at the massive end ($M_* > 10^{10} M_\odot$). Allowing the variation of the BH seed mass, on the other hand, is necessary since that parameter affects how quickly BHs grow and quench their host galaxies, a process that is inherently different in the hybrid AGN feedback model compared to the thermal one. We opted not to vary the AGN heating temperature, ΔT_{AGN} , since it has a widespread impact on nearly all galaxy and BH properties, and there is currently no single relation or property that can be used to constrain its value, especially not when there is additionally a second mode of AGN feedback operating (whose strength is modulated by jet velocities).

The final parameter, the SNIa and CCSN pivot density, $n_{H,pivot}$, is varied since we find AGN feedback to have stronger effects in the low-mass regime ($M_* < 10^{10} M_\odot$) in the hybrid AGN feedback model than in the thermal one (see Section 4.3). It is thus also necessary to change at least one stellar feedback-related parameter, since the interplay between AGN feedback and stellar feedback is different. We chose to vary $n_{H,pivot}$ since the effects of varying it are strongest for intermediate-mass galaxies ($M_* \approx 10^{10} M_\odot$), where we also found the differences between the thermal and hybrid AGN feedback models to be maximal in terms of the observables used for calibration.

4.2 Calibration relations and observational data

The observables considered for calibration are largely the same as for the thermal AGN feedback model (E. Chaikin et al. 2025). We calibrate on (1) the $z = 0$ GSMF, (2) the $z = 0$ projected galaxy stellar half-mass radius–stellar mass relation, and (3) the $z = 0$ black hole mass–stellar mass relation. However, to this list we add a fourth observable: (4) the bolometric AGN luminosity function (LF) (also at $z = 0$). This is added since the accretion efficiency that we have introduced has a strong effect on the growth history of BHs and how many of them are in the thick disc, radiatively inefficient regime ($f_{\text{Edd,d}} < 0.01$) at $z = 0$, and how many in the thin disc, radiatively efficient regime ($f_{\text{Edd,d}} > 0.01$). This parameter thus strongly affects the balance of accretion disc states and feedback modes, and this is inherently tied to the AGN LF (see Section 4.3.2).

The observational data, which we introduce in the following subsections, are all converted to the cosmology we adopted for COLIBRE, and to the Chabrier IMF, also used in COLIBRE. For calibration against observational data, we only consider galaxies with stellar masses above 100 resolution elements, which corresponds to $M_* \approx 10^9, 10^8, \text{ and } 10^7 M_\odot$ for m7, m6 and m5, respectively. Galaxies with fewer stellar particles are poorly resolved, and we do not consider the simulation predictions to be reliable in that stellar mass regime.¹⁰ We do not consider observational comparisons at $M_* \gtrsim 10^{11.3} M_\odot$, since the effects of aperture sizes, projection, and observational sensitivity become more important for central galaxies of galaxy groups and clusters. When plotting any of our calibration relations that involve the stellar mass, we use fixed bin widths of 0.25 dex. Next, we detail individual observational datasets used for calibration.

4.2.1 Galaxy stellar mass function

The GSMF data that we compare with is from S. P. Driver et al. (2022) (their table 6), which corresponds to the fourth data release of the GAMA survey. We increase the normalization of their GSMF by 0.0807 dex (as recommended by the authors) to account for two effects: (1) the data is at a median redshift of $z = 0.03$, whereas we are interested in comparing at $z = 0$, and (2) the small GAMA volume is biased and not representative of the cosmic mean.

Galaxy masses determined observationally exhibit both random (statistical), as well as systematic uncertainties. While we cannot correct for systematic uncertainties, it is important for our comparison to account for the effects of Eddington bias that appear in these observations, which can shift the GSMF towards different values (J. Schaye et al. 2025). The stellar masses in GAMA have a 0.1 dex random uncertainty. We thus apply a 0.1 dex random lognormal scatter to our simulated stellar masses, which mimics Eddington bias, i.e. the effects of random observational uncertainties in determining stellar masses. We refer the reader to fig. D1 in J. Schaye et al. (2025) for the effects of applying different amounts of this scatter on the main observables, using

¹⁰The actual particle limit down to which a simulation prediction is reliable depends on what is being considered. For example, we find that the galaxy stellar mass function can be correctly predicted with as few particles as 10 or even less. On the other hand, galaxy sizes may only start being reliable once resolved with a few hundred or even 1000 particles. For simplicity, we use a single threshold of 100 particles for all quantities considered here, but discuss the role of the threshold where relevant.

the thermal AGN feedback model. The effects are very similar for the hybrid model.

4.2.2 Galaxy size–mass relation

We calibrate galaxy sizes by considering the galaxy size–stellar mass relation, which we compare with J. A. Hardwick et al. (2022) (their table C1), who reported projected half-mass galaxy sizes from a mass-selected sample of galaxies in the xGASS spectroscopic survey (B. Catinella et al. 2018). We define galaxy sizes as the radii that contain, in projection (along a random direction), half of the 50 kpc 3D mass of a given galaxy. We apply 0.1 dex lognormal scatter to the stellar masses, as for the GSMF.

4.2.3 Black hole mass–stellar mass relation

To ensure the simulated BH masses are realistic, we calibrate using the BH mass–stellar mass relation at $z = 0$. We primarily compare with data from A. W. Graham & N. Sahu (2023) (their table 1, with stellar masses corrected as per the erratum A. W. Graham & N. Sahu 2024), who performed morphological decompositions and defined galaxies as elliptical, lenticular or discy. They found that the BH mass–stellar mass relation was largely similar for ellipticals and lenticulars, but differs from that of discs, which have significantly lower BH masses at fixed stellar mass. We construct an overall median BH mass–stellar mass relation, representative of the entire galaxy population, by binning the data from A. W. Graham & N. Sahu (2023), using fractions of different morphological types of galaxies as weights. This procedure is described in more detail in Appendix G, where we also provide, for convenience, the binned median BH masses with the associated uncertainties on both the median stellar and BH masses in the bins. While only ≈ 150 (in total) BHs were included in these bins, this is sufficient to constrain the median BH masses to within an uncertainty of ≈ 0.1 dex per stellar mass bin. For reference, we also compare with the power law fit from N. J. McConnell & C.-P. Ma (2013), which is based on a smaller sample (with fewer spirals), and which was obtained without weighting the contributions of different morphological types. Therefore, this fit is closer to describing the $M_{\text{BH}}-M_*$ relation for spheroid-dominated galaxies than for the overall population, and is not expected to agree with our binned weighted median points based on A. W. Graham & N. Sahu (2023), nor with our simulations.

The simulated BH mass used to compare with the above data is the subgrid mass of the most massive BH found within a 50 kpc 3D aperture around a given galaxy. We apply 0.15 dex lognormal scatter to the simulated stellar masses of the galaxies hosting these BHs, which is slightly larger than for the GSMF and size comparisons, but more consistent with the stellar mass uncertainties found by A. W. Graham & N. Sahu (2023). We only compare at $M_* > 10^{10} M_\odot$, since the observational data used here only extends down to such masses.

4.2.4 Bolometric AGN luminosity function

For the AGN bolometric LF we use the compilation from X. Shen et al. (2020),¹¹ who estimated bolometric luminosities from soft X-ray, hard X-ray and mid-infrared AGN observations. We compare to their $z = 0.2$ data, using the simulation predictions at the

¹¹Taken directly from the data repository at <https://github.com/gkulkarni/QLF/blob/master/Data/allqlfs.dat>.

same redshift. We also compare with their ‘global’ fit (i.e. using redshifts $0 < z < 6$ to constrain their fitting function). The AGN luminosities, used to construct our predictions for the bolometric AGN LF, are computed using equation (22). We add 0.3 dex random lognormal scatter to these luminosities, mimicking typical observational uncertainties in inferring bolometric luminosities. When comparing with the observed (inferred) AGN bolometric LF, we only consider luminosities¹² $L_{\text{AGN,bol}} > 10^{42} \text{ erg s}^{-1}$.

This lower limit is imposed for three reasons. First, observational inferences of the AGN LF are likely very uncertain at these low luminosities, and affected by systematics (X. Shen et al. 2020). Secondly, we find that nearly all of the BHs that contribute to the bolometric AGN LF at $L_{\text{AGN,bol}} > 10^{42} \text{ erg s}^{-1}$ are in the thin disc state. On the other hand, nearly all BHs in the thick disc state ($f_{\text{Edd,d}} < 0.01$) have luminosities below this threshold value. We are confident about our predictions for the luminosities from the thin disc, since they use radiative efficiencies of order 10 per cent (but dependent on BH spin) implied by standard accretion disc theory (I. D. Novikov & K. S. Thorne 1973). On the other hand, thick disc radiative efficiencies (we use those from R. Mahadevan 1997, appendix D) are very uncertain. Thirdly, from a practical standpoint, we find that our predicted AGN LFs converge with resolution for luminosities larger than the above-mentioned critical value, while they increase with higher resolution at values below it (see Section 4.6).

4.3 Single parameter variations

Here we vary most of the parameters that are used in our calibration procedure (see E. Chaikin et al. 2025 for the effects of variations of the other parameters). These variations are summarized in Table 2. We perform all variations at m7 resolution, using $(50 \text{ Mpc})^3$ simulation volumes (i.e. L050). The effects of these variations are generally similar at higher resolutions, so we do not show them here. The fiducial hybrid model used as the baseline in these comparisons is that found by calibration as described in Sections 4.4 and 4.5. Here, we do not discuss how well the fiducial model fits the data; this is discussed in Section 4.5.

We vary: the jet velocity normalization, $v_{\text{jet},0}$ (Section 4.3.1), the accretion efficiency parameter, $r_{\text{tr},0}$ (Section 4.3.2), and the thermal feedback coupling efficiency, ϵ_f (Section 4.3.3). We vary the first two parameters since they are new parameters introduced in the hybrid AGN feedback model, and since they impact nearly all galaxy and BH properties. While the coupling efficiency is not a new parameter, its variation has more significant consequences than it did in the OWLs and EAGLE models (C. M. Booth & J. Schaye 2010; J. Schaye et al. 2015). We briefly discuss the impact of varying the BH seed mass and the stellar feedback pivot density in Section 4.3.4. They control the strength of AGN and stellar feedback, respectively.

4.3.1 The impact of varying jet velocities

In Fig. 7, we show the impact of varying the jet velocities, specifically the velocity normalization parameter $v_{\text{jet},0}$, on the calibration observables that we described in the previous subsection. The velocity normalization, $v_{\text{jet},0}$, is varied up and

¹²Even this limit could be too low. Observed X-ray LFs are typically not reliable below $L_X = 10^{42} \text{ erg s}^{-1}$, due to contamination from the host galaxy. This limit corresponds to approximately $L_{\text{bol}} = 10^{43} \text{ erg s}^{-1}$, and it may be redshift-dependent.

Table 2. The simulations used for our parameter variation study (all in 50^3 Mpc^3 volume boxes and at m7 resolution). The first column is the simulation name, while the remaining columns are the parameters being varied: (1) the jet velocity normalization (note that $v_{\text{jet},0} \propto \sqrt{M_{\text{BH}}}$, see equation 42), (2) the accretion efficiency parameter $r_{\text{tr},0}$ (note that $\epsilon_{\text{acc,thick}} \propto f_{\text{Edd,d}} / \sqrt{r_{\text{tr},0}}$ for $f_{\text{Edd,d}} < 0.01$, see equation 16), (3) the thermal feedback coupling efficiency in the thin disc ($0.01 < f_{\text{Edd,d}} < 1$). Other parameters are kept fixed at their fiducial values found through calibration (Table 3). The first row is the fiducial thermal AGN feedback simulation (for reference), the second is the fiducial hybrid AGN feedback simulation, while other groups of rows indicate different simulations where a single parameter is varied in the hybrid model (with the varied values highlighted in bold).

Simulation	$v_{\text{jet},0}$ (km s ⁻¹)	$r_{\text{tr},0}$	ϵ_f
L050m7 (thermal)	N/A	N/A	0.1
L050m7h (fiducial)	$10^{4.5}$	10^4	0.03
L050m7h_low_ V_j	10^4	10^4	0.03
L050m7h_high_ V_j	10^5	10^4	0.03
L050m7h_const_ V_j	N/A^a	10^4	0.03
L050m7h_const_ V_j _recal ^b	N/A^a	10^4	0.03
L050m7h_low_ ϵ_{acc}	$10^{4.5}$	10^6	0.03
L050m7h_high_ ϵ_{acc}	$10^{4.5}$	10^2	0.03
L050m7h_const_ ϵ_{acc}	$10^{4.5}$	N/A^c	0.03
L050m7h_const_ ϵ_{acc} _recal ^b	$10^{4.5}$	N/A^c	0.03
L050m7h_low_ ϵ_f	$10^{4.5}$	10^4	0.01
L050m7h_high_ ϵ_f	$10^{4.5}$	10^4	0.1

Notes.^aThese simulations did not use the $v_{\text{jet},0}$ parameter; they instead used constant jet velocities equal to 10^4 km s^{-1} .

^bThese simulations had further parameters (the BH seed mass and the stellar feedback pivot density) varied, beyond the ones shown here.

^cThese simulations did not use the $r_{\text{tr},0}$ parameter; they instead had constant accretion efficiencies $\epsilon_{\text{acc}} = 1$ in the thick disc.

down by factors of ≈ 3.16 (0.5 dex) relative to the fiducial value ($10^{4.5} \approx 31\,600 \text{ km s}^{-1}$). This corresponds to ranges of energy per feedback event differing by an order of magnitude. These variations change how explosive,¹³ intermittent, and well-sampled the feedback is. For reference, we performed a run with a constant velocity $v_{\text{jet}} = 10^4 \text{ km s}^{-1}$, as opposed to the fiducial $v_{\text{jet}} \propto M_{\text{BH}}^{0.5}$ scaling. Note that comparisons of our fiducial model with these higher- and lower-velocity runs, as well as with the one using a constant velocity, cannot be used to rule out these models. A recalibration should first be attempted on such models. The low-velocity variation shown in Fig. 7 was our initial attempt at finding a best-fitting model. We have attempted to recalibrate the constant-velocity run; this is discussed in Appendix H. As we show there, a constant-velocity model that matches all the observed calibration data can indeed be found. We opted for variable velocities, however, to better sample AGN feedback in low-mass galaxies, and to avoid an unphysical BH mass–stellar mass relation, as explained in Section 2.4.5.

Increasing jet velocities results in a slight reduction of the GSMF (top left) at $M_* > 10^{10.5} M_\odot$. This is due to the more explosive nature of feedback, which tends to reduce radiative energy losses and make feedback more efficient. If the jet velocities are reduced, the number of galaxies in this mass range increases. Somewhat counterintuitively, larger velocities lead to an increase in the GSMF at $10^9 < M_* < 10^{10} M_\odot$, probably

¹³By ‘explosive’ we refer to how much energy a single gas particle receives as part of a feedback event. More explosive feedback events are less prone to numerical overcooling and drive stronger outflows.




Figure 7. The impact of varying the jet velocities (see Section 3.10) on galaxy and BH observables considered for calibration, in L050m7 simulations. From left to right and top to bottom we show, in order the (1) galaxy stellar mass function, (2) median galaxy projected size–stellar mass relation, (3) median BH mass–stellar mass relation, and (4) AGN bolometric LF. The first three are shown at $z = 0$, while the AGN LF is shown at $z = 0.2$. The grey dashed line shows the thermal AGN model, the black solid line shows the fiducial hybrid AGN model (with the shaded region showing uncertainties around that curve, which are Poisson errors in the case of the galaxy stellar mass and AGN LF functions, and standard errors of the median in the case of galaxy sizes and BH masses), and coloured curves show variations relative to the fiducial hybrid AGN simulation (see legend). Lines are dotted at $M_* < 10^9 M_\odot$, which corresponds to roughly 100 star particles, indicating limitations due to resolution. The velocity normalization was varied up and down by factors of 3.16 (0.5 dex; orange and blue, respectively), corresponding to increasing and decreasing the energies of individual feedback events by a factor of 10. The variations were done relative to the fiducial value of $v_{\text{jet},0} = 10^{4.5} \approx 31\,600 \text{ km s}^{-1}$. We also performed a simulation with a constant velocity of $v_{\text{jet},0} = 10^4 \text{ km s}^{-1}$ (green). Observational data are taken from (1) S. P. Driver et al. (2022), (2) J. A. Hardwick et al. (2022), (3) A. W. Graham & N. Sahu (2023) and N. J. McConnell & C.-P. Ma (2013), and (4) X. Shen et al. (2020). Further details are discussed in Section 4.2.

because AGN feedback is sampled more poorly and is thus less effective, despite being more explosive. Lower velocities similarly reduce the number of low-mass galaxies. Small differences in galaxy masses between the different runs persist down to $M_* \approx 10^8$, indicating that AGN feedback affects dwarf galaxies. The run that has a constant $v_{\text{jet}} = 10^4 \text{ km s}^{-1}$ features a similar number of high-mass galaxies as the other runs, with the exception of the low-normalization one, but the number of galaxies around the knee ($M_* \approx 10^{10} M_\odot$) is significantly increased. This is probably caused by AGN feedback being poorly sampled, and therefore less effective, in low-mass galaxies.

Differences in the GSMF at lower masses could merely indicate a horizontal shift; if there is a greater number of more massive galaxies, there must also be a smaller number of less massive galaxies. In Appendix I, we show the effect of the same jet velocity variations on the stellar mass–halo mass relation. There, it is clear that higher jet velocities cause more effective feedback in haloes more massive than $10^{12} M_\odot$, leading to decreased stellar masses at fixed halo mass. The opposite is true for lower-mass haloes.

Galaxy sizes (top right) show a picture consistent with stellar masses in terms of jet velocity variations: increasing (decreasing) jet velocities results in massive galaxies, $M_* > 10^{10} M_\odot$, being

larger (smaller) due to more effective AGN feedback, and the opposite is true for less massive galaxies. Galaxy sizes show differences between the different simulations down to a mass of $M_* \approx 10^8 M_\odot$, as for the GSMF. The constant-velocity case has the largest galaxies, despite not having the largest velocities at high masses. This is likely due to the assembly of galaxies and their cores occurring at high redshifts, when the BH mass, and therefore the jet velocity (in the fiducial, variable-velocity case) is not large.

BH masses (bottom left) are higher (lower) at fixed stellar mass, if higher (lower) jet velocities are used. However, this effect is driven by the fact that galaxy masses are changed when jet velocities are changed, shifting the same BHs horizontally in the $M_{\text{BH}}-M_*$ diagram. We checked the BH mass–halo mass relation and found it to be unaffected by the velocity variation. This suggests that BH self-regulation is tight and independent of jet velocity, at least to the degree that we varied it here. The case of constant $v_{\text{jet}} = 10^4 \text{ km s}^{-1}$ has similar BH masses as the other runs (except the run with low velocities).

Increasing jet velocities causes a slight increase (decrease) in the bolometric AGN LF (bottom right) if jet velocities are increased (decreased), but only at $L_{\text{AGN,bol}} < 10^{44} \text{ erg s}^{-1}$. The AGN LF is also slightly higher in the case of constant $v_{\text{jet}} = 10^4 \text{ km s}^{-1}$. These changes are probably related to how well AGN feedback is sampled.

The impact of these variations demonstrates that jet velocities cannot be calibrated by considering only a single one of the calibration observables. This partly motivates the calibration approach that we outline in Section 4.4.

4.3.2 The impact of varying the accretion efficiency in the thick disc state

Fig. 8 has the same format as Fig. 7, but shows the effects of varying the thick disc accretion efficiency ϵ_{acc} through variations of the dimensionless size of the disc ($r_{\text{tr},0}$) at $f_{\text{Edd},d} = 0.01$. We vary $r_{\text{tr},0}$ by factors of 100, which changes the accretion efficiency normalization by factors of 10, relative to the fiducial scaling $\epsilon_{\text{acc}} = 0.0316(f_{\text{Edd},d}/0.01)$. Similar to what we did for the velocities, we also show for reference a simpler case that uses no accretion efficiency (i.e. $\epsilon_{\text{acc}} = 1$ in the thick disc). In this figure, we do not show a 100 per cent efficiency model that has been recalibrated, but instead show the first-order effects of setting the accretion efficiency to unity. An attempt at recalibration is shown in Appendix H, where we show that a recalibrated model with a 100 per cent accretion efficiency cannot reproduce the AGN LF.

The GSMF (top left) is significantly affected by changes to the accretion efficiency only at $M_* > 10^{9.5} M_\odot$, although small differences exist at lower masses. The number of intermediate-mass galaxies at the knee of the GSMF ($M_* = 10^{10} M_\odot$) drops monotonically with increasing ϵ_{acc} . The number of massive galaxies ($M_* \approx 10^{11} M_\odot$) is similar in all cases, except for the $\epsilon_{\text{acc}} = 1$ case, where the number of such galaxies is slightly lower. As shown in Appendix I, at fixed halo mass, an increase in the accretion efficiency causes a decrease in stellar masses.

The galaxy sizes (top right) show a picture consistent with stellar masses, with massive galaxies unaffected by variations of the accretion efficiency. For lower-mass galaxies ($M_* < 10^{11} M_\odot$) there is a monotonic trend with accretion efficiency, where higher values lead to larger sizes (all the way to $M_* = 10^8 M_\odot$). This shows that thick disc feedback is important in low-mass galaxies, and not just in massive ones.

BH masses (bottom left) are largely unaffected, although higher accretion efficiencies lead to a small drop in the normalization of the relation.

The bolometric AGN LF (bottom right) is strongly affected by variations of the accretion efficiency. This is caused by the accretion efficiency also changing the effectiveness of jet feedback in the thick disc, which in turn changes the likelihood that a BH is in the thick disc state ($f_{\text{Edd},d} < 0.01$) rather than in the radiatively efficient thin disc state ($f_{\text{Edd},d} > 0.01$). This appears counterintuitive, but can be explained by self-regulation. In the self-regulated growth and feedback picture, BHs will, on average, tend to have the same feedback powers in different simulations (C. M. Booth & J. Schaye 2009; C. M. Booth & J. Schaye 2010; H. Li et al. 2025). As a result, if the accretion efficiency is decreased, BHs adjust by increasing their accretion rates in order to have the same feedback powers. However, this makes them more likely to enter the thin disc state. BHs in the thin disc state are visible as AGN, and are virtually the only ones contributing to the observed AGN LF at $z = 0$, at least at $L_{\text{AGN,bol}} > 10^{42} \text{ erg s}^{-1}$. The sensitivity of the bolometric AGN LF to the accretion efficiency was used to calibrate the accretion efficiency parameter, as illustrated in the figure. The strong disagreement of the fully efficient accretion in the thick disc ($\epsilon_{\text{acc}} = 1$ for $f_{\text{Edd},d} < 0.01$) with the observations is one of the main motivations for our introduction of the accretion efficiency.

4.3.3 The impact of varying the coupling efficiency

In Fig. 9, we show the effects on the calibration relations of varying the coupling efficiency of radiation from the thin disc, at $0.01 < f_{\text{Edd},d} < 1$, which is the main channel of thermal AGN feedback in our simulations (Section 6.2). We vary the value by approximately 0.5 dex up and down relative to the fiducial value, which is $\epsilon_f = 0.03$ at m7 resolution. We find that the effects of varying it on galaxy properties are widespread, persisting down to $M_* \approx 10^8 M_\odot$, and they also have a mild impact on the AGN LF. A higher thin disc coupling efficiency results in higher stellar masses (see also the effects of this variation on the stellar mass–halo mass relation, shown in Appendix I).

The effects on galaxy properties may appear to be counterintuitive; increasing the coupling efficiency (i.e. making feedback stronger) results in higher stellar masses and smaller galaxies, which one would intuitively expect to happen when making feedback weaker. However, these effects are an indirect consequence of the main impact of varying the coupling efficiency, which is a change in BH masses (visible in the bottom left panel of Fig. 9). Increasing the coupling efficiency reduces BH masses, and vice versa. This is again a result of self-regulation (e.g. C. M. Booth & J. Schaye 2009; C. M. Booth & J. Schaye 2010; H. Li et al. 2025). In short, feedback from BHs adjusts to yield approximately the same outflow powers (on average) in different simulations.¹⁴ If one uses lower coupling efficiencies, BHs have to grow more

¹⁴In previous simulations, which used constant heating temperatures for AGN feedback, feedback was not sampled well in low-mass galaxies. Self-regulation was thus only important for massive galaxies ($M_* > 10^{10} M_\odot$), where AGN feedback was dominant. In COLIBRE, especially in the hybrid AGN feedback model, AGN feedback is important for a wider range of galaxy masses. The $M_{\text{BH}}-M_*$ relation is not as steep around $M_* \approx 10^{10} M_\odot$ as it was in EAGLE (R. G. Bower et al. 2017), and self-regulation applies more generally in COLIBRE, not just for massive galaxies dominated by AGN feedback.

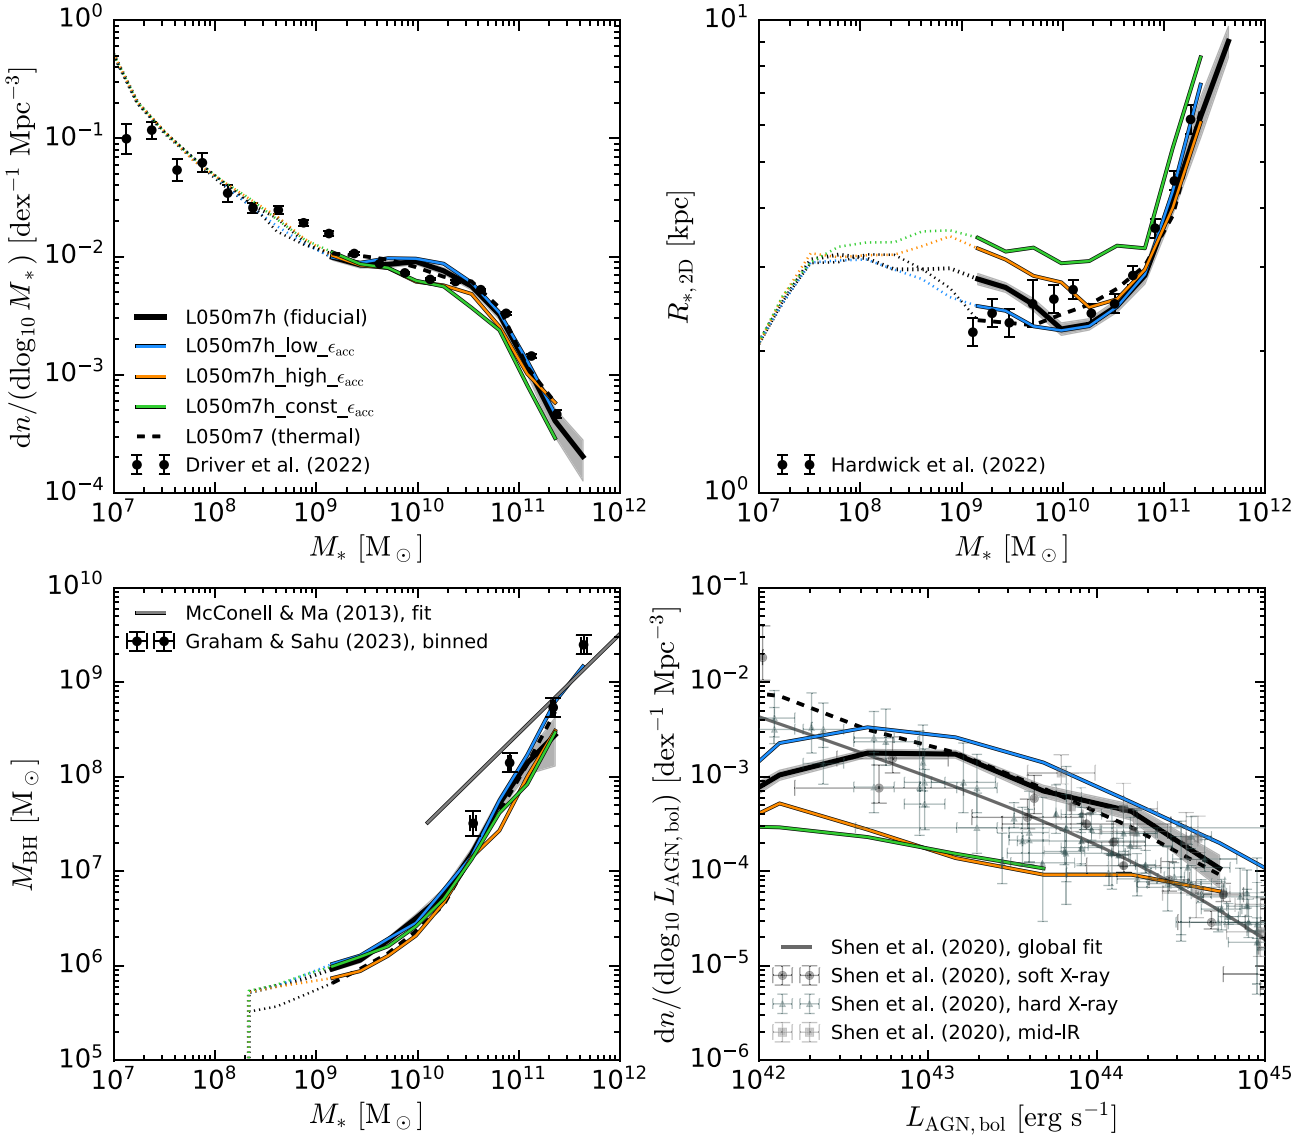


Figure 8. As Fig. 7, but showing the impact of varying the thick disc ($f_{\text{Edd},d} < 0.01$) accretion efficiencies through variations of the parameter $r_{\text{tr},0}$ (note that $\epsilon_{\text{acc,th}} \propto f_{\text{Edd},d}/\sqrt{r_{\text{tr},0}}$, see Section 3.3.1, equation 18). The fiducial value of ($r_{\text{tr},0} = 10^4$) yields the scaling for $f_{\text{Edd},d} < 0.01$. Here, we show the effect of varying $r_{\text{tr},0}$ up and down by factors of 100 (orange and blue, respectively), which changes the normalization of the $\epsilon_{\text{acc}} - f_{\text{Edd},d}$ relation by factors of 10. We also compare with a case of fully efficient accretion in the thick disc, i.e. where $\epsilon_{\text{acc}} = 1$ (green).

in order to have the same feedback powers. The change in BH masses then indirectly impacts galaxy properties, since we tie AGN heating temperatures and jet kick velocities explicitly to BH mass (Section 3.10). Increasing the coupling efficiency decreases the BH masses, which in turn makes AGN feedback weaker on galaxy scales. This demonstrates that in our model, in terms of parameters, changing feedback efficiencies is not dominant, as BHs tend to adjust to inject as much energy as they need to regulate their own growth. Instead, the efficiency of AGN feedback (in terms of how efficiently it quenches galaxies) is mostly governed by how explosive it is, which is instead controlled by heating temperatures and jet velocities.

4.3.4 The impact of varying other calibration parameters

We vary neither the BH seed mass nor stellar feedback pivot densities here, as the effects of those variations are similar in the hybrid and thermal AGN feedback models. We refer the reader to E. Chaikin et al. (2025) for a detailed discussion of such variations. In short, increasing the BH seed mass causes earlier AGN growth and feedback, and therefore earlier quenching of galaxies (making them less massive and larger). It is most relevant for intermediate-mass and massive galaxies ($M_* > 10^{10} M_\odot$). Increasing the pivot density decreases the heating temperatures of SN feedback, making it less effective at low-to-intermediate galaxy masses ($10^9 < M_* < 10^{10.5} M_\odot$), thus increasing their

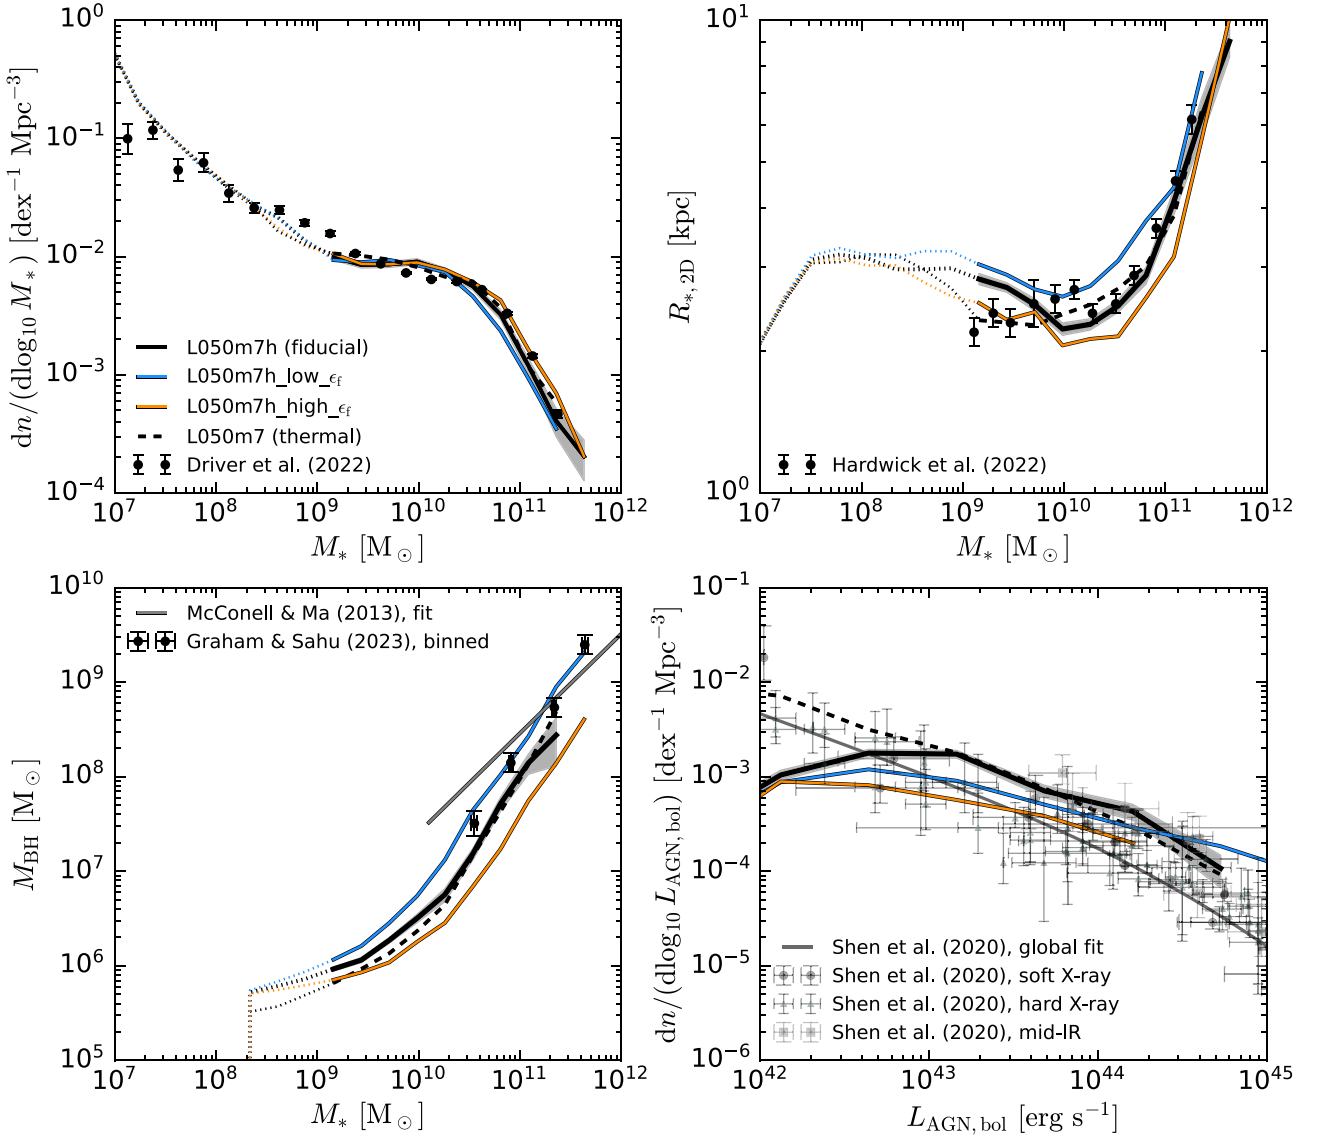


Figure 9. As Fig. 7, but showing the impact of varying the thin disc ($0.01 < f_{Edd,d} < 1$) coupling efficiency, ϵ_f , used for thermal feedback (see Section 2.4.5), on the galaxy and BH properties considered for calibration. The fiducial value of the coupling efficiency at m7 resolution is $\epsilon_f = 0.03$; we vary this value by ≈ 0.5 dex to $\epsilon_f = 0.01$ and $\epsilon_f = 0.1$.

masses and decreasing their sizes. Variations of both of these parameters have very little effect on the BH mass–stellar mass relation and on the AGN LF.

4.4 Calibration procedure

In Section 4.3, we showed how varying the jet velocities, the accretion efficiencies and the coupling efficiencies affects the main galaxy and BH properties that we calibrate on (Section 4.2). We also use information on how the variations of the other two parameters (Section 4.1), the BH seed mass and the stellar feedback pivot density, affect the galaxy and BH properties (Section 4.3.4). Informed by this knowledge, our calibration procedure consists of the following steps:

Step 1: Choose a value of the jet velocity normalization, $v_{jet,0}$.

Step 2: Given a choice of $v_{jet,0}$, calibrate the coupling efficiency ϵ_f and the accretion efficiency parameter $r_{tr,0}$. These are

calibrated on the M_{BH} – M_* relation and the AGN LF at $z = 0$, respectively. More details of this are given below.

Step 3: Given a $v_{jet,0}$ and the newly calibrated values of ϵ_f and $r_{tr,0}$, calibrate the seed mass $M_{BH,seed}$ and the SNIa and CCSN pivot density $n_{H,pivot}$, using the $z = 0$ GSMF and stellar size–mass relation.

Step 4: Repeat Steps 1–3 for different values of $v_{jet,0}$. The fully calibrated models with different choices of $v_{jet,0}$ are then compared, and the one that best fits all of the data is chosen as the final model.

This procedure was performed first at m7 resolution, using (25 Mpc) 3 volumes (L025). The final models (in Step 4) were compared using larger (L050) simulations, since they provide better statistics and include more massive galaxies. The best-fitting set of parameter values found at m7 was used as the initial guess for m6, which was then recalibrated. The same was then repeated for m5, using the set of parameter values found for m6 as the initial guess. For m6 and m5, the full procedure was not repeated. We did

not recalibrate $r_{\text{tr},0}$ (in Step 2) since we found it to be unnecessary. ϵ_f was only slightly adjusted (in Step 2) at m6 compared to m7, but remained the same at m5 compared to m6. Only the BH seed mass and SN pivot density were significantly varied for m6 and m5 (in Step 3). The above procedure was done with ≈ 30 simulations at m7, and ≈ 10 each at m6 and m5.

The calibration was split into the above steps for the following reasons. The jet velocity affects nearly all observables that we consider in the calibration (see Section 4.3.1). Models with different values of $v_{\text{jet},0}$ thus cannot be calibrated by comparing to a observed relation alone; we must instead fully calibrate such models. There is also no relation or quantity that would uniquely constrain jet velocities.

Variations of the coupling efficiency, ϵ_f (Section 4.3.3), and the accretion efficiency parameter, $r_{\text{tr},0}$ (Section 4.3.2), affect both galaxy and BH properties, while the BH seed mass, $M_{\text{BH,seed}}$, and the SN pivot density, $n_{\text{H,pivot}}$, affect mostly galaxy properties, with almost no effect on the high-mass end of the $M_{\text{BH}} - M_*$ relation or on the AGN LF (at $z \approx 0$). We thus first use the $M_{\text{BH}} - M_*$ relation and AGN LF to constrain the values of ϵ_f and $r_{\text{tr},0}$, respectively, and as a subsequent step use the remaining calibration observables (GSMF and galaxy half-mass size–stellar mass relation at $z = 0$) to calibrate $M_{\text{BH,seed}}$ and $n_{\text{H,pivot}}$.

The calibration of the coupling efficiency ϵ_f was done in the same way as for the thermal AGN feedback model in COLIBRE. Increasing the value of ϵ_f decreases the normalization of the $M_{\text{BH}} - M_*$ relation at the massive end ($M_* > 10^{10} M_\odot$), and vice versa, due to self-regulation (C. M. Booth & J. Schaye 2010). While changing ϵ_f affects galaxy properties in COLIBRE alongside BH masses, it is important to calibrate the coupling efficiency on the $M_{\text{BH}} - M_*$ relation first, and only then calibrate other parameters that directly affect galaxy properties (specifically the BH seed mass and stellar feedback pivot density).

We found that changing accretion efficiencies strongly affects the likelihood that BHs will be in a given accretion disc state (radiatively inefficient thick disc with $f_{\text{Edd},d} < 0.01$ or radiatively efficient thin disc with $f_{\text{Edd},d} > 0.01$), as well as changing galaxy properties, but it does not strongly affect BH masses (see Section 4.3.2). The accretion efficiency parameter $r_{\text{tr},0}$ (we remind the reader that $\epsilon_{\text{acc,th}} \propto f_{\text{Edd},d} / \sqrt{r_{\text{tr},0}}$ at $f_{\text{Edd},d} < 0.01$, see equation 18) can thus be calibrated largely independently of the coupling efficiency ϵ_f . Increasing the accretion efficiency decreases the normalization of the AGN LF, since more growth and feedback in the thick disc state (mostly jet feedback) makes it more likely that BHs will have $f_{\text{Edd},d} < 0.01$ and thus not contribute to the AGN LF. The opposite is true if the accretion efficiency is decreased.

Having chosen ϵ_f and $r_{\text{tr},0}$, the remaining step of choosing the seed mass $M_{\text{BH,seed}}$ and pivot density $n_{\text{H,pivot}}$ was relatively simple. We found that the seed mass needed to increase by roughly a factor of two (at all resolutions) in comparison with the thermal AGN feedback model in order for galaxy sizes at the massive end to match observations. The stellar feedback pivot density had to be reduced, decreasing the heating temperature used for SN feedback, thereby weakening it. This was done to counter the effects of a larger seed mass in the hybrid AGN model compared to the thermal one, which results in a relatively more important role of AGN feedback in low-mass and intermediate-mass galaxies ($M_* \lesssim 10^{10} M_\odot$).

4.5 Calibration results

Here, we discuss the final parameter choices for our hybrid AGN feedback simulations at m7, m6 and m5 resolution, as given in Table 3. The jet velocity scaling is the same at all resolutions, and is given by the parameter $v_{\text{jet},0} = 10^{4.5} \text{ km s}^{-1}$, yielding $v_{\text{jet}} \approx 31600 \sqrt{M_{\text{BH}}/10^9 M_\odot} \text{ km s}^{-1}$. Jet feedback is thus very explosive, with massive BHs launching particle pairs at a velocity slightly exceeding $0.1c$. This is not far from the observed values of kpc-scale jets, which are typically in the transrelativistic regime ($0.3\text{--}0.5c$, N. N. Jetha, M. J. Hardcastle & I. Sakelliou 2006; L. M. Mullin & M. J. Hardcastle 2009), or even slower (A. Shulevski et al. 2019). The energy associated with each of these kicks corresponds to an equivalent temperature increase of $\approx 2.4 \times 10^{10} \text{ K}$ (assuming a mean molecular weight of 1.2, appropriate for ionized hydrogen). This is ≈ 7.5 times higher than the maximum heating temperatures used for the thermal feedback channel (in both the thermal and hybrid AGN feedback models) at m7 ($\Delta T_{\text{AGN,max}} = 10^{9.5} \approx 3.1 \times 10^9 \text{ K}$) and 2.4 times higher than the maximum used at m6 and m5 ($\Delta T_{\text{AGN,max}} = 10^{10} \text{ K}$). Note also that this difference is even more stark if one considers an individual jet event to correspond to two particles, since we always kick in pairs. In this case, the total feedback AGN energy for each jet event is a factor ≈ 15 and ≈ 4.5 higher than in a thermal event for m7 and m6/m5, respectively.

Our initial choice for the parameter was $v_{\text{jet},0} = 10^4 \text{ km s}^{-1}$, which made jet feedback approximately a factor of two less energetic per feedback event than thermal AGN feedback, at fixed BH mass. This was motivated by a desire to resolve jets as well as possible, while ensuring they are still highly supersonic relative to the circumgalactic/ICM in their host haloes (jets need to be sufficiently supersonic, $v_{\text{jet}} \gtrsim 10c_s$, to inflate lobes, see F. Huško et al. 2024). However, we found it impossible to calibrate models to a satisfactory degree with such velocities. In particular, massive haloes ($M_{\text{halo}} > 10^{13} M_\odot$) hosted galaxies that were too massive, star-forming and small compared to observations at $z = 0$.¹⁵ We then attempted a value of $v_{\text{jet},0} = 10^{4.25} \text{ km s}^{-1}$, which mitigated but did not fully resolve the problem. With our current value, $v_{\text{jet},0} = 10^{4.5} \text{ km s}^{-1}$, galaxy sizes are in agreement with the data at m7 resolution, but possibly too small at m6 resolution at the high-mass end (see Section 4.6).

The fact that the hybrid AGN feedback model has to have such explosive AGN feedback can be explained by it featuring delayed BH growth in massive galaxies at high redshift ($z > 3$) (Huško et al. in preparation). Since heating temperatures and jet velocities scale with BH mass in COLIBRE, delayed BH growth implies less effective feedback for the same host galaxies at the same cosmic time. By making feedback more explosive (at fixed BH mass), we likely compensated for this effect. However, in hindsight, increasing the heating temperatures for the thermal AGN channel would probably have been at least as effective, if not more so, since quenching at high redshifts may be driven more by the thermal channel of feedback.

The highest values of jet velocities we use, $10^{4.5} \text{ km s}^{-1}$, are significantly higher compared to other cosmological galaxy formation simulations that use jets (S. Kaviraj et al. 2017; R. Davé et al.

¹⁵See the blue curve in the top right panel in Fig. 7; this disagreement with observational data could not be countered by e.g. increasing the BH seed mass.

Table 3. The final calibrated values of the parameters in the hybrid AGN feedback model (in parentheses we give the values of these parameters in the thermal AGN feedback model, if they exist). In order, these are 1) the jet velocity normalization (note that $v_{\text{jet}} \propto v_{\text{jet},0} \sqrt{M_{\text{BH}}}$, see equation 42), (2) the accretion efficiency parameter $r_{\text{tr},0}$ (note that $\epsilon_{\text{acc,th}} \propto f_{\text{Edd,d}} / \sqrt{r_{\text{tr},0}}$ for $f_{\text{Edd,d}} < 0.01$, see equation 18), (3) the thermal feedback coupling efficiency, (4) the BH seed mass, and (5) the SN feedback pivot density.

	(km s ⁻¹)	$r_{\text{tr},0}$	ϵ_f	$M_{\text{BH,seed}}(M_{\odot})$	$n_{\text{H,pivot}}(\text{cm}^{-3})$
m7	10 ^{4.5}	10 ⁴	0.03 (0.1)	5 × 10 ⁵ (3.16 × 10 ⁵)	1.5 (0.6)
m6	10 ^{4.5}	10 ⁴	0.02 (0.05)	5.62 × 10 ⁴ (3 × 10 ⁴)	0.75 (0.5)
m5	10 ^{4.5}	10 ⁴	0.02 (0.05)	4 × 10 ⁴ (2 × 10 ⁴)	1.2 (1.0)

2019; Y. Dubois et al. 2021), which use values up to 10⁴ km s⁻¹. However, our maximum jet velocities are only reached at high BH masses, $M_{\text{BH}} \geq 10^9 M_{\odot}$. Galaxies with such high BH masses have already gone through their strongest gas outflow phases and been quenched, which occurs when their progenitor galaxies have a mass of $M \approx 10^{10.5} M_{\odot}$ and their BHs a mass of $M_{\text{BH}} \approx 10^8 M_{\odot}$. These galaxies have jet velocities more similar to 10⁴ km s⁻¹, used by other galaxy formation simulations. It should be noted, however, that at high redshifts, AGN jets in our simulations are predominantly from the thin disc state at high Eddington ratios (see Section 6), unlike other simulations.

Our final choice for the accretion efficiency parameter was $r_{\text{tr},0} = 10^4$, which yields the scaling (see Section 3.3.1, equation 18) $\epsilon_{\text{acc,th}} = 0.0316(f_{\text{Edd,d}}/0.01)$ for $f_{\text{Edd,d}} < 0.01$. We find this same scaling to work at all three resolutions: m7, m6, and m5. As is visible, the accretion efficiencies are of order one per cent or lower; this is necessary to significantly dampen the strength of jets from the thick disc state at $f_{\text{Edd,d}} < 0.01$. Our calibrated coupling efficiency is equal to $\epsilon_f = 0.03$ at m7 and slightly lower, $\epsilon_f = 0.02$, at m6 and m5. This is a factor of 3 (2.5) times lower than the coupling efficiencies used in the thermal AGN feedback model (0.1 and 0.05).

Our final BH seed masses are approximately two times larger in the hybrid AGN feedback model than in the thermal AGN feedback model at all resolutions. These values are necessary in order to reproduce the stellar masses and sizes of galaxies in massive haloes ($M_{\text{halo}} > 10^{13} M_{\odot}$). This implies that the hybrid AGN feedback model is slightly less effective at quenching these massive galaxies, given the same BH seed mass. Our final pivot densities are higher than the values used in the thermal AGN feedback model at all resolutions. This implies that SN feedback is weaker, which is necessary to reproduce stellar masses and sizes of galaxies around the peak of the stellar mass–halo mass relation ($M_{\text{halo}} \approx 10^{12} M_{\odot}$). We had to weaken SN feedback (relative to the thermal AGN feedback model) since AGN feedback is stronger in that mass regime due to the higher seed mass.

4.6 Convergence

In Fig. 10, we show the calibration relations for our biggest simulations at each resolution that were run with the hybrid AGN feedback model and that have reached $z = 0$ (L200m7h, L100m6h, and L025m5h, with L100m7h also included for a box-size comparison), each having been calibrated independently. These are to be compared with observations and with each other. The figure has the same format as Figs 7–9, but we extend the axes ranges somewhat, since we are also interested in testing the convergence with resolution. Since the simulations are (re)calibrated at higher resolution, this constitutes a weak convergence test in the terminology of J. Schaye et al. (2015). For clarity, we do not

show the comparison with the thermal AGN feedback model (see Figs 25 and 26 in J. Schaye et al. 2025 for that comparison). We note here, however, that the main observables shown here are reproduced to a similar degree of accuracy as with the thermal AGN feedback model. This means that the abundance of galaxies and their sizes, as well as the BH masses, are similar between the two models. This agreement between the two AGN feedback implementations also holds for many other predicted galaxy scaling relations that were not considered during the calibration. Both the thermal and hybrid AGN models at the m5 resolution perform slightly worse than their m6 and m7 counterparts, which is an outcome of more crude calibration due to wall-clock time constraints of running m5 simulations.

4.6.1 Galaxy stellar mass function

In the top left of Fig. 10, we compare the simulated and observed GSMFs. We find that the model predicts similar GSMFs for m7 and m6 at all masses where the simulation predictions are reliable for both resolutions ($M_* \gtrsim 10^9 M_{\odot}$), with the m6 model having a slightly higher GSMF at $10^9 < M_*/M_{\odot} < 10^{10}$. All three GSMFs are in fairly good agreement with the observational data at all stellar masses, within 0.2 dex (this is comparable to the systematic uncertainty in the observationally inferred stellar masses). The simulated GSMF is ≈ 0.15 dex higher at the knee of the GSMF ($M_* \approx 10^{10} M_{\odot}$) than the data, and ≈ 0.2 dex lower than the data at $M_* \approx 10^9 M_{\odot}$ for m7 and m6, but in good agreement with the data at m5. At lower masses, the simulated GSMF is in agreement with the increasingly uncertain observational data, even down to $M_* = 10^7 M_{\odot}$, where the m7 resolution is not reliable. The simulated GSMF is higher than the last data point for the most massive galaxies from S. P. Driver et al. (2022) ($M_* \approx 10^{11.6} M_{\odot}$), but the comparison is not reliable here (see Section 4.2).

4.6.2 Galaxy size–mass relation

The top right panel of Fig. 10 compares simulated galaxy sizes to observations. Both m7 and m6 are generally in good agreement with the data at most stellar masses (within 1σ , where σ is the standard error of the median, given by the error bars), while m5 shows slightly too large sizes (at $M_* \approx 10^{10} M_{\odot}$), with some of the disagreement stemming from the L025m5h box not being volume-converged. The m7 model shows signs of disagreement with the data at $M_* = 10^9 M_{\odot}$, but the sizes are affected by resolution effects at these masses. The m6 model slightly underestimates the observed sizes at the high-mass $M_* > 10^{11} M_{\odot}$ end (at a level of more than 2σ), by ≈ 0.1 dex. However, the size–mass relation is very steep at these masses. A shift in stellar masses by 0.2 dex could alleviate this tension (by moving the entire relation horizontally), and such a shift would be similar in magnitude to

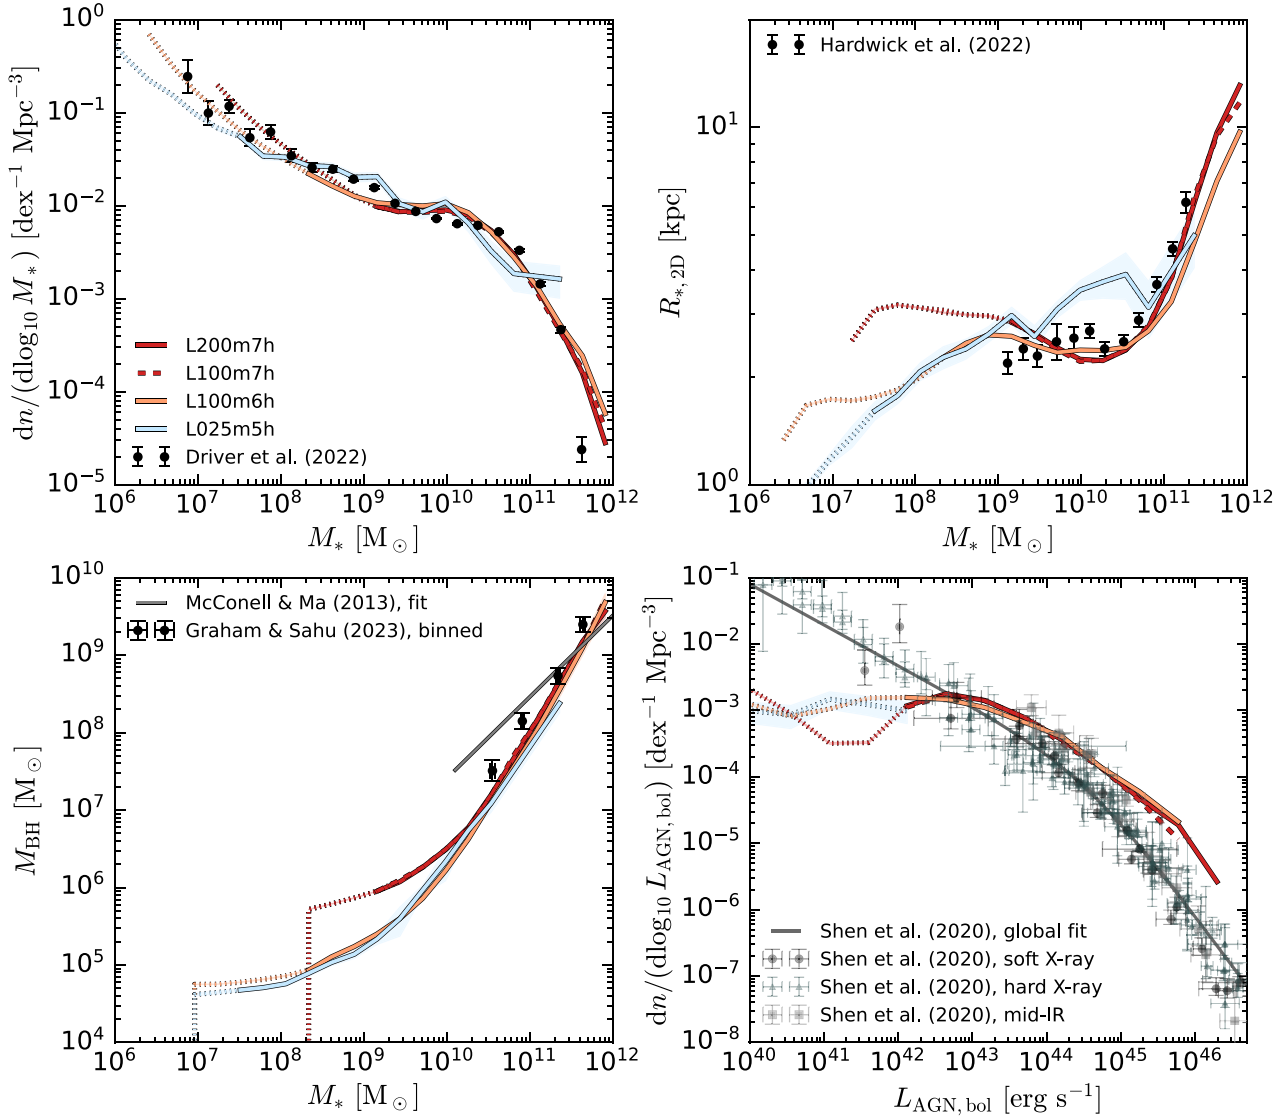


Figure 10. The calibration observables for our final calibrated models at m7, m6 and m5 resolutions in the largest simulated volumes that have reached $z = 0$ (L200, L100, and L025, respectively). The simulations were calibrated independently at each resolution. The panels that have stellar mass on the x -axis extend down to $M_* = 10^6 M_\odot$, but we stop the curves where we resolve a galaxy with a single stellar particle ($\approx 10^7$, $\approx 10^6$, and $\approx 10^5 M_\odot$, respectively), and make them dotted where the simulation predictions are unreliable (at a stellar mass corresponding to ≈ 100 particles for each resolution in the first 3 panels, and at an AGN luminosity of $L_{\text{bol}} = 10^{42} \text{ erg s}^{-1}$ for the last panel). At the high-mass (or luminosity) end we stop the curves at the first bin that has fewer than 10 objects in it. For other details of the comparison, see the caption of Fig. 7.

the systematic uncertainties in observational determinations of stellar masses.

4.6.3 Black hole mass–stellar mass relation

The bottom left panel of Fig. 10 compares BH masses with observations. The BH masses are similar at all three resolutions for $M_* > 10^{10} M_\odot$, and in close agreement (within the 1σ uncertainty) with the data from A. W. Graham & N. Sahu (2023). The simulated $M_{\text{BH}}-M_*$ relation has the same slope as the relation from A. W. Graham & N. Sahu (2023), which is much steeper than the power-law fit from N. J. McConnell & C.-P. Ma (2013). This is because that fit is more appropriate for spheroid-dominated galaxies, which have higher BH masses at fixed stellar mass. Below $M_* \approx 10^{10} M_\odot$, the simulation predictions begin to

diverge. At low stellar masses, the BH mass tends smoothly to the value of the seed mass, $M_{\text{BH, seed}}$. These values correspond to $5 \times 10^5 M_\odot$, $\approx 5.63 \times 10^4 M_\odot$, and $4 \times 10^4 M_\odot$ at m7, m6, and m5, respectively.

4.6.4 Bolometric AGN luminosity function

In the bottom right panel of Fig. 10, we compare the AGN bolometric LF to the observational data. We find that the results are well converged above $L_{\text{AGN, bol}} = 10^{42} \text{ erg s}^{-1}$, the value below which the observational comparison is not reliable (Section 4.2). At values below this, the AGN LF for m6 and m5 is significantly higher than m7, and closer to the observed data. At $L_{\text{AGN, bol}} > 10^{43} \text{ erg s}^{-1}$, the predicted AGN LF is 0.3–0.7 dex too high compared with the data. This disagreement implies that our adopted

accretion efficiencies in the thick disc state ($f_{\text{Edd},d} < 0.01$) are too low. The disagreement is caused by two effects. First, for practical reasons, during calibration we almost always plotted our predictions at $z = 0$ (we often did not output snapshots at redshifts other than $z = 0$ and $z = 2$ during calibration), as for most observational comparisons, the data is at $z \approx 0$. However, the data for the AGN bolometric LF is at a median redshift $z \approx 0.2$. Mean AGN luminosities grow by nearly 0.3 dex between these redshifts. Secondly, we mostly used L025 or L050 boxes during calibration, and the AGN LF is not fully converged in terms of volume for these boxes. Despite these mild disagreements with the data, our AGN LF is still in better agreement with it than had we not used sub-unity accretion efficiencies (Fig. 8), and this was our main motivation to calibrate on the AGN LF.

4.7 The role of AGN feedback in low-mass galaxies

As shown in Section 4.3 (see also Appendix I), the hybrid AGN feedback simulations feature significant AGN feedback effects in low-mass galaxies. There is no strong observational evidence of AGN feedback operation in low-mass galaxies (e.g. A. C. Fabian 2012), but some evidence in selected low-mass galaxies has begun to emerge in recent years (e.g. S. J. Penny et al. 2018; D. S. Carr et al. 2025; S. Salehirad, A. E. Reines & M. Molina 2025). Around half of the injected AGN energy in the hybrid AGN feedback model is in the form of jets (see Section 6), which could explain the increased role of AGN feedback in the hybrid model in low-mass galaxies compared to the purely thermal one. This interpretation could be consistent with observations that find evidence of jets in low-mass galaxies (M. Mezcua, H. Suh & F. Civano 2019). Alternatively, the stronger effects of AGN feedback in the hybrid simulations could be due to the use of a higher BH seed mass, which results in overall higher BH masses at fixed halo mass, and therefore stronger feedback. This is consistent with previous simulation studies that find the mass of the BH to be the primary driver of quenching at lower BH masses (e.g. R. S. Sharma et al. 2020; S. Koudmani, N. A. Henden & D. Sijacki 2021; S. Wellons et al. 2023; E. Arjona-Gálvez, A. Di Cintio & R. J. J. Grand 2024). C. M. Booth & J. Schaye (2013) and S. Koudmani, D. Sijacki & M. C. Smith (2022) have shown that a reduced efficiency of SN feedback can also make AGN feedback stronger by allowing more BH growth; this is also consistent with our findings, since the hybrid AGN feedback simulations use lower SN feedback heating temperatures by necessity.

5 BLACK HOLE SPIN–MASS RELATION

In this work we largely refrain from providing predictions on galaxy and BH properties from simulations run with the hybrid AGN feedback model. Some of these are presented in J. Schaye et al. (2025) and others will be explored in future studies. However, we present in Fig. 11 the predicted BH spin–BH mass relation at $z = 0$ (note that this shows the spin magnitude) for the L200m7h, L100m7h, L100m6h, and L025m5h simulations, which use the final calibrated hybrid AGN feedback model presented in Section 4. BH spin is a new and important BH property introduced in the hybrid AGN model, and Fig. 11 illustrates the typical magnitudes of BH spin for BHs of different masses.

From the left panel, which shows the relation for all BHs, we see that median BH spins are generally high, $|a| > 0.8$, for intermediate-mass BHs ($M_{\text{BH}} \in [10^6, 10^8] M_{\odot}$) at m7 and m6, with peak median values of $|a| = 0.96$ and $|a| = 0.92$, reached at

a BH mass of $M_{\text{BH}} = 10^{6.5} M_{\odot}$ at both resolutions. For BH masses outside the range $M_{\text{BH}} \in [10^6, 10^8] M_{\odot}$, the magnitude of BH spin decreases towards even lower or even higher BH masses on either side of this range. At m5, the BH spin grows monotonically up to $M_{\text{BH}} \approx 10^8 M_{\odot}$, the highest BH mass at which we can constrain BH spin using a L025 volume, where it reaches a value $|a| \approx 0.6$, which is significantly lower than at m7 and m6. For a discussion of the reasons behind this lack of convergence, see below.

The shape of the BH spin–mass relation is a result of a balance between gas accretion, BH–BH mergers and jet-induced spin-down. At intermediate masses ($M_{\text{BH}} \in [10^6, 10^8] M_{\odot}$), BHs are most likely to be in the thin disc state (see Fig. 3, which shows that BHs with $M_{\text{BH}} = 10^7 M_{\odot}$ are the ones most likely to be in the thin disc state, and this is close to the BH mass at which the median spin magnitude is highest). Spin-up in the thin disc state is very efficient due to high gas accretion rates, very quick spin redirections, and negligible jet-induced spin-down. Mergers tend to result in BH spins significantly smaller than $|a| = 1$ (e.g. $|a| \approx 0.7$, E. Barausse et al. 2012), which can explain lower BH spins for massive BHs ($M_{\text{BH}} > 10^8 M_{\odot}$). These BHs are hosted by quenched, gas-poor galaxies whose growth is mainly driven by mergers. Jet-induced spin-down due to prolonged accretion in the thick disc state may also play a role in driving the median BH spin to values as low $|a| \approx 0.3$ at $M_{\text{BH}} > 10^{9.5} M_{\odot}$, which is hard to achieve through mergers alone. BH seeds in our simulations mostly grow through mergers. However, since we seed BHs with $|a| \approx 0$, these initial mergers are unlikely to lead to significant spin-up. The increase in BH spin with BH mass at the low-mass end ($M_{\text{BH}} < 10^6 M_{\odot}$), visible in Fig. 11, is thus likely to be due to gas accretion. Since the BH seed mass is a calibration parameter that changes with resolution, our prediction for the BH spin–mass relation is likely resolution-dependent at these low BH masses, as is visible in the figure. However, the large difference between the seed mass at m6 and m7 is mostly due to the two resolutions using a different halo threshold mass for BH seeding.

In the right panel of Fig. 11, we compare the predicted BH spin magnitudes with observational measurements of BH spin. These measurements rely on X-ray reflection spectroscopy and thermal continuum fitting (see the review by C. S. Reynolds 2021). They should be treated with care, since they rely on uncertain models (e.g. S. Hagen & C. Done 2023). In addition to the compilation of BH spin measurements provided by C. S. Reynolds (2021), which span a wide range of BH masses, we also include measurements for low-mass BHs ($M_{\text{BH}} \approx 10^6 M_{\odot}$) from L. Mallick et al. (2022), who used X-ray reflection spectroscopy. We supplement these datasets with our own compilation of individual measurements from the following studies, using largely but not exclusively X-ray reflection spectroscopy: D. J. Walton et al. (2013), M. J. Valtonen et al. (2016), R. V. Vasudevan et al. (2016), Event Horizon Telescope Collaboration et al. (2019), J. Jiang et al. (2019), S. Laine et al. (2020), C. Bambi et al. (2021), R. Ghosh & S. Laha (2021), D. J. Walton et al. (2021), and J. Sisk-Reynés et al. (2022). Many of these observational measurements are lower limits rather than true measurements of BH spin, so we do not bin the data.

All of the above spin measurements are for BHs accreting at moderate rates (C. S. Reynolds 2021). For this reason, we restrict our comparison to BHs in the thin disc state ($0.01 < f_{\text{Edd},d} < 1$), where the AGN spectrum can be distinguished from that of the host galaxy. The inferred BH masses in the observed sample have a median uncertainty of ≈ 0.4 dex. We therefore add 0.4 dex lognormal scatter to our simulated BH masses to replicate Ed-

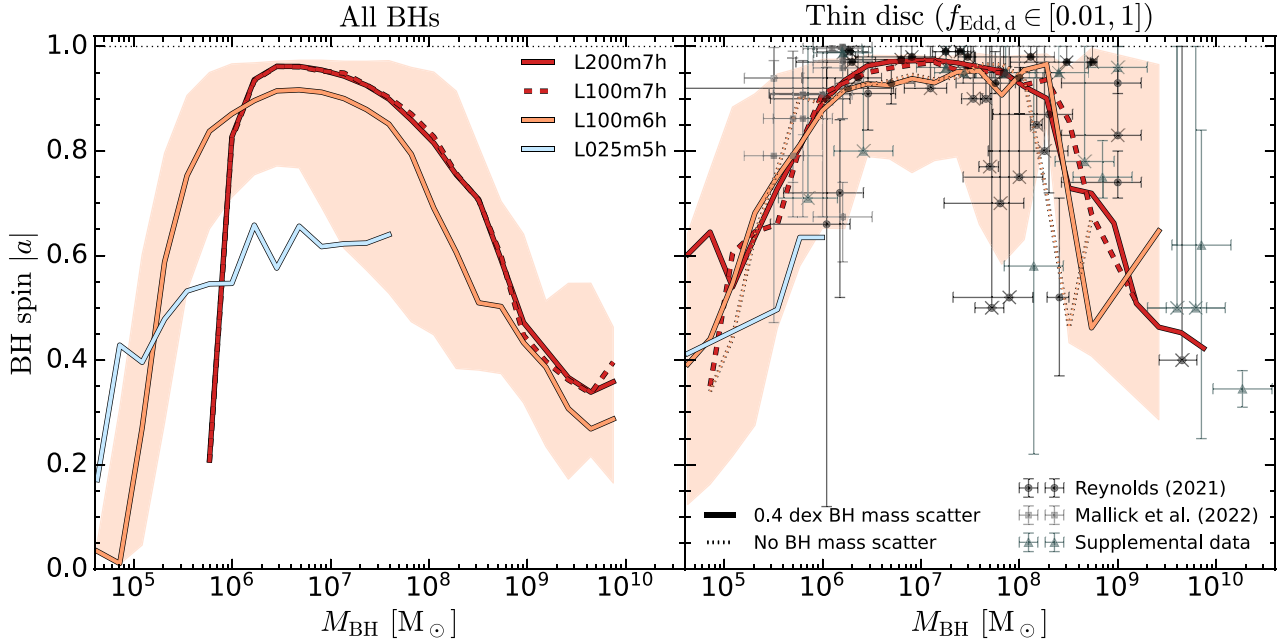


Figure 11. The BH spin–mass relation for the L200m7h (red solid), L100m7h (red dashed), L100m6h (orange solid), and L025m5h (light blue solid) simulations. The orange shaded region shows the 16–84 percentile range for L100m6h. The left panel shows the relation for all BHs, while the right panel shows it only for BHs in the thin disc state, where we have also added 0.4 dex lognormal scatter to BH masses. The exception in this panel is the dotted line, which shows the case of no added scatter for L100m6h. The spins of BHs in the thin disc state are compared with observations of moderately accreting BHs (C. S. Reynolds 2021 and L. Mallick et al. 2022, as well as our compilation of other studies; see the text for further details). For these observed BHs, many measured spins are lower limits. We mark these measurements with crosses, show no lower error bar and extend the upper one to $|a| = 1$.

dington bias (for reference, we also show the case with no scatter added using a dotted line for L100m6h—the difference between the two cases is small). Comparing the simulation lines in the right panel with those in the left panel of Fig. 11, the latter of which are for all BHs, we see that applying this selection results in higher BH spins at fixed BH mass, at least at intermediate BH masses ($M_{\text{BH}} \in [10^6, 10^8] M_{\odot}$). This is indicative of those BHs spinning up more efficiently, which is expected for the thin disc state.

Our predictions for both m7 and m6 resolution are broadly consistent with the observed trend of the BH spin–mass relation, while m5 seems to predict too low BH spins for intermediate-mass BHs. At a mass of $M_{\text{BH}} \approx 10^7 M_{\odot}$, where the relation peaks, the predicted spin magnitudes are in agreement with the observations at m7 and m6, all of which are close to $|a| = 1$. At lower BH masses, our predictions are also close to the observations. The scatter is reproduced there, while the low-spin percentile range of the scatter is possibly too large at $M_{\text{BH}} \approx 10^7 M_{\odot}$. We predict massive BHs ($M_{\text{BH}} > 10^{8.5} M_{\odot}$) to have slightly lower BH spin than the dozen or so observed ones in this mass range, at least at m7. If this is not a result of observational bias, and is instead true disagreement, it could be indicative of overly strong jet spin-down of massive BHs, since BHs cannot otherwise obtain low spins. However, at m6, we predict BH spin to be somewhat higher at the high-mass end ($M_{\text{BH}} > 10^8 M_{\odot}$) than at m7, at least up to where predictions are available at m6 from the L100m6h simulation ($M_{\text{BH}} \approx 10^{8.5} M_{\odot}$). Larger volumes are needed to probe higher masses at m6, in order to determine whether the hybrid AGN feedback model predicts correct BH spins in massive galaxies. Finally, we do not compare in detail

with other simulations, but our predictions for the relation are broadly similar to those of other recent cosmological hydrodynamical simulations that include BH spin modelling (S. Bustamante & V. Springel 2019; L. Sala et al. 2024; R. S. Beckmann et al. 2025).

We stress that BH spins were not used during the calibration of the hybrid AGN model, and are thus true predictions. However, during calibration we found that the prediction for BH spins is intimately tied to the radiative activity of BHs. In particular, variations of the accretion efficiency of the thick disc, which strongly influence the AGN LF (Section 4.3.2), also had a large impact on the predicted BH spin–mass relation. Larger values of the accretion efficiency result in lower BH spins. This is a consequence of BHs more frequently being in the thick disc rather than in the thin disc state near $z = 0$, which results in less efficient spin-up. Interestingly, we found that models that do not match the AGN LF also do not match the observed BH spin–mass relation.

The good agreement of predicted BH spins at m7 and m6 resolutions, as well as that of S. Bustamante & V. Springel (2019) and L. Sala et al. (2024), all of which are relatively low-resolution ($m_g > 10^6 M_{\odot}$), is intriguing, and somewhat unexpected, especially given the too low BH spins at m5 resolution. This can be explained as follows. Efficient spin-up of BHs requires: 1) that a sufficient fraction of the mass accretion should occur in the thin disc state, where spin-up is easy, and 2) good alignment of BH spin with the surrounding gas on resolved scales, i.e. the BH spin, a , needs to be positive often enough. We claim that the first condition is fulfilled by matching the observed AGN LF. The second condition is easy to fulfill at low resolutions.

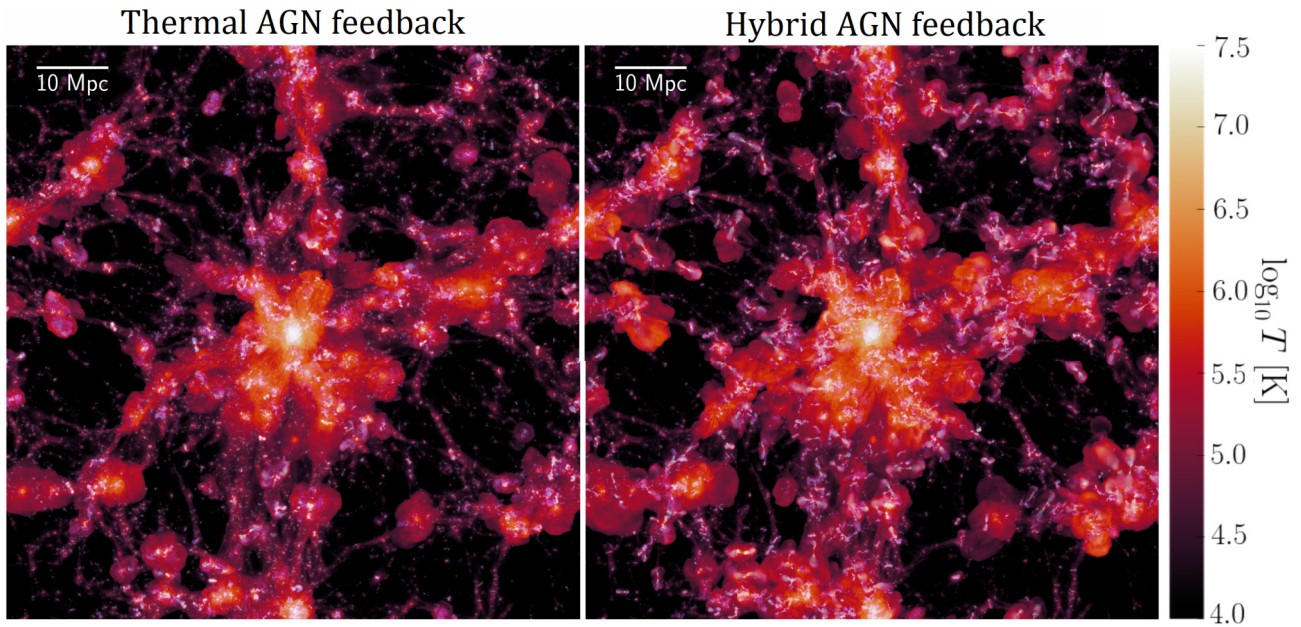


Figure 12. Temperature in slices of 30 cMpc at $z = 0.2$ through the L100m6 (left) and L100m6h (right) simulations, using thermal and hybrid AGN feedback, respectively. The images are in proper coordinates. Red to yellow colours encode the mass-weighted projected temperature, while blue to white colours highlight particles affected by thermal or jet feedback up to 5 Gyr in lookback time. For these particles, the surface density is used to determine the opacity, while brighter (white) colours indicate more recent AGN events.

Low-resolution simulations, even with a multiphase ISM, do not resolve the turbulent nature of the gas as well as high-resolution simulations. Furthermore, BH kernels are larger at low resolution, encompassing a larger fraction of the galaxy and thus probing better the average angular momentum of the gas, which is well-defined for disc galaxies. This results in the angular momentum of the gas accreting on to a subgrid BH being aligned with the large-scale angular momentum of gas in the galaxy disc than it might be at higher resolution. These considerations imply that spin-up should be less efficient at higher resolutions, which we can indeed see from Fig. 11. We interpret this as a result of gas turbulence being resolved better and BH kernels being smaller.

These considerations show that predicting BH spins is a bigger challenge in high-resolution simulations than low-resolution ones, in a model that assumes that the subgrid accretion disc is aligned with the direction of the angular momentum in the BH kernel on resolved scales of the simulation. This strong assumption works in low-resolution simulations, where turbulence is not resolved well and where BH kernels encompass a substantial fraction of the galaxy. The assumption breaks down at higher resolutions, where the subgrid accretion disc is not resolved, nor angular momentum flows on to it, while the gas turbulence is resolved and the BH kernels are small. In high-resolution simulations, the angular momentum of gas accreting on to the BH is stochastic and often misaligned with respect to a galaxy disc (R. Levine, N. Y. Gnedin & A. J. S. Hamilton 2010; P. F. Hopkins et al. 2012; D. Anglés-Alcázar et al. 2017). This is also consistent with observations that show nuclear molecular gas discs that are misaligned with respect to their host galaxies (J. E. Greene et al. 2013; I. Ruffa et al. 2020). On the other hand, observations of radio jets indicate broad alignment with galaxy discs (S. Kaviraj et al. 2015; L. A. Vazquez Najjar & H. Andernach 2019; X. Zheng, Y. Zhang & H. Röttgering 2024), which may imply alignment in a time-averaged sense.

6 AGN FEEDBACK IN THE HYBRID SIMULATIONS

In this section, we analyse in more detail how AGN feedback operates in practice in the COLIBRE simulations. We focus on jet feedback and we only consider basic properties; more detailed analyses will be presented in future studies. In Section 6.1, we highlight the qualitative differences between the thermal and hybrid AGN feedback simulations using visualizations of the IGM and large-scale AGN activity. We also quantify the far-reaching effects of jets. In Section 6.2, we compare the feedback energies in the hybrid and thermal AGN feedback simulations, focusing on how much energy is injected as a function of cosmic time. For the hybrid AGN feedback simulations, we discuss the relative importance of the two feedback channels and the three accretion disc states.

6.1 Large-scale outflows

In Fig. 12, we show the temperature slices through the L100m6 and L100m6h simulations (run with the thermal and hybrid AGN feedback models, respectively) at $z = 0.2$. We also overlay particles directly heated or kicked by AGN feedback (see caption for details). AGN feedback drives outflows and heats the IGM through shocks. These are frequently bipolar even in the thermal AGN feedback case (which is nominally isotropic, but bipolar outflows tend to form as gas escapes galaxies; P. D. Mitchell et al. 2020), but more often so in the hybrid AGN feedback case, due to the explicit two-sided jet activity. Individual jets may be discerned in the right panels as well as the bow shocks ahead of the jets. The jets appear to push out gas to larger radii compared to thermal feedback. This is consistent with large-scale jet activity in the SIMBA simulations (R. Davé et al. 2019; J. Borrow, D. Anglés-Alcázar & R. Davé 2020).

We now quantify some of our discussion from above. When kicking particles into jets, we tag them by the ID of the BH that kicked those particles. We do not do this for thermal feedback, so we restrict the following analysis to jet feedback only. In addition to recording the ID of the BH particle, we record how much energy was given to the gas particles, and the time when the kick occurred. This can be used to track individual jet particles and their positions at any given time. Using this information, for every central halo in L200m7h at $z = 0$, we find all particles that were kicked as part of jet feedback by the current most massive BH (in a 50 kpc 3D aperture centred on the most bound particle) hosted by that halo. We then compute the distance to both the centre of mass and centre of energy of these particles, where for the latter, the weight is $m_g v_{\text{jet}}^2$, v_{jet} being the target jet velocity at kick time. We compute mean redshifts at kick time for the centre of energy case. For every BH, we also find the farthest kicked particle.

Fig. 13 shows the results of this analysis, as a function of halo mass, defining the haloes here as objects whose mean density is 200 times larger than the critical density at $z = 0$. In the top panel, we show the raw distances, while the bottom panel shows distances normalized by the virial radius. The median jet distances increase with halo mass from ≈ 0.5 Mpc (≈ 6 virial radii) at $M_{200,c} = 10^{11} M_\odot$ to ≈ 4 Mpc (≈ 10 virial radii) at $M_{200,c} \approx 10^{13} M_\odot$. They then drop off, both in raw terms and when normalized by virial radius. This is true for both the centre of mass and centre of energy cases. The median jet distance for the most massive haloes in L200m7h, with $M_{200,c} \approx 10^{15} M_\odot$, is 2 Mpc (≈ 1 virial radius).

This decrease in the median jet distance is likely caused by two effects. First, AGN activity is most efficient at expelling gas around the $M_{200,c} \approx 10^{13} M_\odot$ mass scale (see the scaling of gas fraction with halo mass; Davies et al. in preparation). For these haloes, jet particles are therefore most likely to end up at many virial radii away. Secondly, more massive haloes have recurrent jet activity at lower redshifts, which can be seen in Fig. 13 in terms of the mean kick redshift for the most massive haloes ($z \approx 1.5$ for $M_{200,c} \approx 10^{15} M_\odot$ compared to $z \approx 3$ for $M_{200,c} \approx 10^{11} M_\odot$). These jet particles kicked at lower redshifts are less likely to be located as far away as the ones kicked at high redshift.

In Fig. 13, we also show the median distance to the farthest kicked jet particle. This saturates to 10 Mpc, which corresponds to ≈ 20 virial radii for $M_{200,c} = 10^{13} M_\odot$ and ≈ 5 virial radii for $M_{200,c} = 10^{15} M_\odot$. While not shown here, in the case of individual haloes, the farthest jet particles can be located up to 20 Mpc away (largely at $M_{200,c} \leq 10^{13}$), which is, incidentally, similar to SIMBA (J. Borrow et al. 2020; M. Gebhardt et al. 2024), but less extreme than IllustrisTNG (M. Ayromlou, D. Nelson & A. Pillepich 2023). Note that even the centre of energy can be as far away as ≈ 100 virial radii for low-mass haloes ($M_{200,c} \approx 10^{11} M_\odot$).

We warn the reader that the above discussion does not necessarily pertain to the actual size of currently active jets, nor even jet remnants. This is for the following two reasons. First, BHs may kick particles that exit haloes hosting them and then effectively detach from the halo's motion. Even if we are not considering satellites here, the haloes may move considerably between the time at kick and $z = 0$, and thus the farthest jet particles may be located far away simply due to this motion. Secondly, much of the jet feedback occurs at high redshift, especially for high-mass BHs, hosted by high-mass haloes, that do jet feedback in the thin disc state (see the next subsection). The peak redshift of this activity is at $z \approx 3$. The fact that BHs hosted by more massive haloes kick

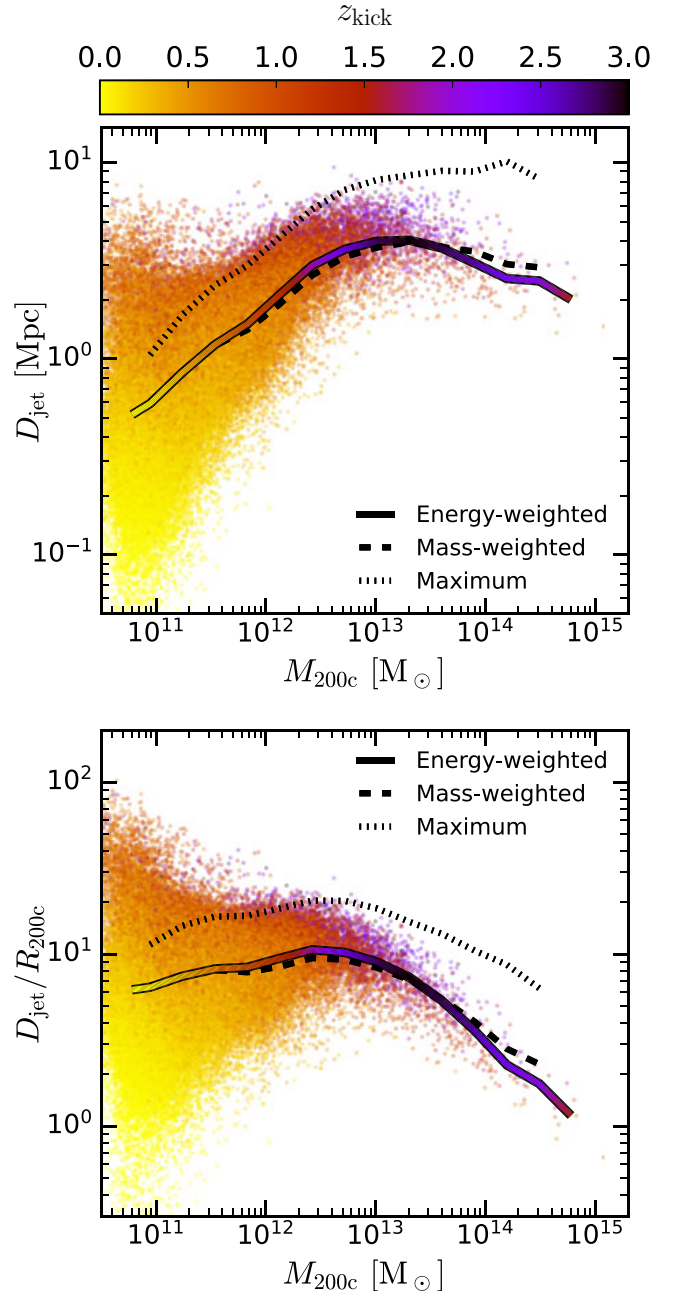


Figure 13. The distances out to which a given halo can affect its surroundings through jet feedback in the L200m7h simulation. We show at $z = 0$ and for central haloes, as a function of halo mass, three distance scales (top: raw scale in physical units, bottom: the same, normalized by the virial radius). The solid line shows the median distance to the centre of energy of jet-kicked particles, where the weight for every particle is $m_g v_{\text{jet}}^2/2$. The coloured dots show these median centres of energy for individual haloes. Colours represent the mean redshift at kick time, weighted in the same way as the distance. Dashed and dotted lines show median distances computed using the centre of mass and the farthest kicked jet particle, respectively. All distances are computed with respect to the halo centre (the most bound particle), and accounting for periodic boundary conditions. The BH used in the analysis is the most massive bound BH located in a 50 kpc exclusive aperture.

more particles at higher redshifts, as seen in Fig. 13, may be an effect of jet feedback being stronger in the thin disc state, but it may also simply be a consequence of more massive haloes having formed earlier. In either case, these jet particles have departed their haloes at relatively early redshifts. Such particles join the Hubble flow and are therefore located even farther from their host haloes than they would otherwise be (compared to a simple scenario of ballistic ejection from a self-gravitating halo, without cosmic expansion).

6.2 Feedback energetics

Here, we focus on injected AGN energies, comparing the thermal AGN feedback simulations to those using the hybrid model, as well as comparing different channels and modes in the hybrid simulations. We also compare with injected SN energies. We study the energetics both globally as a function of redshift, and in terms of cumulative energies as a function of BH mass at $z = 0$. We use the L200m7 and L200m7h simulations since they are converged (volume-wise) in terms of these energy measures, exhibit the smallest variation with time between snapshots, and are the largest simulations using the hybrid model that we have available, in terms of volume. Some of the quantitative analysis in the following discussion depends slightly on resolution, but the qualitative trends that we find also apply for the L100m6 simulations, which we do not show here.

The left column of Fig. 14 shows the energetics in terms of the cosmic evolution of feedback power densities (top row) and cumulative injected feedback energy densities (bottom row). In both the thermal and hybrid AGN feedback simulations, SN feedback dominates above $z \approx 1.5$, after which AGN take over (see top left). At $z = 0$, AGN are injecting 3 times as much energy as SN in the thermal AGN feedback simulation, and an order of magnitude more in the hybrid AGN feedback simulation. By $z = 0$ the total injected AGN energy (bottom left) is nearly equal to the total injected SN energy in the thermal AGN feedback simulation, and equal for the hybrid AGN case.

From the left panels, we can also see that the AGN are injecting similar amounts of energy in the two AGN models up to $z \approx 5$, with the thermal model showing slightly more AGN energy at these early redshifts. This is related to the thermal AGN feedback simulations exhibiting more efficient BH growth at the earliest redshifts (Huško et al. in preparation). AGN start injecting more energy (by ≈ 50 per cent) in the hybrid model compared to the thermal one below $z = 5$ and overall around 25 per cent more energy is injected in the hybrid model by $z = 0$.

The left panels also show the hybrid AGN power and energy density components due to thermal isotropic and kinetic jet feedback. Jet feedback is dominant at low redshift $z \leq 0.5$. This is in agreement with the canonical picture of radiatively driven ‘quasar’ feedback dominating at high redshifts, and mechanical feedback dominating at low redshifts. Thermal feedback is, however, non-negligible even at $z = 0$. Finally, jet feedback is again more important than thermal feedback at $z > 2$. By $z = 1$, both feedback modes have injected nearly equal amounts of energy, and this remains the case down to $z = 0$. Our findings are in good agreement with the observationally motivated empirical model of P. Mocz, A. C. Fabian & K. M. Blundell (2013).

In the middle column of Fig. 14, we show the individual contributions to the two feedback modes (thermal and jet) from the three different accretion disc states in the hybrid AGN simulations. In general, three feedback channels are important: (1) thick

disc jets (from BHs with $f_{\text{Edd},d} < 0.01$) at low redshift ($z < 1$), (2) thin disc thermal feedback (from BHs with $0.01 < f_{\text{Edd},d} < 1$) at all redshifts, and (3) thin disc jets at $z > 1$. The dominance of jets over thermal feedback at low ($z < 0.5$) and high redshifts ($z > 2$) is thus due to different types of discs. By $z = 0$, thin disc jets have injected a very similar amount of energy as thin disc thermal feedback. The total thick disc jet energy injected by $z = 0$ is a factor of ≈ 2 lower than the thin disc jet and thermal components. Winds from the thick and slim discs ($f_{\text{Edd},d} > 1$ for the latter), implemented using thermal feedback in the hybrid model, are generally negligible in terms of global energetics with the exception of slim disc winds at $z > 5$, when they are energetically comparable to thin disc jets and thermal feedback. Jets from the slim disc are always negligible, although we suspect that this may not be the case in a more realistic model.¹⁶

In the right column of Fig. 14, we show the effective feedback efficiencies in the same simulations. These are computed by dividing the power (energy) densities by the rest mass-energy accretion rate densities (rest accreted mass-energy densities) on to all BHs. By construction, the thermal run has a total efficiency of $\epsilon = \epsilon_{\text{t}} \epsilon_{\text{rad}} = 0.01$. Note that jet efficiencies here do not have the same meaning as the ‘true’ jet efficiencies defined in Section 3, since we are here dividing the energies that may be injected in any accretion state with masses that may have been accreted in other accretion states. Most BH accretion occurs in our model at moderate Eddington ratios, $f_{\text{Edd},d} = 0.01 - 0.3$, while most jet energy is launched at Eddington ratios below (in the thick disc state) or above those values, with little growth occurring at such Eddington ratios, largely due to low values of the accretion efficiency.

The total instantaneous efficiency in the hybrid AGN feedback simulation is slightly lower (by ≈ 0.1 dex) than in the thermal AGN feedback simulations at $z < 1$, despite it having a larger power density at these redshifts. The total instantaneous efficiency is larger in the hybrid AGN feedback simulations at $z > 2$, and reaches a maximum of ≈ 3 per cent at $z = 5$. The jet channel has a mean instantaneous efficiency of 0.3-0.6 per cent at $z < 2$ (increasing to ≈ 2 per cent at $z = 5$, likely due to thin disc jets becoming important). The thermal instantaneous efficiency increases monotonically with redshift from 0.3 per cent at $z = 0$ to 0.6 per cent at $z = 3$. This increase with redshift is probably due to the median BH spin also increasing with redshift (Huško et al. in preparation), which also increases the radiative efficiencies. The cumulative total efficiency (in the bottom panel) is almost always higher in the hybrid AGN feedback model compared to the thermal one, with the exception of very high redshifts ($z > 8$).

The left panel Fig. 15 shows the BH mass dependence of the cumulative injected AGN feedback energies per BH by $z = 0$ for the L200m7 thermal and hybrid AGN feedback simulations. In the middle and right panels, we show the same quantity but

¹⁶Low jet efficiencies are a consequence of our use of constant and low accretion efficiencies ($\epsilon_{\text{acc}} = 0.01$) in the slim disc. Given the fact that the disc-scale accretion rate cannot exceed 100 times the Eddington rate, the horizon-scale Eddington ratio in the slim disc cannot be above the Eddington rate. This means that the dimensionless magnetic flux is always low in the slim disc, and thus also the jet efficiency (see equations 19 and 25). A more realistic choice for ϵ_{acc} , e.g. one where it scales with $f_{\text{Edd},d}$ but such that the horizon-scale accretion rate does not drop below Eddington, would ensure that jets are at least as important in the slim disc state as they are in the thin disc state.

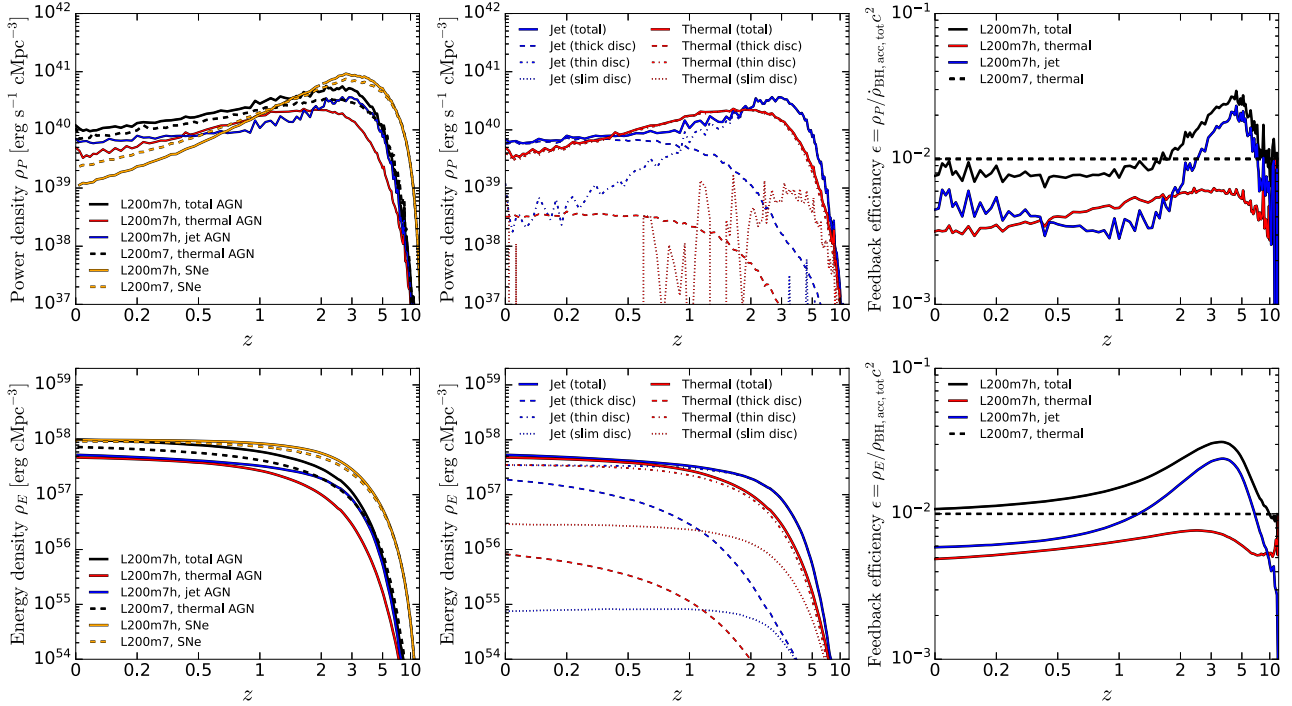


Figure 14. Cosmic evolution of feedback energetics in the L200m7 (dashed) and L200m7h (solid) simulations, using thermal and hybrid AGN feedback, respectively. The top row shows the feedback power densities, while the bottom row shows the cumulative injected feedback energy densities. Yellow lines show energies from SNe, while black lines show total AGN energies. The thin red and blue lines show the thermal and kinetic jet components of AGN feedback in the hybrid simulation. The middle panels show the thermal and jet components in the hybrid AGN feedback simulation split on to individual contributions from the three accretion disc states, as given by the legend. In the right panels, the top (bottom) shows the effective feedback efficiency, computed as the total power (energy) density divided by the total rest mass-energy accretion rate density (accreted rest mass-energy density).




Figure 15. *Left:* BH mass dependence of the cumulative injected AGN and SN feedback energies at $z = 0$ in the L200m7 (thermal) and L200m7h (hybrid) AGN feedback simulations. *Middle:* the cumulative injected feedback energies in units of $0.01 M_{\text{BH,acc}} c^2$. *Right:* as the middle panel but splitting the jet and thermal contributions in the hybrid AGN feedback simulation on to contributions from different accretion disc states.

divided by $0.01 M_{\text{BH,acc}} c^2$ in order to reduce the dynamic range, and since we expect the total AGN energy to roughly scale with the accreted BH mass. We also show energies injected by SNe of the host galaxies of these BHs. We find that the AGN-injected energies are comparable to SN-injected ones at all BH masses, and for both the thermal and hybrid AGN feedback models. This suggests that AGN feedback is relevant in both low-mass and high-mass galaxies, which is consistent with the effects we see when varying feedback parameters (Section 4.3), since these variations impact galaxies of all masses. This is true for parameter variations that affect either of the two AGN feedback modes.

Cumulative AGN-injected energies clearly overtake SN-injected ones only above $M_{\text{BH}} \approx 10^9 M_{\odot}$ (see middle panel of Fig. 15). The reason that this does not occur at lower mass, e.g. at $M_{\text{BH}} \approx 10^7 - 10^8 M_{\odot}$, where AGN-induced quenching begins to occur, is due to the fact that some of the SN energy does not do useful work since it is injected at low heating temperatures ($\Delta T < 10^{7.5}$ K) and is thus more prone to cooling losses. Furthermore, a clear transition at $M_{\text{BH}} \approx 10^7 - 10^8 M_{\odot}$ may be visible when comparing instantaneous feedback energies, but in Fig. 15 we are comparing cumulative energies injected over the entire history of a given BH (galaxy).

The most massive BHs appear to be injecting more feedback energy in the thermal AGN simulation than in the hybrid one, or at least to have injected it when integrated over time. In the hybrid AGN simulation, approximately half of the cumulative injected energy is in thermal and half in kinetic jet form by $z = 0$, regardless of BH mass. As can be seen in the right panel, the thermal component is almost entirely due to thin disc feedback, while for the jet channel, the thick disc contribution is dominant below $M_{\text{BH}} \approx 10^7 M_{\odot}$ and the thin disc one above that mass, up to $M_{\text{BH}} \approx 10^{10} M_{\odot}$. For the most massive BHs, thin and thick disc jets are approximately equally important in terms of total energy injected by $z = 0$, although this energy was injected at very different times for the two feedback channels.

7 CONCLUSIONS

We introduce a new, hybrid AGN feedback model for cosmological, hydrodynamical galaxy formation simulations. The model includes two feedback modes: thermal isotropic and kinetic jets. The thermal mode represents the effects of AGN disc winds shocked on subgrid scales and/or radiative heating on subgrid or larger scales. The jet mode represents the effects of relativistic jets driven by the R. D. Blandford & R. L. Znajek (1977) mechanism. The model includes three accretion disc states (Section 3.2): the thick, thin and slim disc states. BHs occupy these states depending on their Eddington ratio, with the three states populated with BHs that have low ($f_{\text{Edd},d} < 0.01$), moderate ($0.01 < f_{\text{Edd},d} < 1$) and super-Eddington ($f_{\text{Edd},d} > 1$) accretion rates, respectively. Here, $f_{\text{Edd},d}$ is the Eddington ratio defined using the accretion rate on to a subgrid accretion disc. Both feedback modes are used in all three accretion states, but with different efficiencies. We also include indirect effects of mass loss due to accretion disc winds in the thick and slim disc states (Section 3.3.1), through a new parameter: the accretion efficiency, ϵ_{acc} .

The model includes a subgrid prescription for BH spin (Sections 3.6 and 3.7), which evolves due to (1) gas accretion, (2) BH–BH mergers, (3) jet-induced spin-down, and (4) Lense–Thirring torques. This prescription is based on accretion disc theory, although it mostly uses numerical fitting functions obtained from general-relativistic (radiation) magnetohydrodynamical simulations. The BH spin affects the feedback efficiencies of all feedback modes, as well as jet directions (Section 3.5).

The thermal isotropic and kinetic jet feedback modes are implemented through BH heating and kicking events, respectively (Section 3.9), with the heating temperature, ΔT_{AGN} , and kick velocity, v_{jet} , modulating the explosiveness and frequency of AGN feedback (Section 3.10). The energy of each feedback event scales in proportion with the BH mass. In addition to ΔT_{AGN} and v_{jet} , our model has one free parameter that is directly related to AGN feedback: the coupling efficiency, ϵ_f , of radiation to thermal feedback in the thin disc state. This same parameter is present in single-mode AGN feedback models often used in galaxy formation simulations. However, parameters related to the accretion efficiency, introduced in our model, also have a strong indirect impact on feedback.

The hybrid AGN feedback model is implemented in the SWIFT code (M. Schaller et al. 2024) used for galaxy formation simulations, and has been coupled to the COLIBRE galaxy formation simulation suite (J. Schaye et al. 2025). While the flagship simulations of the COLIBRE suite used a thermal-only AGN feedback model (C. M. Booth & J. Schaye 2009), a set of runs use the hybrid AGN feedback model presented here. These runs were

performed at all three resolutions included in the COLIBRE suite (m7, m6, m5), but in boxes up to 8 times smaller in volume than the flagship runs, at each resolution.

In Section 4, we demonstrated the impact of varying three model parameters: the jet velocity normalization (Fig. 7), the thick disc accretion efficiency (Fig. 8), and the thermal feedback coupling efficiency (Fig. 9). We also discussed how the final model was calibrated. As shown there, AGN feedback has important effects in intermediate-mass and low-mass galaxies, all the way down to $M_* \approx 10^8 M_{\odot}$, and not just in massive galaxies as typically assumed. We confirmed this by considering the energies injected by AGN and comparing them to those injected by SNe (Figs 14 and 15).

New parameters introduced in the hybrid AGN feedback model require calibration, especially those that are more numerical in nature (e.g. jet velocities). Parameters that are also present in the thermal AGN feedback model, as well as SN-related parameters, require recalibration when used with the new hybrid AGN feedback model. Instead of an *ab initio* calibration of the entire set of stellar and AGN feedback parameters, as an initial guess we used the fiducial set of parameters from the thermal AGN COLIBRE simulations (E. Chaikin et al. 2025). We then varied these parameters as necessary to reproduce the same calibration galaxy and BH properties as the thermal AGN model (the GSMF, the stellar mass–size relation, and the stellar mass–BH mass relation, all at $z = 0$), but we also calibrate on the bolometric AGN luminosity function (at $z = 0.2$) since it holds useful information on how many BHs are radiatively efficient, corresponding to our thin disc state ($0.01 < f_{\text{Edd},d} < 1$). The calibration was done first for m7, and the parameter values were then slightly adjusted for m6 and m5. The parameters we (re)calibrate for the hybrid AGN feedback model are (1) the jet velocity normalization, $v_{\text{jet},0}$; (2) the accretion efficiency-related parameter, $r_{\text{tr},0}$; (3) the coupling efficiency ϵ_f ; (4) the BH seed mass, $M_{\text{BH,seed}}$; and (5) the stellar feedback pivot density $n_{\text{H,pivot}}$. The final, calibrated values of the parameters for the three resolutions are given in Table 3.

The calibration of the hybrid AGN feedback model presented in this study likely did not yield the best-fitting set of parameter values, particularly in relation to the slim accretion disc state (used for super-Eddington accretion). In the future, it will be useful to perform a calibration wider in scope (utilizing more parameters) and more efficient (utilizing emulators). Despite the more limited approach to the calibration, the hybrid AGN feedback simulations generally reproduce the observational data well (Fig. 10).

The BH spin–mass relation (Fig. 11), arguably the most important new prediction from the hybrid AGN model in relation to BHs, is in broad agreement with observations. The predicted relation has a peak for intermediate-mass BHs ($M_{\text{BH}} \in [10^6, 10^8] M_{\odot}$). These BHs typically have near-maximal spin, $|a| \approx 1$. Lower- and higher-mass BHs have lower spins. This shape of the BH spin–mass relation is linked to the relevance of (1) accretion in the thin disc state, which tends to spin-up BHs towards $|a| \approx 1$, (2) BH–BH mergers, which tend to result in spins significantly less than $|a| \approx 1$, and (3) jet-induced spin-down, which can almost completely spin down BHs.

Finally, in Section 6 we presented the first results on the operation of AGN feedback in the hybrid AGN feedback model, as found in the larger, production runs included in the COLIBRE suite. By $z = 0$, jet particles are able to reach large distances (of order a dozen Mpc, in some cases 100 virial radii) from their host galaxies and impact the IGM on larger scales than in the

purely thermal AGN model (Figs 12 and 13). This has implications for cosmological studies, where the impact of baryonic effects on matter power spectra are of great interest. The total amount of injected AGN energy (Fig. 14) is larger in the hybrid model than in the thermal one at all but the highest ($z > 5$) redshifts, by nearly 50 per cent. Jets provide roughly half of the total AGN energy (independent of BH mass; Fig. 15), with the thick disc state ($f_{\text{Edd},d} < 0.01$) dominating the jet channel at lower redshift ($z < 1$) and the thin disc state ($0.01 < f_{\text{Edd},d} < 1$) dominant at higher redshifts ($z > 1$, peaking at $z \approx 3$). Thermal isotropic feedback from the thin disc is important at all redshifts and dominates over jet feedback at $z \approx 0.5$ – 1.5 , but only slightly. Other feedback channels contribute negligibly to cosmic energy averages.

The hybrid AGN simulations use a different AGN feedback model compared to the fiducial COLIBRE simulations, but they are calibrated to the same observational data. This provides a useful estimate of the impact of different implementations of AGN feedback on: 1) galaxy properties that were not calibrated, 2) CGM, IGM and ICM properties, and 3) baryonic effects on cosmology, among others. These simulations can also be used to study various aspects of AGN feedback, including AGN jets.

The hybrid AGN feedback simulations will be a powerful tool facilitating studies focusing on BH and accretion physics in a fully cosmological context, especially since our model includes BH spin evolution and three distinct accretion disc states. In particular, the simulations can be used to study the statistics of spin for the BH population, as well as BH spin evolution. The inclusion of BH spin will also facilitate more accurate predictions of gravitational wave emission due to BH–BH mergers. While not discussed here, the outputs of the simulations include quantities necessary for the calculation of the radio emission of jets and lobes in post-processing. Predictions using this synthetic radio emission will be discussed in future studies.

ACKNOWLEDGEMENTS

FH thanks Nina Karačić for her work on Fig. 2, Bert Vandebroucke for help with implementing the hybrid AGN feedback model into SWIFT and COLIBRE, and Rob McGibbon for help with the processing and analysis of the hybrid AGN feedback simulations. The research in this paper made use of the *SWIFT* open-source simulation code (<http://www.swiftsim.com>, M. Schaller et al. 2024) version 1.0.0. The *SWIFT* simio Python library was used to analyze and visualize the data from the simulations (J. Borrow & A. Borrisov 2020, J. Borrow & A. J. Kelly 2021). FH acknowledges support from the Science Technology Facilities Council through a CDT studentship (ST/P006744/1). CGL acknowledges support from STFC consolidated grants ST/T000244/1 and ST/X001075/1. JT acknowledges support of a STFC Early Stage Research and Development grant (ST/X004651/1). ABL acknowledges support by the Italian Ministry for Universities (MUR) program ‘Dipartimenti di Eccellenza 2023–2027’ within the Centro Bicocca di Cosmologia Quantitativa (BiCoQ), and support by UNIMIB’s Fondo Di Ateneo Quota Competitiva (project 2024-ATEQC-0050). SP acknowledges support from the Austrian Science Fund (FWF) through grant number V 982-N. This project has received funding from the Netherlands Organization for Scientific Research (NWO) through research programme Athena 184.034.002. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility

(www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The data underlying the figures shown in this article will be provided upon reasonable request to the corresponding author. The SWIFT code is public (<https://gitlab.cosma.dur.ac.uk/swift/swiftsim>), while the COLIBRE modules will be made available after the public data release. People interested in using the hybrid AGN feedback simulations before that point are encouraged to contact the corresponding author.

REFERENCES

- Abbott T. M. C. et al., 2022, *Phys. Rev. D*, 105, 023520
 Abramowicz M. A., Czerny B., Lasota J. P., Szuszkiewicz E., 1988, *ApJ*, 332, 646
 Aird J., Coil A. L., Georgakakis A., Nandra K., Barro G., Pérez-González P. G., 2015, *MNRAS*, 451, 1892
 Akiyama K. et al., 2022, *ApJ*, 930, L16
 Ali-Haïmoud Y., Bird S., 2013, *MNRAS*, 428, 3375
 Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017, *MNRAS*, 464, 2840
 Anglés-Alcázar D. et al., 2021, *ApJ*, 917, 53
 Arjona-Gálvez E., Di Cintio A., Grand R. J. J., 2024, *A&A*, 690, A286
 Avara M. J., McKinney J. C., Reynolds C. S., 2016, *MNRAS*, 462, 636
 Ayromlou M., Nelson D., Pillepich A., 2023, *MNRAS*, 524, 5391
 Bahé Y. M. et al., 2022, *MNRAS*, 516, 167
 Bambi C. et al., 2021, *Space Sci. Rev.*, 217, 65
 Barausse E., 2012, *MNRAS*, 423, 2533
 Barausse E., Rezzolla L., 2009, *ApJ*, 704, L40
 Barausse E., Morozova V., Rezzolla L., 2012, *ApJ*, 758, 63
 Bardeen J. M., 1970, *Nature*, 226, 64
 Bardeen J. M., Petterson J. A., 1975, *ApJ*, 195, L65
 Bauer A. E. et al., 2013, *MNRAS*, 434, 209
 Beckmann R. S. et al., 2025, *MNRAS*, 536, 1838
 Begelman M. C., Scepi N., Dexter J., 2022, *MNRAS*, 511, 2040
 Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, *MNRAS*, 488, 3143
 Benítez-Llambay A. et al., 2026, *MNRAS*, 546, stag268
 Bennett J. S., Sijacki D., Costa T., Laporte N., Witten C., 2024, *MNRAS*, 527, 1033
 Biermann P. L., Strittmatter P. A., 1987, *ApJ*, 322, 643
 Bird S., Ni Y., Di Matteo T., Croft R., Feng Y., Chen N., 2022, *MNRAS*, 512, 3703
 Birzan L., Rafferty D. A., McNamara B. R., Wise M. W., Nulsen P. E. J., 2004, *ApJ*, 607, 800
 Blandford R. D., Begelman M. C., 1999, *MNRAS*, 303, L1
 Blandford R. D., Königl A., 1979, *ApJ*, 232, 34
 Blandford R. D., Znajek R. L., 1977, *MNRAS*, 179, 433
 Boehringer H., Voges W., Fabian A. C., Edge A. C., Neumann D. M., 1993, *MNRAS*, 264, L25
 Bondi H., 1952, *MNRAS*, 112, 195
 Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
 Booth C. M., Schaye J., 2010, *MNRAS*, 405, L1
 Booth C. M., Schaye J., 2013, *Scientific Reports*, 3, 1738
 Borrow J., Borrisov A., 2020, *The J. Open Source Softw.*, 5, 2430
 Borrow J., Kelly A. J., 2021, preprint ([arXiv:2106.05281](https://arxiv.org/abs/2106.05281))
 Borrow J., Anglés-Alcázar D., Davé R., 2020, *MNRAS*, 491, 6102
 Borrow J., Schaller M., Bower R. G., Schaye J., 2022, *MNRAS*, 511, 2367
 Bourne M. A., Sijacki D., 2021, *MNRAS*, 506, 488
 Bourne M. A., Yang H.-Y. K., 2023, *Galaxies*, 11, 73

- Bourne M. A., Fiacconi D., Sijacki D., Piotrowska J. M., Koudmani S., 2024, *MNRAS*, 534, 3448
- 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
- Bryan G. L., Norman M. L., 1998, *ApJ*, 495, 80
- Bustamante S., Springel V., 2019, *MNRAS*, 490, 4133
- Carr D. S., Kannappan S. J., Hutchens Z. L., Polimera M. S., Norris M. A., Eckert K. D., Moffett A. J., 2025, *ApJ*, 985, 25
- Catinella B. et al., 2018, *MNRAS*, 476, 875
- Cenci E., Sala L., Lupi A., Capelo P. R., Dotti M., 2021, *MNRAS*, 500, 3719
- Chabrier G., 2003, *PASP*, 115, 763
- Chaikin E., Schaye J., Schaller M., Bahé Y. M., Nobels F. S. J., Ploekinger S., 2022, *MNRAS*, 514, 249
- Chaikin E., Schaye J., Schaller M., Benítez-Llambay A., Nobels F. S. J., Ploekinger S., 2023, *MNRAS*, 523, 3709
- Chaikin E. et al., 2025, preprint (arXiv:2509.04067)
- Chang Y.-Y., van der Wel A., da Cunha E., Rix H.-W., 2015, *ApJS*, 219, 8
- Cheung E. et al., 2016, *Nature*, 533, 504
- Cho H., Prather B. S., Su K.-Y., Narayan R., Natarajan P., 2024, *ApJ*, 977, 200
- Choi E., Ostriker J. P., Naab T., Johansson P. H., 2012, *ApJ*, 754, 125
- Cielo S., Antonuccio-Delogu V., Macciò A. V., Romeo A. D., Silk J., 2014, *MNRAS*, 439, 2903
- Costa T., Sijacki D., Haehnelt M. G., 2014, *MNRAS*, 444, 2355
- Costa T., Rosdahl J., Sijacki D., Haehnelt M. G., 2018, *MNRAS*, 473, 4197
- Crenshaw D. M., Kraemer S. B., George I. M., 2003, *ARA&A*, 41, 117
- Croton D. J. et al., 2006, *MNRAS*, 365, 11
- D'Silva J. C. J., Driver S. P., Lagos C. D. P., Robotham A. S. G., Summers J., Windhorst R. A., 2023, *ApJ*, 959, L18
- Dalla Vecchia C., Schaye J., 2012, *MNRAS*, 426, 140
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieeferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
- Davis T. A., Greene J. E., Ma C.-P., Blakeslee J. P., Dawson J. M., Pandya V., Veale M., Zabel N., 2019, *MNRAS*, 486, 1404
- Dehnen W., 2014, *Computational Astrophysics and Cosmol.*, 1, 1
- Delvecchio I. et al., 2014, *MNRAS*, 439, 2736
- Di Matteo T., Croft R. A. C., Springel V., Hernquist L., 2003, *ApJ*, 593, 56
- Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
- Done C., Gierliński M., Kubota A., 2007, *A&A Rev.*, 15, 1
- Driver S. P. et al., 2022, *MNRAS*, 513, 439
- Dubois Y., Devriendt J., Slyz A., Teyssier R., 2012, *MNRAS*, 420, 2662
- Dubois Y., Volonteri M., Silk J., 2014, *MNRAS*, 440, 1590
- Dubois Y. et al., 2021, *A&A*, 651, A109
- Eckert D., Gaspari M., Gastaldello F., Le Brun A. M. C., O'Sullivan E., 2021, *Universe*, 7, 142
- Edge A. C., 2001, *MNRAS*, 328, 762
- Eke V. R., Cole S., Frenk C. S., 1996, *MNRAS*, 282, 263
- Event Horizon Telescope Collaboration, 2019, *ApJ*, 875, L5
- Event Horizon Telescope Collaboration, 2021, *ApJ*, 910, L13
- Fabian A. C., 1994, *ARA&A*, 32, 277
- Fabian A. C., 1999, *MNRAS*, 308, L39
- Fabian A. C., 2012, *ARA&A*, 50, 455
- Falle S. A. E. G., 1991, *MNRAS*, 250, 581
- Fan X. et al., 2001, *AJ*, 122, 2833
- Fanidakis N., Baugh C. M., Benson A. J., Bower R. G., Cole S., Done C., Frenk C. S., 2011, *MNRAS*, 410, 53
- Feruglio C., Maiolino R., Piconcelli E., Menci N., Aussel H., Lamastra A., Fiore F., 2010, *A&A*, 518, L155
- Fiacconi D., Sijacki D., Pringle J. E., 2018, *MNRAS*, 477, 3807
- Fiore F. et al., 2017, *A&A*, 601, A143
- Forouhar Moreno V. J., Helly J., McGibbon R., Schaye J., Schaller M., Han J., Kugel R., 2025, preprint (arXiv:2502.06932)
- Fraser-McKelvie A., Brown M. J. I., Pimblet K. A., 2014, *MNRAS*, 444, L63
- Gaspari M., Ruszkowski M., Oh S. P., 2013, *MNRAS*, 432, 3401
- Gebhardt M. et al., 2024, *MNRAS*, 529, 4896
- Ghisellini G., Tavecchio F., Maraschi L., Celotti A., Sbarrato T., 2014, *Nature*, 515, 376
- Ghosh R., Laha S., 2021, *ApJ*, 908, 198
- Graham A. W., Sahu N., 2023, *MNRAS*, 518, 2177
- Graham A. W., Sahu N., 2024, *MNRAS*, 530, 3429
- Greene J. E. et al., 2013, *ApJ*, 771, 121
- Greengard L., Rokhlin V., 1987, *J. Comput. Phys.*, 73, 325
- Greenstein J. L., Schmidt M., 1964, *ApJ*, 140, 1
- Griffin A. J., Lacey C. G., Gonzalez-Perez V., Lagos C. d. P., Baugh C. M., Fanidakis N., 2019, *MNRAS*, 487, 198
- Gull S. F., Northover K. J. E., 1973, *Nature*, 244, 80
- Guo M., Stone J. M., Kim C.-G., Quataert E., 2023, *ApJ*, 946, 26
- Hagen S., Done C., 2023, *MNRAS*, 525, 3455
- Hahn O., Michaux M., Rampf C., Uhlemann C., Angulo R. E., 2020, *Astrophysics Source Code Library*, record ascl:2008.024
- Han J., Cole S., Frenk C. S., Benítez-Llambay A., Helly J., 2018, *MNRAS*, 474, 604
- Hardwick J. A., Cortese L., Obreschkow D., Catinella B., 2022, *MNRAS*, 516, 4043
- Harrison C. M., Costa T., Tadhunter C. N., Flütsch A., Kakkad D., Perna M., Vietri G., 2018, *Nat. Astron.*, 2, 198
- Heckman T. M., Best P. N., 2014, *ARA&A*, 52, 589
- Hirschmann M., Dolag K., Saro A., Bachmann L., Borgani S., Burkert A., 2014, *MNRAS*, 442, 2304
- Hlavacek-Larrondo J., Fabian A. C., Edge A. C., Ebeling H., Sanders J. S., Hogan M. T., Taylor G. B., 2012, *MNRAS*, 421, 1360
- Hobbs A., Power C., Nayakshin S., King A. R., 2012, *MNRAS*, 421, 3443
- Hopkins P. F., 2025, *The Open J. Astrophys.*, 8, 56
- Hopkins P. F., Quataert E., 2011, *MNRAS*, 415, 1027
- Hopkins P. F., Hernquist L., Hayward C. C., Narayanan D., 2012, *MNRAS*, 425, 1121
- Hopkins P. F. et al., 2024, *The Open Journal of Astrophysics*, 7, 19
- Hoyle F., Lyttleton R. A., 1939, *Proc. Cambridge Philosophical Society*, 35, 405
- Hoyle F., Lyttleton R. A., 1941, *MNRAS*, 101, 227
- Hudson D. S., Mittal R., Reiprich T. H., Nulsen P. E. J., Andernach H., Sarazin C. L., 2010, *A&A*, 513, A37
- Huško F., Lacey C. G., 2023a, *MNRAS*, 520, 5090
- Huško F., Lacey C. G., 2023b, *MNRAS*, 521, 4375
- Huško F., Lacey C. G., Schaye J., Schaller M., Nobels F. S. J., 2022, *MNRAS*, 516, 3750
- Huško F., Lacey C. G., Schaye J., Nobels F. S. J., Schaller M., 2024, *MNRAS*, 527, 5988
- Huško F., Lacey C. G., Roper W. J., Schaye J., Briggs J. M., Schaller M., 2025, *MNRAS*, 537, 2559
- Jetha N. N., Hardcastle M. J., Sakelliou I., 2006, *MNRAS*, 368, 609
- Jiang J. et al., 2019, *MNRAS*, 489, 3436
- Jiang Y.-F., Stone J. M., Davis S. W., 2019, *ApJ*, 880, 67
- Kaiser C. R., Alexander P., 1997, *MNRAS*, 286, 215
- Kaviraj S., Shabala S. S., Deller A. T., Middelberg E., 2015, *MNRAS*, 454, 1595
- Kaviraj S. et al., 2017, *MNRAS*, 467, 4739
- Kerr R. P., 1963, *Phys. Rev. Lett.*, 11, 237
- Khandai N., Di Matteo T., Croft R., Wilkins S., Feng Y., Tucker E., DeGraf C., Liu M.-S., 2015, *MNRAS*, 450, 1349
- King A. R., Lubow S. H., Ogilvie G. I., Pringle J. E., 2005, *MNRAS*, 363, 49
- King A. R., Pringle J. E., Livio M., 2007, *MNRAS*, 376, 1740
- Kino M., Kawakatu N., 2005, *MNRAS*, 364, 659
- Koudmani S., Henden N. A., Sijacki D., 2021, *MNRAS*, 503, 3568
- Koudmani S., Sijacki D., Smith M. C., 2022, *MNRAS*, 516, 2112
- Koudmani S., Somerville R. S., Sijacki D., Bourne M. A., Jiang Y.-F., Profit K., 2024, *MNRAS*, 532, 60
- Kroupa P., 2001, *MNRAS*, 322, 213
- Krumholz M. R., McKee C. F., Klein R. I., 2006, *ApJ*, 638, 369
- Kuo C. Y. et al., 2014, *ApJ*, 783, L33
- Lagos C. D. P., Cora S. A., Padilla N. D., 2008, *MNRAS*, 388, 587
- Lagos C. D. P., Padilla N. D., Cora S. A., 2009, *MNRAS*, 395, 625
- Laine S. et al., 2020, *ApJ*, 894, L1

- Lense J., Thirring H., 1918, *Physikalische Zeitschrift*, 19, 156
- Levine R., Gnedin N. Y., Hamilton A. J. S., 2010, *ApJ*, 716, 1386
- Li H., Chen Y., Wang H., Mo H., 2025, *MNRAS*, 543, 1878
- Liu H., Wu Q., 2013, *ApJ*, 764, 17
- Lodato G., 2007, *Nuovo Cimento Rivista Serie*, 30, 293
- Lodato G., Price D. J., 2010, *MNRAS*, 405, 1212
- Lodato G., Pringle J. E., 2007, *MNRAS*, 381, 1287
- Ludlow A. D., Schaye J., Schaller M., Richings J., 2019, *MNRAS*, 488, L123
- Ludlow A. D., Fall S. M., Wilkinson M. J., Schaye J., Obreschkow D., 2023, *MNRAS*, 525, 5614
- Maccarone T. J., Gallo E., Fender R., 2003, *MNRAS*, 345, L19
- Madau P., Haardt F., Dotti M., 2014, *ApJ*, 784, L38
- Mahadevan R., 1997, *ApJ*, 477, 585
- Mallick L. et al., 2022, *MNRAS*, 513, 4361
- Markoff S., Falcke H., Fender R., 2001, *A&A*, 372, L25
- Martin R. G., Pringle J. E., Tout C. A., 2007, *MNRAS*, 381, 1617
- Marziani P., Garnica Luna K., Floris A., del Olmo A., Deconto-Machado A., Buendia-Rios T. M., Negrete C. A., Dultzin D., 2025, *Universe*, 11, 69
- Massonneau W., Volonteri M., Dubois Y., Beckmann R. S., 2023, *A&A*, 670, A180
- McConnell N. J., Ma C.-P., 2013, *ApJ*, 764, 184
- McDonald M. et al., 2015, *ApJ*, 811, 111
- McGibbon R., Helly J. C., Schaye J., Schaller M., Vandenbroucke B., 2025, *J. Open Source Softw.*, 10, 8252,
- McKinney J. C., Tchekhovskoy A., Blandford R. D., 2012, *MNRAS*, 423, 3083
- McNamara B. R., Nulsen P. E. J., Wise M. W., Rafferty D. A., Carilli C., Sarazin C. L., Blanton E. L., 2005, *Nature*, 433, 45
- Merloni A., Heinz S., 2008, *MNRAS*, 388, 1011
- Mezcua M., Suh H., Civano F., 2019, *MNRAS*, 488, 685
- Michaux M., Hahn O., Rampf C., Angulo R. E., 2021, *MNRAS*, 500, 663
- Mitchell P. D., Schaye J., Bower R. G., Crain R. A., 2020, *MNRAS*, 494, 3971
- Mocz P., Fabian A. C., Blundell K. M., 2013, *MNRAS*, 432, 3381
- Moderski R., Sikora M., 1996, *MNRAS*, 283, 854
- Moffett A. J. et al., 2016, *MNRAS*, 462, 4336
- Mullin L. M., Hardcastle M. J., 2009, *MNRAS*, 398, 1989
- Murray N., Chiang J., Grossman S. A., Voit G. M., 1995, *ApJ*, 451, 498
- Narayan R., Yi I., 1994, *ApJ*, 428, L13
- Narayan R., Yi I., 1995, *ApJ*, 452, 710
- Narayan R., Igumenshchev I. V., Abramowicz M. A., 2003, *PASJ*, 55, L69
- Narayan R., Sądowski A., Penna R. F., Kulkarni A. K., 2012, *MNRAS*, 426, 3241
- Narayan R., Chael A., Chatterjee K., Ricarte A., Curd B., 2022, *MNRAS*, 511, 3795
- Nelson D. et al., 2019, *Computat. Astrophys. Cosmol.*, 6, 2
- Nemmen R. S., Tchekhovskoy A., 2015, *MNRAS*, 449, 316
- Ni Y. et al., 2022, *MNRAS*, 513, 670
- Nobels F. S. J., Schaye J., Schaller M., Bahé Y. M., Chaikin E., 2022, *MNRAS*, 515, 4838
- Nobels F. S. J., Schaye J., Schaller M., Ploeckinger S., Chaikin E., Richings A. J., 2024, *MNRAS*, 532, 3299
- Noda H., Done C., 2018, *MNRAS*, 480, 3898
- Nomura M., Ohsuga K., Takahashi H. R., Wada K., Yoshida T., 2016, *PASJ*, 68, 16
- Novikov I. D., Thorne K. S., 1973, in *Black Holes (Les Astres Occlus)*. p. 343
- Ogilvie G. I., 1999, *MNRAS*, 304, 557
- Papaloizou J. C. B., Pringle J. E., 1983, *MNRAS*, 202, 1181
- Penny S. J. et al., 2018, *MNRAS*, 476, 979
- Ploeckinger S., Schaye J., 2020, *MNRAS*, 497, 4857
- Ploeckinger S., Richings A. J., Schaye J., Trayford J. W., Schaller M., Chaikin E., 2025, *MNRAS*, 543, 891
- Pringle J. E., 1992, *MNRAS*, 258, 811
- Quera-Bofarull A., Done C., Lacey C. G., Nomura M., Ohsuga K., 2023, *MNRAS*, 518, 2693
- Rafferty D. A., McNamara B. R., Nulsen P. E. J., Wise M. W., 2006, *ApJ*, 652, 216
- Rennehan D., Babul A., Moa B., Davé R., 2024, *MNRAS*, 532, 4793
- Reynolds C. S., 2021, *ARA&A*, 59, 117
- Ricarte A., Narayan R., Curd B., 2023, *ApJL*, 954, L22
- Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349
- Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780
- Roy N. et al., 2024, *ApJ*, 970, 69
- Ruffa I., Laing R. A., Prandoni I., Paladino R., Parma P., Davis T. A., Bureau M., 2020, *MNRAS*, 499, 5719
- Russell H. R., McNamara B. R., Edge A. C., Hogan M. T., Main R. A., Vantyghem A. N., 2013, *MNRAS*, 432, 530
- Sabater J. et al., 2019, *A&A*, 622, A17
- Sądowski A., Narayan R., Penna R., Zhu Y., 2013, *MNRAS*, 436, 3856
- Saikia D. J., 2022, *JA&A*, 43, 97
- Sala L., Valentini M., Biffi V., Dolag K., 2024, *A&A*, 685, A92
- Salehirad S., Reines A. E., Molina M., 2025, *ApJ*, 979, 26
- Salim S. et al., 2007, *ApJS*, 173, 267
- Sarazin C. L., 1986, *Rev. Mod. Phys.*, 58, 1
- Sbarrato T., Ghisellini G., Giovannini G., Giroletti M., 2021, *A&A*, 655, A95
- Sbarrato T., Ghisellini G., Tagliaferri G., Tavecchio F., Ghirlanda G., Costamante L., 2022, *A&A*, 663, A147
- Schaller M. et al., 2024, *MNRAS*, 530, 2378
- Schaye J. et al., 2010, *MNRAS*, 402, 1536
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schaye J. et al., 2025, preprint ([arXiv:2508.21126](https://arxiv.org/abs/2508.21126))
- Schmidt M., 1968, *ApJ*, 151, 393
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 337, 24
- Shapiro S. L., 2005, *ApJ*, 620, 59
- Sharma P., Quataert E., Hammett G. W., Stone J. M., 2007, *ApJ*, 667, 714
- Sharma R. S., Brooks A. M., Somerville R. S., Tremmel M., Bellovary J., Wright A. C., Quinn T. R., 2020, *ApJ*, 897, 103
- Shen X., Hopkins P. F., Faucher-Giguère C.-A., Alexander D. M., Richards G. T., Ross N. P., Hickox R. C., 2020, *MNRAS*, 495, 3252
- Shi F., Yuan F., Tombesi F., Xie F.-g., 2025, preprint ([arXiv:2502.08138](https://arxiv.org/abs/2502.08138))
- Shin E.-J., Sijacki D., Smith M. C., Bourne M. A., Koudmani S., 2025, *MNRAS*, 544, 2467
- Shulevski A. et al., 2019, *A&A*, 628, A69
- Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
- Sisk-Reynés J., Reynolds C. S., Matthews J. H., Smith R. N., 2022, *MNRAS*, 514, 2568
- Sądowski A., 2009, *ApJS*, 183, 171
- Sądowski A., Narayan R., McKinney J. C., Tchekhovskoy A., 2014, *MNRAS*, 439, 503
- Šlaus B. et al., 2024, *A&A*, 684, A19
- Smolčić V. et al., 2017, *A&A*, 602, A6
- Stanek R., Evrard A. E., Böhringer H., Schuecker P., Nord B., 2006, *ApJ*, 648, 956
- Steffen J. L. et al., 2023, *ApJ*, 942, 107
- Steinborn L. K., Dolag K., Hirschmann M., Prieto M. A., Remus R.-S., 2015, *MNRAS*, 448, 1504
- Stella L., Vietri M., 1998, *ApJ*, 492, L59
- Su K.-Y. et al., 2021, *MNRAS*, 507, 175
- Sun W.-H., Malkan M. A., 1989, *ApJ*, 346, 68
- Taam R. E., Liu B. F., Yuan W., Qiao E., 2012, *ApJ*, 759, 65
- Talbot R. Y., Bourne M. A., Sijacki D., 2021, *MNRAS*, 504, 3619
- Tchekhovskoy A., Narayan R., McKinney J. C., 2010, *ApJ*, 711, 50
- Thorne K. S., 1974, *ApJ*, 191, 507
- Tombesi F., Cappi M., Reeves J. N., Palumbo G. G. C., Yaqoob T., Braito V., Dadina M., 2010, *A&A*, 521, A57
- Trayford J. W. 2026, *MNRAS*, 545, staf2040
- Tremmel M., Karcher M., Governato F., Volonteri M., Quinn T. R., Pontzen A., Anderson L., Bellovary J., 2017, *MNRAS*, 470, 1121
- Urry C. M., Padovani P., 1995, *PASP*, 107, 803
- Valtonen M. J. et al., 2016, *ApJ*, 819, L37
- Vasudevan R. V., Fabian A. C., Reynolds C. S., Aird J., Dauser T., Gallo L. C., 2016, *MNRAS*, 458, 2012

- Vazquez Najar L. A., Andernach H., 2019, preprint (arXiv:1908.09989)
- Vogelsberger M. et al., 2014, *Nature*, 509, 177
- Volonteri M., Madau P., Quataert E., Rees M. J., 2005, *ApJ*, 620, 69
- Walton D. J., Nardini E., Fabian A. C., Gallo L. C., Reis R. C., 2013, *MNRAS*, 428, 2901
- Walton D. J. et al., 2021, *MNRAS*, 506, 1557
- Wang J.-M., Zhou Y.-Y., 1999, *ApJ*, 516, 420
- Weinberger R. et al., 2017, *MNRAS*, 465, 3291
- Wellons S. et al., 2023, *MNRAS*, 520, 5394
- Whitaker K. E., van Dokkum P. G., Brammer G., Franx M., 2012, *ApJ*, 754, L29
- Wiklund T., Combes F., Henkel C., 1995, *A&A*, 297, 643
- Wise M. W., McNamara B. R., Nulsen P. E. J., Houck J. C., David L. P., 2007, *ApJ*, 659, 1153
- Xie F.-G., Yuan F., 2012, *MNRAS*, 427, 1580
- Young L. M. et al., 2011, *MNRAS*, 414, 940
- Yuan F., 2003, *ApJ*, 594, L99
- Yuan F., Narayan R., 2014, *ARA&A*, 52, 529
- Yuan F., Bu D., Wu M., 2012, *ApJ*, 761, 130
- Yuan F., Wang H., Yang H., 2022, *ApJ*, 924, 124
- Zhang Z. et al., 2021, *A&A*, 650, A155
- Zheng X., Zhang Y., Röttgering H., 2024, *A&A*, 686, A169

APPENDIX A: SIMPLIFIED VERSIONS OF THE MODEL

In F. Huško et al. (2022, 2024, 2025), we introduced the aspects of the model that concern the thick, thin and slim disc states, respectively. We studied AGN feedback in these individual accretion disc states separately, but also in hybrid variants. These studies were performed using idealized galaxy clusters in the first two cases, and high-redshift cosmological simulations of a proto-cluster in the latter case. The first study modelled only the thick disc and only kinetic jet feedback, using BH spin evolution. The second study also considered the thin disc, comparing models that apply either the thick disc state and its jets or the thin disc and its thermal feedback at all Eddington ratios. There, a hybrid model where the two modes switched at $f_{\text{Edd},d} = 0.01$ was also considered. In the final study, the super-Eddington slim disc was introduced, with its jet and wind (thermal) feedback, considering only hybrid models where the thin disc transitions to the slim disc state at $f_{\text{Edd},d} = 1$.

In the hybrid model that we introduce here, we combine these three accretion disc states in a similar way as in the above studies. Similar to the last of the above three studies, we include jets in the thin disc, which was not done in the second study. Here we also include wind (thermal) feedback in the thick disc, which was not included in either of the first two studies. Another major modification to the model presented here is that we include accretion efficiencies in the thick and slim disc, both on account of their mass-loss induced by strong winds. Finally, the previous versions of the hybrid AGN feedback model were coupled with the SWIFT-EAGLE galaxy formation model, whereas here we couple it with COLIBRE.

APPENDIX B: THE TRANSITION BETWEEN ACCRETION DISC STATES

The state of the accretion flow (thick, thin or slim disc) is assumed to depend on the dimensionless accretion rate, i.e. the Eddington ratio. However, it is not immediately clear if one should use the disc-scale Eddington ratio, $f_{\text{Edd},d}$, or the horizon-scale one, $f_{\text{Edd},H}$, to determine whether the accretion disc should transition to another state. This choice is discussed in Section B1, whereas

in Section B2 we discuss the choice of critical Eddington ratios at which these transitions occur.

B1 Disc-scale versus horizon-scale Eddington ratios

From a modelling perspective, it is not clear whether accretion discs should transition between different states based on the value of the horizon-scale accretion rate or the inflow rate on to the disc. This is a question that concerns only simulations and not reality. The reason is that the transition between the thick and thin disc is, in reality, probably smooth rather than abrupt. In an ideal scenario, as a BH that is deep in the thick disc (i.e. both $f_{\text{Edd},d} \ll 0.01$ and $f_{\text{Edd},H} \ll 0.01$) progressively increases its accretion rate, its accretion efficiency should also progressively increase until it reaches $\epsilon_{\text{acc}} \approx 1$ at the critical Eddington ratio $f_{\text{Edd}} \approx 0.01$. At this point, it does not matter whether one uses $f_{\text{Edd},d}$ or $f_{\text{Edd},H} = \epsilon_{\text{acc}} f_{\text{Edd},d}$ to decide the transition between the states, since both will give the same outcome. The same is true if the disc is transitioning from the thin to the thick disc state.

Let us now consider the above scenario in the context of our model. If the disc is currently thin, $\epsilon_{\text{acc}} = 1$ and $f_{\text{Edd},H} > 0.01$ (therefore also $f_{\text{Edd},d} > 0.01$), the disc will stay thin. Let us imagine then slightly reducing $f_{\text{Edd},H}$ to $f_{\text{Edd},H} = 0.009$. The disc will become thick, with $\epsilon_{\text{acc}} \ll 1$ and $f_{\text{Edd},d} = \epsilon_{\text{acc}} f_{\text{Edd},H} \ll 0.009$. If the accretion rate again slightly increases (so that $f_{\text{Edd},H} > 0.01$), the disc may or may not transition to the thin state depending on whether we consider $f_{\text{Edd},H}$ or $f_{\text{Edd},d}$. If we consider $f_{\text{Edd},H}$, the accretion rate would need to increase by nearly a factor of 100 (assuming a per cent level accretion efficiency), i.e. $f_{\text{Edd},d}$ would need to be of order unity, for the disc to transition back to the thin state. We consider this to be unrealistic, and instead use $f_{\text{Edd},d}$ to decide the transition. This is supported by our numerical experiments (not shown here) where we attempted to use $f_{\text{Edd},H}$ to decide the transition, but found it impossible to calibrate realistic models with that choice.

The above discussion may point to our accretion efficiency prescription being somewhat unrealistic, more specifically, the fact that accretion efficiencies suddenly drop to very low values even if the disc is only slightly in the thick state. This is an outcome of our calibration rather than an *a priori* assumption we made, with the $r_{\text{tr},0}$ parameter having a large value (see Section 4.3.2). We considered models that have accretion efficiencies of order unity for accretion discs that are only slightly in the thick state, but our attempt at calibrating such models was unsuccessful. We believe that this is related to the jet feedback prescription for the thick disc. We use jet efficiency formulas that depend only on BH spin, i.e. jets are very strong for discs that are only slightly in the thick disc state. In reality, however, jets should be relatively weak at such high Eddington ratios, and radiative feedback stronger (F. Yuan & R. Narayan 2014). Our need to use very low accretion efficiencies therefore probably compensates for the overly strong jets at the point of transition between the thick and thin disc states. Future general-relativistic radiative magnetohydrodynamical (GRRMHD) simulations of the transition between the thick and thin disc state will likely resolve this issue from a theoretical perspective.

B2 Critical Eddington ratios

The transition between the thin disc at moderate Eddington ratios and the slim disc is simply set to occur at $f_{\text{Edd},H} = 1$, in other

words, we assume that the slim disc coincides with the super-Eddington accretion disc state. The transition between the thick (R. Narayan & I. Yi 1994) and thin (N. I. Shakura & R. A. Sunyaev 1973) discs is more subtle. It should occur around $f_{\text{Edd},d} \approx \alpha^2$ (R. Narayan & I. Yi 1995), where α is the accretion disc viscosity parameter. More recent calculations suggest that the properties of the thick disc already begin to change at $f_{\text{Edd},H} = 0.2\alpha^2$, and the transition appears to be complete by $f_{\text{Edd},d} = 0.7\alpha$ (F. Yuan & R. Narayan 2014). Between these two values, the disc takes on a transition state whose properties are not well understood. For conceivable values of α , which may be as low as 0.05 and as high as 0.1–0.4 (A. R. King, J. E. Pringle & M. Livio 2007, F. Yuan & R. Narayan 2014), the transition state may occupy the range 0.001–0.3 in $f_{\text{Edd},d}$. Observations of both X-ray binary (C. Done et al. 2007) and AGN spectra (H. Noda & C. Done 2018) find this transition to occupy a narrower range of $f_{\text{Edd},d} = 0.01$ –0.03. An analysis of the radiative and mechanical powers of AGN shows the transition to span the same range (H. R. Russell et al. 2013). We assume the lower end of this range to be the critical transition rate at which the two accretion disc states (the thin and thick disc) interchange; $f_{\text{Edd},d,\text{crit}} = 0.01$.

Given this choice, we can set a value for the viscosity parameter α , which appears in many of the equations describing accretion disc structure that we will discuss. For this purpose we use the finding of numerical calculations that the transition spans the range between $0.2\alpha^2$ and 0.7α in $f_{\text{Edd},d}$ (F. Yuan & R. Narayan 2014). We assume that the geometric mean of these two boundaries corresponds to $f_{\text{Edd},d,\text{crit},0} = 0.01$, which is true for $\alpha \approx 0.1$, so we set $\alpha = 0.1$ in this model.

APPENDIX C: INNERMOST STABLE CIRCULAR ORBIT

The radius of the ISCO (R. P. Kerr 1963) is given by $R_{\text{ISCO}} = r_{\text{ISCO}}R_G$, where r_{ISCO} is the dimensionless radius of the ISCO. r_{ISCO} depends on BH spin as

$$r_{\text{ISCO}} = 3 + Z_2 \mp \sqrt{(3 - Z_1)(3 + Z_1 + 2Z_2)}, \quad (\text{C1})$$

where the minus and plus sign are for prograde and retrograde accretion, respectively, and Z_1 and Z_2 are functions of BH spin given by

$$Z_1(a) = 1 + (1 - a^2)^{1/3} \left[(1 + |a|)^{1/3} + (1 - |a|)^{1/3} \right] \quad (\text{C2})$$

and

$$Z_2(a) = \sqrt{3a^2 + Z_1^2}. \quad (\text{C3})$$

The specific angular momentum at the ISCO is given by $L_{\text{ISCO}} = M_{\text{BH}}G\ell_{\text{ISCO}}/c$, where ℓ_{ISCO} is a dimensionless function of BH spin:

$$\ell_{\text{ISCO}}(a) = \frac{2}{3\sqrt{3}} \left(1 + 2\sqrt{3r_{\text{ISCO}} - 2} \right). \quad (\text{C4})$$

APPENDIX D: RADIATIVE EFFICIENCY IN THE THICK AND SLIM DISCS

In the thick disc state, most of the thermal energy is advected inwards towards the BH before it has time to radiate away. This results in a low, but non-zero radiative efficiency. This efficiency is mostly governed by radiative cooling rates of electrons through synchrotron, bremsstrahlung, and Compton processes. Numerical simulations have not focused on radiative efficiencies of thick

discs, so we take the results from an early numerical study (R. Mahadevan 1997), in which cooling processes were studied in the context of the original thick disc solution of R. Narayan & I. Yi (1994) (note that more recent results have been obtained by F.-G. Xie & F. Yuan 2012, although their formulae are not as convenient for general numerical implementation).

The numerical study in question (R. Mahadevan 1997) found two different regimes: for $f_{\text{Edd},H} < f_{\text{Edd},\text{crit},\text{visc}}$, viscous heating dominates the heating of electrons, whereas for $f_{\text{Edd},\text{crit},\text{visc}} < f_{\text{Edd},H} < f_{\text{Edd},\text{crit},\text{thick}}$, it is dominated by ion-electron heating. Here, $f_{\text{Edd},\text{crit},\text{visc}}$ is the transitional value between the two thick disc state, and $f_{\text{Edd},\text{crit},\text{thick}} = 0.01$ is the transitional accretion rate which separates thin and thick discs. The radiative efficiency in the viscous heating regime is given by

$$\epsilon_{\text{rad},\text{thick}} = 0.04\epsilon_{\text{rad},\text{thin}} \left(\frac{\delta}{0.1} \right) \left(\frac{1 - \beta}{0.5} \right) \left(\frac{6}{r_{\text{ISCO}}} \right), \quad (\text{D1})$$

and in the ion-electron heating regime by

$$\epsilon_{\text{rad},\text{thick}} = 0.2\epsilon_{\text{rad},\text{thin}} \left(\frac{f_{\text{Edd},H}}{\alpha^2} \right) \left(\frac{\beta}{0.5} \right) \left(\frac{6}{r_{\text{ISCO}}} \right). \quad (\text{D2})$$

Here, β is the ratio of gas pressure and total pressure (which includes the magnetic pressure). A somewhat different parameter, β_m , can be defined as the ratio of gas pressure and magnetic pressure (F. Yuan & R. Narayan 2014). The two parameters are related by $\beta = \beta_m/(1 + \beta_m)$. β_m is not an independent parameter; many simulations have found that $\alpha\beta_m \approx 0.5$ (e.g. F. Yuan & R. Narayan 2014; M. C. Begelman, N. Scepi & J. Dexter 2021), which we adopt. δ represents the fraction of viscous energy dissipation directly transferred to the electrons, and is constrained in theoretical studies between 0.1 and 0.5 (P. Sharma et al. 2007, F. Yuan & R. Narayan 2014). Observations imply a value close to 0.2 (F. Yuan 2003, H. Liu & Q. Wu 2013), which we adopt. The critical Eddington ratio between the two thick disc states can be found by ensuring that both formulae (equations (D1) and (D2)) yield the same radiative efficiency. This gives an accretion rate equal to

$$f_{\text{Edd},\text{crit},\text{visc}} = 0.2 \left(\frac{\delta}{0.1} \right) \left(\frac{1 - \beta}{\beta} \right) \alpha^2, \quad (\text{D3})$$

which is equal to 1.28×10^{-3} for our assumed values of the different parameters.

For the slim disc, we take results based on numerical works (A. Sądowski 2009). We use the following fitting formula adopted for their results by (P. Madau, F. Haardt & M. Dotti 2014):

$$\epsilon_{\text{rad},\text{SD}} = \frac{0.1}{f_{\text{Edd},H}} A(a) \left[\frac{0.985}{1.6/f_{\text{Edd},H} + B(a)} + \frac{0.015}{1.6/f_{\text{Edd},H} + C(a)} \right], \quad (\text{D4})$$

where the three BH spin-dependent functions are given by $A(a) = (0.9663 - 0.9292a)^{-0.5639}$, $B(a) = (4.627 - 4.445a)^{-0.5524}$, and $C(a) = (827.3 - 718.1a)^{-0.7060}$. For accretion that is sufficiently sub-Eddington, this formula yields the radiative efficiency of a thin disc (equation 23). For super-Eddington accretion, the radiative efficiency drops with the Eddington ratio.

APPENDIX E: WARP RADII AND DISC SURFACE DENSITIES

The radius R_{warp} , which separates the inner and outer accretion disc, can be calculated by equating the Lense–Thirring precession time-scale ($t_p = 2\pi/\Omega_p$, with $\Omega_p = 2GJ_{\text{BH}}/c^2R^3$ the precession

rate) and the vertical warp propagation time-scale ($t_{\text{warp}} = R^2/\nu_2$, with ν_2 the kinematic viscosity in the vertical direction) (J. E. Pringle 1992; R. G. Martin, J. E. Pringle & C. A. Tout 2007; S. Cielo et al. 2014). We use the relation $\dot{M} = 3\pi\nu_1\Sigma$ (for $R \gg R_{\text{ISCO}}$, D. Fiacconi et al. 2018) to calculate ν_1 , and therefore ν_2 . For the thick and slim disc, we assume $\nu_2 = \nu_1$. For the thin disc, ν_2 can be related to ν_1 by $\nu_2 = \xi\nu_1$, with ξ a numerical factor (G. Lodato & D. J. Price 2010). The constant parameter ξ is often also expressed in the form α_2/α . Early theoretical calculations predicted $\alpha_2 = 1/2\alpha$ for small α (J. C. B. Papaloizou & J. E. Pringle 1983), which has also been confirmed by simulations (G. Lodato & J. E. Pringle 2007). Later simulations found that higher-order corrections to this prediction may need to be included for realistic values of α (G. Lodato & D. J. Price 2010), such as $\alpha = 0.2$. These numerical results agree with an early theoretical prediction (G. I. Ogilvie 1999) which we assume here:

$$\xi = \frac{\nu_2}{\nu_1} = \frac{\alpha_2}{\alpha} = \frac{2}{\alpha^2} \frac{1 + 7\alpha^2}{4 + \alpha^2}, \quad (\text{E1})$$

which reduces to $1/2\alpha^2$ for small α .

The warp radius, R_{warp} , can be derived for the advection-dominated thick and slim discs as (see G. I. Ogilvie 1999):

$$R_{\text{warp,adv}} = R_G \left(\frac{384|a|}{25(H/R)^2} \right)^{2/5}. \quad (\text{E2})$$

Thus, it depends only on the magnitude of the current BH spin, the aspect ratio of the disc (and the BH mass through R_G). For both the thick and slim disc we assume $H/R = 0.3$, based on general-relativistic magnetohydrodynamical (GRMHD) simulations (R. Narayan et al. 2022; A. Ricarte et al. 2023). This value is somewhat lower than the aspect ratio in the self-similar analytical solutions for these discs ($H/R \approx 0.4\text{--}0.5$ as in R. Narayan & I. Yi 1994; J.-M. Wang & Y.-Y. Zhou 1999, respectively). Regardless of these small differences, equation (E2) yields small warp radii (no larger than $\approx 10R_G$), much smaller than for the thin disc (see below). While we take the aspect ratio H/R from the GRMHD simulations, we assume surface densities based on the analytical solutions since they are more wieldy and easily accessible. These take the same form for the two types of accretion disc:

$$\Sigma_{\text{adv}} = \frac{\dot{M}_{\text{acc,H}}}{2\pi R |\nu_r|}, \quad (\text{E3})$$

where ν_r is the radial inflow velocity of the gas, and the only source of difference between the expression for Σ_{adv} between the two discs. For the slim disc, $\nu_r = -\alpha v_K/\sqrt{5}$ (J.-M. Wang & Y.-Y. Zhou 1999), where $v_K = \sqrt{M_{\text{BH}}G/R}$ is the Keplerian velocity. For the thick disc we take $\nu_r = -\alpha v_0 v_K$ (R. Narayan & I. Yi 1994), where the parameter v_0 is not $1/\sqrt{5} \approx 0.45$ but is instead dependent on α through $v_0 = 3/(5 + 2\varepsilon)$, where $\varepsilon = (5/3 - \gamma)/(\gamma - 1)$, γ is the adiabatic index of the gas, related to the gas-to-total pressure ratio β by $\gamma = (8 - 3\beta)/(6 - 3\beta)$. Finally, β can be connected to α through the findings of GRMHD simulations (F. Yuan & R. Narayan 2014) as $\beta = 1/(1 + 2\alpha)$, which is equivalent to the relation $\alpha\beta_m = 0.5$ mentioned in Appendix D. v_0 depends on α only weakly; for $\alpha = 0.2$, we obtain $v_0 = 0.52$, very close to the value for the slim disc. Having introduced these dependencies, we are also in a position to state the specific angular momentum $L(R)$ for the thick and slim discs. For the slim disc, it is a factor $1/\sqrt{5} \approx 0.45$ of the Keplerian value $L_K = \sqrt{M_{\text{BH}}GR}$, while for the thick disc, it is $\Omega_0 = \sqrt{2\varepsilon/(5 + 2\varepsilon)}$ times the Keplerian value. For $\alpha = 0.2$, $\Omega_0 \approx 0.37$, again close to the slim disc value.

For the thin disc, R_{warp} also depends on the surface density Σ , unlike for the thick and slim discs. The solution for the thin disc (N. I. Shakura & R. A. Sunyaev 1973), assumed in this model, describes three regions: (a) an inner region where radiation pressure dominates the pressure and electron scattering the opacity, which is often unstable and usually does not extend far out, (b) a middle region where gas dominates the pressure and electron scattering dominates the opacity, and (c) an outer region where gas also dominates the pressure, but the opacity is dominated by free-free absorption. Our fiducial choice is to describe the disc using the region (b) from N. I. Shakura & R. A. Sunyaev (1973) (but we found that using region (c) makes little difference). In region (b), the surface density of a thin disc can be expressed as

$$\Sigma_{\text{TD}} = 6.84 \times 10^5 \text{ g cm}^{-2} \alpha^{-4/5} f_{\text{Edd,H}}^{3/5} \left(\frac{\epsilon_{\text{rad}}}{0.1} \right)^{3/5} \left(\frac{M_{\text{BH}}}{10^8 M_{\odot}} \right)^{1/8} \left(\frac{R}{R_S} \right)^{-3/5}. \quad (\text{E4})$$

Using the surface density, the warp radius can be calculated as

$$R_{\text{warp,TD}} = 3410 R_S \alpha^{5/8} \xi^{-5/8} \alpha^{-1/2} f_{\text{Edd,H}}^{-1/4} \left(\frac{\epsilon_{\text{rad}}}{0.1} \right)^{-1/4} \left(\frac{M_{\text{BH}}}{10^8 M_{\odot}} \right)^{1/8} \quad (\text{E5})$$

for region (b) (A. J. Griffin et al. 2019).

At large radii, thin accretion discs can be prone to the effects of self-gravity (see G. Lodato 2007 for a review). The gravity there due to the disc locally becomes comparable to that due to the BH. In order to estimate the radius at which this becomes true, the disc self-gravity radius R_{sg} , we use the Toomre instability parameter, $Q = \Omega c_s/\pi G \Sigma$, and set its value equal to the critical value of 1. This equation, the Toomre instability criterion, be solved for the self-gravity radius, giving

$$R_{\text{sg,TD}} = 6460 R_S \alpha^{28/51} f_{\text{Edd,H}}^{-18/51} \left(\frac{\epsilon_{\text{rad}}}{0.1} \right)^{-18/51} \left(\frac{M_{\text{BH}}}{10^8 M_{\odot}} \right)^{-49/51}. \quad (\text{E6})$$

If $R_{\text{sg}} < R_{\text{warp}}$, we assume that the entire accretion disc is (counter-) aligned and use R_{sg} instead of R_{warp} in all equations where R_{warp} appears.

APPENDIX F: KINETIC JETS AS A REPRESENTATION OF RELATIVISTIC JET FEEDBACK

For use as a feedback mechanism in cosmological simulations of galaxy formation, which have relatively poor resolution, it is not clear how best to represent the effects of relativistic AGN jets. Real jets and the lobes they inflate are affected by many physical processes, including basic (relativistic) hydrodynamics, magnetic fields, instabilities and cosmic ray physics. Many idealized jet simulations have been performed, varying the different physical aspects at play (see the review by M. A. Bourne & H.-Y. K. Yang 2023). While the details change, it appears that the main characteristic of jet physics is that, as the jet material is shocked at the jet head, a hot lobe of material is inflated. The lobe then expands and drives a bow shock that propagates through the ambient medium and heats it anisotropically. This is the main way jets deposit their energy; most or all simulations find that of order half of the injected jet energy is deposited into the medium through the bow shock while the jet is still active. This fraction varies somewhat (in the range 30–70 per cent) with how the jets are simulated. The behaviour of lobes may also differ somewhat depending on whether they constitute nonrelativistic

particles or relativistic ones, on the magnetic fields of the ambient medium, and on how well instabilities are captured in the simulations. However, as at least half of the energy is deposited in the ambient medium independently of the exact contents of the lobes, it appears that for capturing the dominant effects, the precise details of jet and lobe modelling do not matter. The main effects of relativistic AGN jet feedback can thus be represented with a mechanism as simple as non-relativistic kinetic kicks/injections.

The caveat to the above statement is that one needs to resolve the presence of a lobe to capture the driving of a bow shock. With of order dozens or hundreds of particles, we begin to resolve a lobe, so the energy injection should broadly be correct (F. Huško & C. G. Lacey 2023a). With only a dozen or fewer particles representing an entire jet event, this is not true, and instead all of the energy injection into the ambient medium is done through shock heating of individual kicked particles. While this scenario is very common in our simulations, more common than well-resolved jets, this is a resolution limitation that cannot be circumvented by more physical modelling of jets. It instead requires numerical advancements such as adaptive particle splitting/merging (see e.g. M. A. Bourne & D. Sijacki 2020, K.-Y. Su et al. 2021, R. Y. Talbot et al. 2020).

APPENDIX G: THE OBSERVED MEDIAN BLACK HOLE MASS–STELLAR MASS RELATION

We compute the weighted median BH mass as a function of the total stellar mass using the stellar masses and dynamical BH mass measurements from A. W. Graham & N. Sahu (2023), with the stellar masses corrected as per the erratum A. W. Graham & N. Sahu (2024) (a decrease of 0.15 dex relative to A. W. Graham & N. Sahu 2023). The stellar masses from A. W. Graham & N. Sahu (2023) were inferred from deep infrared observations with *Spitzer* using colour-dependent mass-to-light ratios. We confirmed that these are consistent with our own mass-to-light ratios, after converting from their assumed P. Kroupa (2001) to our adopted G. Chabrier (2003) IMF. The dynamical BH mass measurements collected in A. W. Graham & N. Sahu (2023) may be biased towards galaxies of a particular morphological type, at least in the case of some of the methods of BH mass measurement. For this reason, we decided to compute weighted medians that take into account the expected and realized numbers of galaxies of different morphological types in different stellar mass bins.

Our binning procedure consists of the following steps. We compute weighted median BH masses in four bins that each contain ≈ 30 –40 galaxies, with the limits between bins chosen by hand so that no bin has extreme outliers. The weights used for the medians for a given type of galaxy (spheroidal or disk) are equal to $f_{\text{exp}}/f_{\text{obs}}$, where f_{obs} is the actual fraction found for a given galaxy type in a given stellar mass bin in the A. W. Graham & N. Sahu (2023) data, while f_{exp} is the expected fraction in that same stellar mass bin. For this purpose, ellipticals and lenticulars are both considered spheroidal, since they appear to follow the same $M_{\text{BH}}\text{--}M_*$ relation in the A. W. Graham & N. Sahu (2023) data. For f_{exp} , we use the morphological fractions (and their dependence on stellar mass) from A. J. Moffett et al. (2016). The binned stellar and BH masses that we obtain as a result of this procedure are provided in Table G1.

Table G1. The binned median galaxy stellar and BH masses (and their associated errors) from the A. W. Graham & N. Sahu (2023) dataset, corrected for morphological sampling effects. These represent the median BH mass–stellar mass relation for the overall galaxy population, although one that includes Eddington bias, rather than the intrinsic median relation. See the text for a description of the binning procedure and the morphological sampling correction.

$\log_{10} M_* [M_\odot]$	$\log_{10} M_{\text{BH}} [M_\odot]$
10.50 ± 0.03	7.44 ± 0.14
10.85 ± 0.02	7.95 ± 0.10
11.37 ± 0.02	8.85 ± 0.10
11.70 ± 0.03	9.40 ± 0.13

APPENDIX H: SIMPLER RECALIBRATED MODELS

In Figs 7 and 8, we showed the effects of varying the jet velocities and the thick disc accretion efficiencies, respectively. There we also showed models that had constant values of the associated parameters, $v_{\text{jet}} = 10^4 \text{ km s}^{-1}$ and $\epsilon_{\text{acc}} = 1$, respectively. However, showing that such models do not perform as well as our fiducial one when compared against calibration data is not sufficient to rule them out. This is because such parameter variations were not coupled to simultaneous changes of other parameter values. In Fig. H1, we show our best attempts at recalibrating such models. Both of the models were recalibrated by reducing the BH seed masses (since AGN feedback is more effective in lower-mass galaxies/earlier in both runs, this had to be offset by delaying BH growth). The constant-velocity model has a decreased SN pivot density (making SN feedback more effective), while the opposite is true for the constant-efficiency model. That model also has a lower thermal AGN feedback coupling efficiency.

Both recalibrated models perform well when compared against the $z = 0$ GSMF and galaxy size–mass relation (see the top row in Fig. H1). The constant-velocity model better reproduces galaxy sizes at $M_* < 10^{9.5} M_\odot$, and reproduces BH masses and the AGN LF equally well (or even slightly better for the latter) than the fiducial model. Despite these small benefits, the variable-velocity model was chosen for reasons explained in Section 2.4.5. The fact that these models all produce the main observables to a similar degree of accuracy means that these observables cannot be used to discriminate between them. Other galaxy properties, which remain true predictions, must be used.

The model with a constant accretion efficiency falls slightly short on the BH–mass relation at the massive end (see the bottom left in Fig. H1). This could have been remedied by further reducing its coupling efficiency, but note that it already has a very low value of $\epsilon_f = 0.01$. As a result, jet AGN feedback is almost completely dominant over thermal AGN feedback at all masses, and even over SN feedback. The AGN LF (see the bottom right in Fig. H1) is too low, and we were not able to reproduce it with any parameter variations for the model in question. This is indicative of AGN jet feedback from the thick disc accretion state being overly dominant in such a model. Our remedy was the introduction of thick disc accretion efficiencies $\epsilon_{\text{acc}} \ll 1$.




Figure H1. As Figs 7–9, but showing the comparison (in L050 boxes at m7 resolution) between the fiducial hybrid AGN feedback model and two simpler models that, separately, used constant accretion efficiencies in the thick disc (of 100 per cent) or constant jet velocities (of 10^4 km s^{-1}). Both of the simpler models have been recalibrated to reproduce the data shown here as best they can.

APPENDIX I: IMPACT OF PARAMETER VARIATIONS ON THE STELLAR MASS–HALO MASS RELATION

In Section 4.3, we showed the effects of varying different parameters in terms of the impact on two main galaxy observables: the GSMF and galaxy sizes as a function of stellar mass. Differences in the GSMF can be hard to interpret, since changes in stellar mass at fixed halo mass produce horizontal shifts in the GSMF.

To make the impact of these parameter variations on galaxy stellar masses easier to interpret, in Fig. 11 we show the effects of the three parameter variations included in Section 4.3 (the jet velocity normalization, the thick disc accretion efficiency, and the thin disc thermal feedback coupling efficiency) on the stellar mass–halo mass relation for central subhaloes at $z = 0$. We use bins of width 0.3 dex in halo mass. For reference, we also show

the P. Behroozi et al. (2019) semi-empirical model, which reproduces the observed GSMF). For consistency with P. Behroozi et al. (2019), we define the halo mass as the virial mass using the definition from G. L. Bryan & M. L. Norman (1998), which is a fitting function to the results of V. R. Eke, S. Cole & C. S. Frenk (1996).

The left panel of Fig. 11 shows that higher jet velocities lead to a decrease in stellar masses in haloes more massive than $M_{\text{BN}98} \approx 10^{12} M_{\odot}$, likely due to more explosive feedback. In lower-mass haloes, the effect is opposite, likely due to differences in the sampling of AGN feedback. Lower jet velocities, however, also lead to higher stellar masses in low-mass haloes. It thus appears that the fiducial velocity model is close to optimal in terms of the balance between sampling effects and the how efficient (explosive) the feedback is. The middle panel shows that higher accretion efficiencies in the thick disc state lead to lower stellar

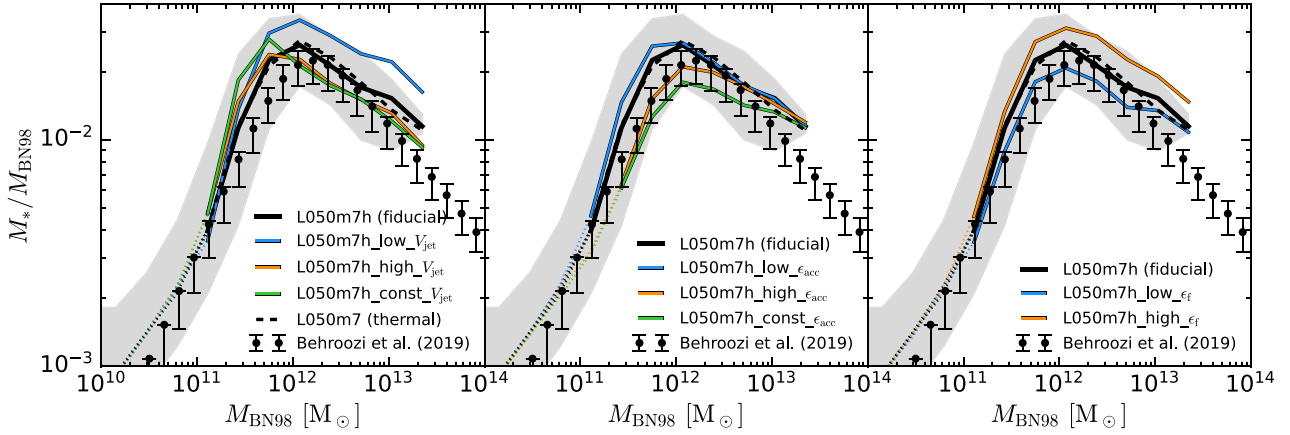


Figure 11. The impact of variations of the jet velocity normalization (left), the thick disc accretion efficiency (middle) and the thin disc thermal coupling efficiency (right) on the stellar mass–halo mass relation for central subhaloes at $z = 0$. The lines show the median from different simulations as given by the legends, with the black shaded region corresponding to the 16–84 percentile range for the fiducial hybrid simulation. The lines become dotted at a halo mass where the median stellar mass is $\leq 10^9 M_\odot$, corresponding to ≈ 100 star particles. The variations shown here are the same as presented in Figs 7–9. For reference, we show the prediction from the P. Behroozi et al. (2019) semi-empirical model. Halo masses follow the G. L. Bryan & M. L. Norman (1998) definition.

masses in all haloes down to $M_{\text{BN98}} \approx 10^{11} M_\odot$, likely due to more effective jet feedback from the thick disc. The right panel shows that higher thin disc thermal coupling efficiencies lead to higher stellar masses in all haloes down to $M_{\text{BN98}} \approx 10^{11} M_\odot$. This is likely a consequence of the coupling efficiency mainly changing BH masses due to self-regulation, which in turn changes the AGN

heating temperatures in those galaxies due to different BH masses (equation 11).

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