THE SPIN EVOLUTION OF MASSIVE BLACK HOLES ACROSS COSMIC TIME

Emily Whitaker

UNIVERSITY COLLEGE CORK
SCHOOL OF PHYSICS

Thesis submitted for the degree of
Master of Science in Physics

August 24th 2025

Supervisor: Dr. Michael Tremmel
Head of School: Prof. Paul Callanan

Contents

“This is to certify that the work I am submitting is my own and has not been submitted for another degree, either at University College Cork or elsewhere. All external references and sources are clearly acknowledged and identified within the contents. I have read and understood the regulations of University College Cork concerning plagiarism and intellectual property.”

Emily Whitaker	24/08/2025
Emily Whitaker	Date

Abstract

Massive black hole spin governs radiative efficiency, jet power, and gravitational-wave signatures, making it a key variable in galaxy-black hole co-evolution. Cosmological simulations, however, cannot resolve disc-scale physics and typically omit spin evolution. Yet, how unresolved accretion geometry–especially during off-centre phases–and unconstrained early growth rates shape cosmic spin histories remains poorly quantified, and observed spin samples are luminosity biased. Here we develop and apply a post-processing spin model that combines the accretion-driven prescriptions Dubois et al. (2014) and Ricarte et al. (2023) to the Romulus25 simulation, evolving both spin magnitude and direction through thin-disk and hot/magnetically arrested disk (MAD) regimes, including spin-dependent radiative and jet efficiencies. We take advantage of Romulus uniquely following realistic massive black hole (MBH) orbits that often wander off-centre to bracket unresolved physics with three off-centre accretion prescriptions (assumed central, chaotic, gas adhesion), and two seed masses ($10^6 M_\odot$ and $10^5 M_\odot$), yielding six models. We find that (i) high spins are built primarily at early times ($z \gtrsim 3$) via prolonged, coherent thin-disk accretion; (ii) declining gas supply and increasing hot-mode duty cycles drive a broad low-spin distribution by $z = 0$; (iii) mergers, occurring between largely aligned progenitors, produce moderate remnant spins and play a secondary role in setting the ensemble distribution; (iv) lowering the seed mass and therefore increasing early effective Eddington ratios increases spin-up, but this imprint fades once black holes grow beyond $\sim 10^{7-8} M_\odot$; and (v) imposing an observational luminosity cut better reproduces the high spins seen in current X-ray samples, revealing strong selection bias relative to the full population. These results imply that present high-spin measurements do not reflect the cosmic mean, but rather a luminosity-weighted

subset with sustained coherent accretion. Strong spin alignment and moderate remnant spins predict low recoil velocities and a Laser Interferometer Space Antenna (LISA) population dominated by weakly precessing binaries. More generally, our off-centre experiments set quantitative bounds on how wandering suppresses spin growth, highlighting the need to model displaced accretion explicitly.

List of Figures

List of Tables

Symbols and Constants

z: redshift

a: dimensionless spin parameter

c: speed of light, $c = 2.99792458^8 \mathrm{m\,s^{-1}}$

J: angular momentum

G: Newton's gravitational constant, $G = 6.67430 \times 10^{-11} \mathrm{m^3\,kg^{-1}\,s^{-1}}$

M: mass

χ_{eff}: effective spin parameter

σ: velocity dispersion

ε_r: radiative efficiency

ε_g: gravitational softening length

r_{ISCO}: the radius of the innermost stable circular orbit around a black hole

n: number density

m_p: mass of a proton, equivalent to $\sim 1.6726 \times 10^{-27}$kg

m_*: star particle mass

p: probability

m_{gas}: gas mass

c^*: star formation efficiency

t_{form}: dynamical formation timescale

ε_{SN}: supernova coupling efficiency

Z: metallicity

n_{th}: star formation threshold number density

T: temperature

ρ: gas density

c_s: speed of sound in a medium

v_θ: rotational relative velocity

v_{bulk}: bulk relative velocity

α: accretion density boost factor

L: luminosity

$\dot{M}$: mass accretion rate

L_{Edd}: Eddington luminosity

$\dot{M}_{Edd}$: Eddington mass accretion rate

ε_{couple}: feedback coupling efficiency

E: Energy

a_{DF}: acceleration due to dynamical friction

M_{BH}: black hole mass

$\rho(< v)$: density of gas moving slower than a velocity v

v: velocity

$\ln\Lambda$: Coulomb logarithm

r: black hole binary separation

$\Delta\mathbf{r}$: relative position vector

$\Delta\mathbf{a}$: relative acceleration vector

$\Delta\mathbf{v}$: relative velocity vector

M_{vir}: galaxy virial mass

R_{vir}: galaxy virial radius

Δ_c: overdensity

ρ_{crit}: critical density of the Universe

H: Hubble Constant at a redshift

σ_T: Thomson cross-section

η: radiative efficiency

f_{Edd}: Eddington fraction; the mass accretion rate as a fraction of the Eddington mass accretion rate

f_{crit}: critical Eddington fraction at which the accretion mode transitions from hot to thin-disk

s: spin-up parameter

s_{HD}: the hydrodynamic component of the spin-up parameter

s_{EM}: the electromagnetic component of the spin-up parameter

l_{ISCO}: specific angular momentum at the innermost circular stable orbit radius

e_{ISCO}: specific energy at the innermost circular stable orbit radius

s_{thin}: spin-up parameter for a thin-disk
l_{thin}: specific angular momentum at the marginally stable orbit radius of a thin-disk
e_{thin}: specific energy at the marginally stable orbit radius of a thin-disk
s_{min}: spin-up parameter for a non-radiative system
k: relative angular frequency of magnetic field lines
Ω_H: angular velocity of the event horizon
r_H: radius of the event horizon
η_{EM}: electromagnetic jet efficiency
ϕ: magnetisation parameter
s_{MAD}: spin-up parameter for a magnetically arrested disk
ΔT: temporal size of a simulation timestep
N_{sub}: number of timesteps
δt: temporal size of integration timesteps
a_T: Thorne spin limit
θ: misalignment angle
α_t: viscosity parameter
ν: viscosity
ℓ: normalised orbital angular momentum of a black hole binary
l: orbital angular momentum of a black hole binary prior to coalescence
L: orbital angular momentum of a black hole binary when widely separated
q: mass ratio
r_{com}: centre of mass position
v_{com}: centre of mass velocity
T_{50}: half-mass assembly time
$T_{Universe}$: age of the Universe
L_{bol}: bolometric luminosity

Units

$M_\odot$: solar mass, equivalent to the mass of the sun
m : meters
s : seconds
kg : kilograms

pc : parsec, equivalent to $\sim 3.0857 \times 10^{16}$m
kpc : kiloparsec
cMpc : comoving Megaparsec
K : Kelvin
yr: years
Gyr: Gigayears
J: Joules
erg: equivalent to 10^{-7}J

Acronyms and Abbreviations

AGN: Active Galactic Nuclei
DF: Dynamical Friction
GRMHD: General Relativistic MagnetoHydroDynamic
GRRMHD: General Relativistic Radiative MagnetoHydroDynamic
ISCO: Innermost Stable Circular Orbit
LISA: Laser Interferometer Space Antenna
MAD: Magnetically Arrested Disk
MBH: Massive Black Hole
ms: marginally stable
RK4: Fourth-order Runge-Kutta
SMBH: SuperMassive Black Hole
SN: Supernova

Acknowledgements

First and foremost, I owe a huge thanks to my supervisor, Dr. Michael Tremmel. Your guidance, and patience have shaped every stage of this project. Thank you for the time you invested in reading drafts, for pushing me to be better, and for the encouragement that made hard problems feel manageable. I could not have asked for a better mentor.

I am also grateful to Dr. Mark Kennedy and Prof. Bertram Bitsch for their support, insightful questions, and many helpful discussions in group meetings and journal clubs. Your feedback improved both the science and the writing.

To my besties and office-mates, Sorcha and Ben: thank you for keeping me sane, for being excellent rubber ducks for my code, for the steady stream of memes when we felt unproductive, and for putting up with all my stressful days throughout this master's. You made the day-to-day fun and the tough days survivable.

To my best friend Caoimhe: thank you for listening to me talk about black holes, often at length, and not always coherently, and for always cheering me on anyway.

Finally, to my parents: thank you for your unwavering support, patience, and encouragement. This thesis exists because you believed in me from the start.

Chapter 1

Introduction

This study focuses on the modelling of MBH spin evolution in cosmological galaxy formation simulations. Numerical frameworks such as ChaNGa (Menon et al. 2015) have previously been employed to investigate black hole growth, spin, and feedback in evolving galaxies (e.g. Tremmel et al. 2017; Beckmann et al. 2025). In this work, a new post-processing model for black hole spin evolution is developed and analysed, combining the spin alignment framework of Dubois et al. (2014) and the accretion-mode-dependent spin-up prescriptions of Ricarte et al. (2023). Applied to black hole populations in the Romulus simulations, this model tracks spin evolution across cosmic time and investigates its relationship with black hole dynamics, accretion environments, and host galaxy properties. This chapter introduces black hole accretion and feedback, and the physical role of MBH spin,and motivates the integration and testing of spin models in a cosmological setting.

1.1 Supermassive Black Holes

Found at the centres of most massive galaxies, supermassive black holes (SMBHs) have masses ranging from a few million to tens of billions of solar masses (Kormendy & Ho 2013; McConnell & Ma 2013). They are believed to form early in the Universe's history and grow over cosmic time through gas accretion and black hole mergers (Volonteri 2010; Inayoshi et al. 2020; Greene et al. 2020).

Black holes are regions of spacetime where gravity is so strong that nothing–not even light–can escape. The boundary beyond which escape becomes impossible is known as the event horizon. Unlike stellar remnants such as white dwarfs or neutron stars, black holes have no internal pressure support to balance gravity; instead, general relativity predicts their formation when a critical density is surpassed (Oppenheimer & Snyder 1939). For SMBHs, the formation pathway is less well understood than for their lower-mass counterparts. They may form from the direct collapse of primordial gas clouds (Loeb & Rasio 1994; Bromm & Loeb 2003), the collapse of massive Population III stars (Madau & Rees 2001), or the successive mergers of smaller black holes (Devecchi & Volonteri 2009).

SMBHs were first hypothesised in the 1960s to explain the extreme luminosities of quasars (Salpeter 1964; Lynden-Bell 1969), because accretion onto a $\gtrsim 10^{6-9} M_\odot$ compact object can convert a large fraction of rest-mass energy into radiation, easily producing the observed quasar luminosities and far exceeding what stellar processes could supply. Observational evidence for their existence has since accumulated from a variety of sources, including high-resolution stellar dynamics near galactic centres (Ghez et al. 2008; Gillessen et al. 2009), the scaling relations between black hole mass and host galaxy properties (Magorrian et al. 1998; Ferrarese & Merritt 2000), and the detection of active galactic nuclei (AGN) powered by accreting SMBHs (Rees 1984; Kormendy & Richstone 1995). Most recently, the Event Horizon Telescope provided direct imaging of the shadow of a SMBH in M87 (Event Horizon Telescope Collaboration et al. 2019).

Since their discovery, SMBHs have become a central topic in astrophysics because the energy they release during active phases is capable of heating, expelling, or preventing the cooling of galactic gas. In practice, AGN feedback is now a necessary ingredient in galaxy-formation models: without it, cosmological simulations vastly overproduce stars, fail to quench massive galaxies, and cannot reproduce the observed stellar mass function, colour bimodality, or hot intracluster gas properties (Silk & Rees 1998; Springel et al. 2005; Fabian 2012; Schaye et al. 2015; Weinberger et al. 2017). Thus SMBHs are not merely by-products of galaxy assembly, but key regulators that couple small-scale accretion physics to galaxy-wide evolution. Understanding how SMBHs form, evolve, and interact with their environment remains one of the key challenges in theoretical and observational astrophysics.

1.2 Massive Black Hole Mergers

MBH mergers are key events in the hierarchical formation of structure in the Universe. As galaxies collide and merge, their central black holes are expected to sink to the centre of the newly formed system via dynamical friction, eventually forming a bound binary and, in many cases, coalescing into a single, more massive, black hole (Begelman et al. 1980; Chapon et al. 2013; Volonteri et al. 2003). These mergers not only grow MBHs but also significantly influence the evolution of their host galaxies by imparting gravitational-wave recoil kicks that can displace or even eject the remnant, altering where and how AGN feedback is deposited, by driving powerful, merger-triggered accretion bursts and jet reorientations that inject energy into the surrounding gas, and by extracting angular momentum from stars and circumbinary gas via three-body scattering and disc torques, which both evacuate nuclear cores and funnel fresh gas inward (Blecha et al. 2011; Kelley et al. 2017; Springel et al. 2005; Bogdanovic et al. 2007).

The final stages of MBH mergers emit gravitational waves in the millihertz regime–frequencies that are inaccessible to ground-based detectors like LIGO and Virgo, but ideal for space-based observatories. LISA, scheduled for launch in the mid-2030s, is designed to detect these low-frequency gravitational waves (Group et al. 2023). LISA will open a new observational window on MBH binaries across a wide range of masses $\left(\sim 10^4 - 10^7 M_\odot\right)$ and redshifts (up to and beyond $z \sim 20$), telling about their dynamical evolution and cosmological assembly history (Sesana et al. 2005; Klein et al. 2016; Barausse et al. 2020). Its expected sensitivity across mass and redshift ranges–illustrated in Figure 1.1, from Colpi et al. (2019)–highlights LISA's unique capability to detect mergers throughout the Universe, including MBH binaries and some extreme mass-ratio inspirals.

LISA detections will provide unprecedented insights into MBH demographics, including their mass and spin distributions, merger rates, and formation pathways (Katz et al. 2020). In particular, the pre-merger spin orientations of MBHs–driven by processes like gas accretion and torques–can leave detectable imprints on the gravitational waveform (Gerosa & Berti 2017). This offers a powerful method for testing theoretical models of MBH spin evolution and their connection to galactic environments (Bogdanovic et al. 2007; Ricarte & Natarajan 2018; Schnittman 2011).

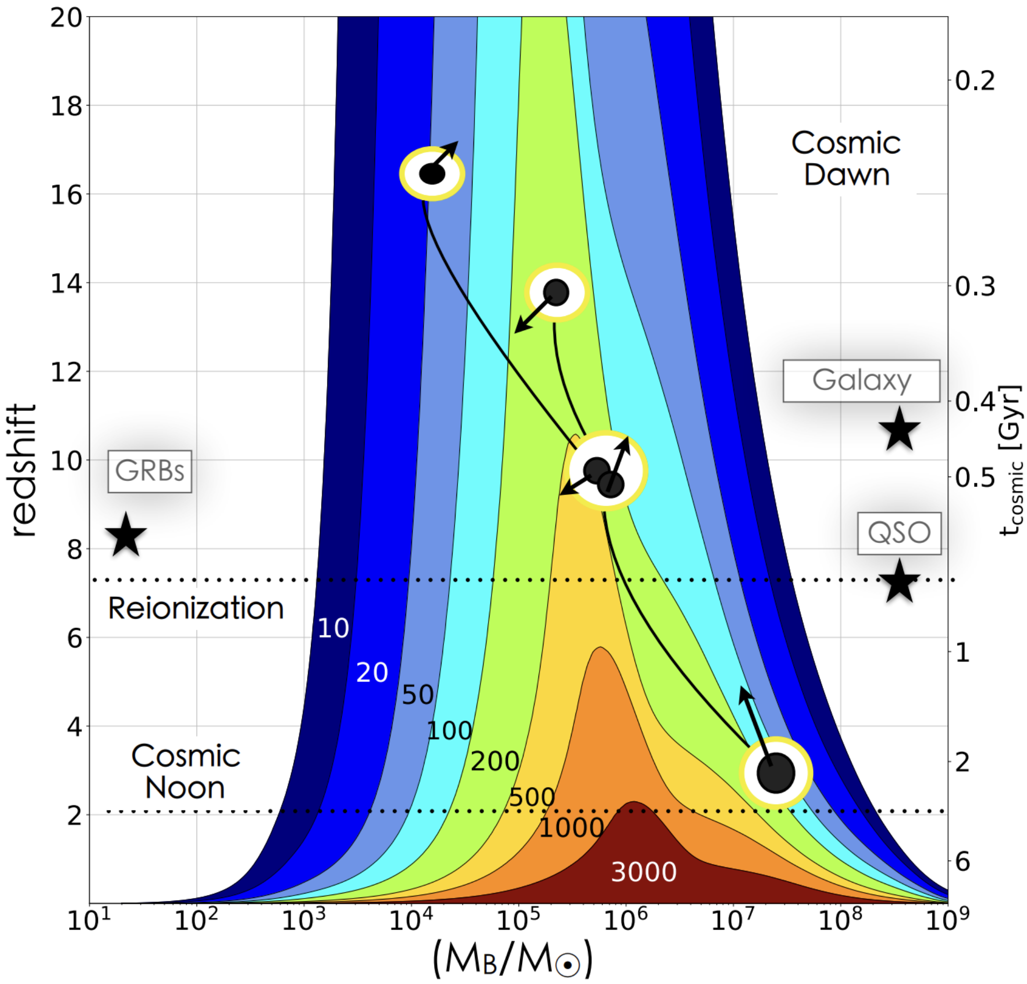

Figure 1.1: *LISA Sensitivity. Signal-to-noise ratio contours as a function of redshift/cosmic time and binary mass. Overplot is an example of a black hole binary evolution and coalescence (black circles and lines) and the most distant long Gamma-Ray Burst host, quasar and galaxy detected so far (black stars). Taken from Colpi et al. (2019).*

Understanding MBH mergers is therefore not only essential for charting black hole growth over cosmic time, but also for interpreting LISA's future discoveries.

1.3 Black Hole Accretion

Accretion is one of the primary processes through which black holes gain mass and angular momentum. As gas falls into a black hole's gravitational potential, it forms an accretion flow whose structure and radiative properties depend on the mass accretion rate and the black hole's environment (Shakura & Sunyaev 1973; Narayan & Yi 1994; Yuan & Narayan 2014). Accretion not only governs black hole growth, but also plays a central role in regulating feedback processes that influence galaxy evolution (Fabian 2012; Kormendy & Ho 2013).

Black holes have a theoretical maximum steady accretion rate set by the balance between gravity and radiation pressure–the Eddington limit. The Eddington luminosity is

$$L_{Edd} = 1.26 \times 10^{38} \frac{M_{BH}}{M_\odot} \mathrm{erg\,s^{-1}}$$

for Thomson scattering in ionised hydrogen. The corresponding accretion rate is

$$\dot{M}_{Edd} = \frac{L_{Edd}}{\eta c^2}$$

where η is the radiative efficiency. Accretion significantly above this limit is expected to become geometrically thick and radiatively inefficient ("super-Eddington") owing to photon trapping and powerful outflows.

At high accretion rates–typically above a few percent of the Eddington limit–black holes are thought to accrete via a geometrically thin, optically thick disk, often referred to as the thin disk regime. In this state, gas efficiently cools via radiation, allowing the disk to remain thin and to radiate a large fraction of the gravitational potential energy as electromagnetic emission (Shakura & Sunyaev 1973). This regime is associated with quasar activity and is highly radiatively efficient, typically yielding strong observational signatures in X-ray and UV bands (Hopkins et al. 2006). Thin-disk accretion is also the dominant regime responsible for spinning up black holes, as the angular momentum of the accreted gas is efficiently transferred to the black hole (Bardeen 1970; Volonteri et al. 2005). An example of such a structure can be seen in Figure 1.2, which shows a

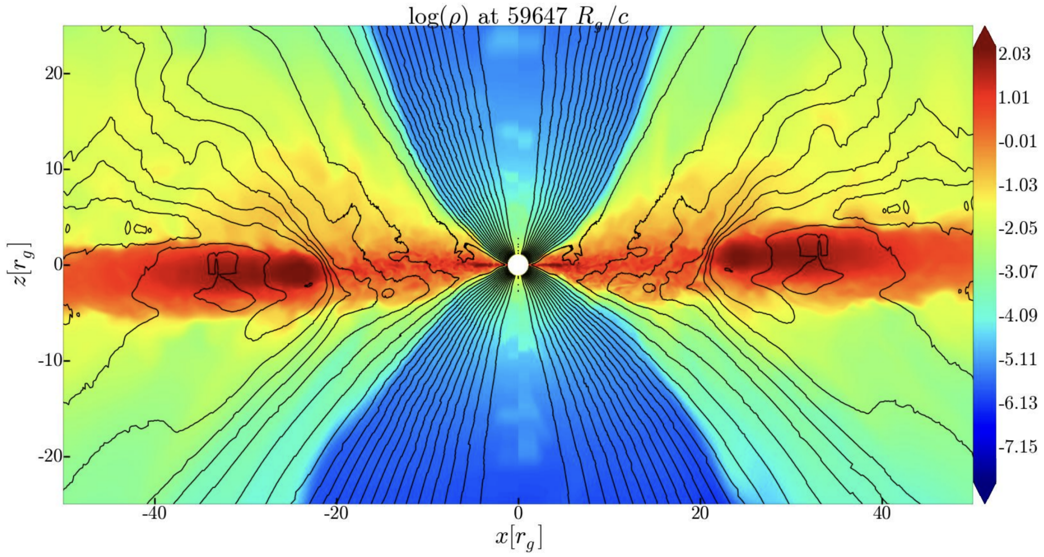

Figure 1.2: *Vertical slice through density in a 3D general relativistic magnetohydrodynamic simulation of a thin disc. Taken from Liska et al. (2019)*

high-resolution general relativistic magnetohydrodynamic (GRMHD) simulation of a thin, magnetised accretion disk launching relativistic jets (Liska et al. 2019).

In contrast, at low accretion rates–well below the Eddington limit–black holes are believed to enter the hot accretion regime, also known as the advection-dominated accretion flow regime (Narayan & Yi 1994; Yuan & Narayan 2014). In this case, the accretion flow becomes geometrically thick and optically thin, and much of the energy is advected into the black hole rather than radiated away. This regime is typically associated with low-luminosity AGN and the quiescent states of MBHs (Ho 2009). While hot accretion flows are inefficient at growing black hole mass or spinning them up, they are thought to be efficient at launching jets and driving mechanical feedback into the host galaxy (Tchekhovskoy et al. 2011; Yuan et al. 2015). Figure 1.3 illustrates this hot accretion flow regime with a simulation showing the vertically extended, turbulent structure of the flow (Yuan & Narayan 2014).

In cosmological simulations, accretion regimes are often distinguished based on the Eddington ratio (the accretion rate relative to the Eddington rate). Many subgrid models (i.e. analytic, resolution-dependent prescriptions that stand in for unresolved physics below the grid/particle scale) implement separate feedback

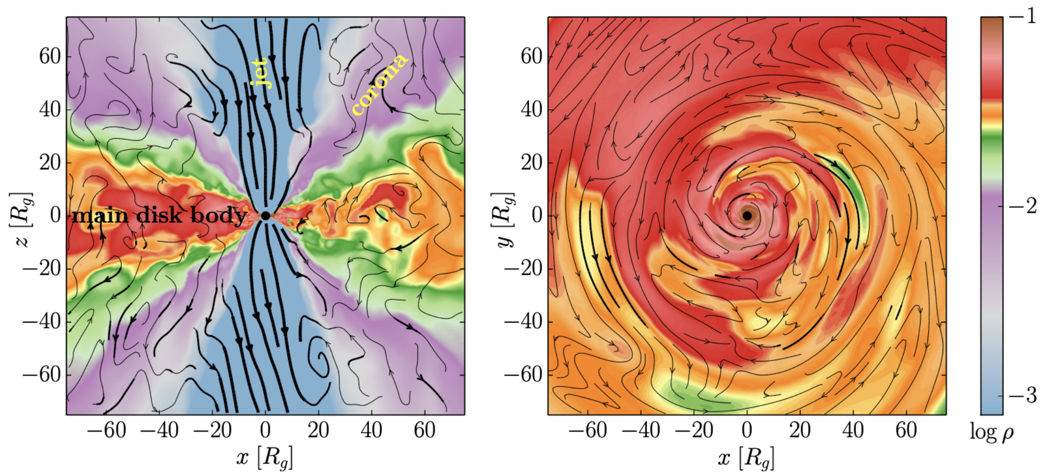

Figure 1.3: *A poloidal slice at a fixed azimuthal angle (left) and a slice through the equatorial plane (right) from a 3D general relativistic magnetohydrodynamic simulation of a hot accretion flow around a spinning black hole. Taken from Yuan & Narayan (2014)*

modes for thin-disk and hot accretion flows to capture the different ways in which black holes interact with their surroundings (Sijacki et al. 2007; Dubois et al. 2012; Weinberger et al. 2017). Accurately modelling these regimes is essential to understanding the growth of black holes over cosmic time and the co-evolution of black holes and their host galaxies.

1.4 Black Hole Spin

The spin of a black hole is a fundamental parameter that encodes information about its formation history and interaction with its environment. For astrophysical black holes, spin is typically described by the dimensionless spin parameter $a = \frac{cJ}{GM^2}$, where J is the angular momentum, M is the mass, G is the gravitational constant, and c is the speed of light. This parameter ranges from 0 (non-rotating Schwarzschild black hole) to 1 (maximally spinning Kerr black hole).

Black hole spin evolves primarily through two channels: gas accretion and black hole mergers. In gas-rich environments, prolonged accretion of material with coherent angular momentum can efficiently spin up a black hole to high values (Volonteri et al. 2005; Berti & Volonteri 2008). In contrast, accretion via randomly oriented episodes—as might occur in a chaotic or clumpy interstellar medium—can lead to moderate or even reduced spin (King et al. 2008; Dotti et al. 2012). Mergers between black holes also influence spin, with the final spin depending on the mass ratio, initial spins, and orientations of the merging pair (Rezzolla et al. 2008; Barausse & Rezzolla 2009).

Spin measurements, although challenging, are becoming increasingly accessible. Spin leaves a measurable imprint on both the structure of the accretion disc and the photon or gravitational-wave signatures we observe. In the electromagnetic band, two complementary techniques are widely used.

Relativistic reflection spectroscopy models the broad, skewed Fe-Kα line produced when hard X-rays irradiate the inner disc. The degree of red-wing extension depends on how close the disc extends to the black hole, which is assumed to be the radius of the innermost stable circular orbit (ISCO). Because the ISCO shrinks from $6GMc^2$ for a non-spinning (Schwarzschild) black hole hole to $1.24GMc^2$ for a maximally prograde Kerr black hole ($a \sim 0.998$), measuring the red wing constrains a (Reynolds 2014).

Continuum fitting ; In thermal-dominant states of stellar-mass systems (and, in principle, AGN with reliable inclination and distance), continuum fitting compares the multi-colour blackbody disc spectrum to relativistic thin-disc models; the inner edge temperature and normalisation again encode the ISCO radius and thus the spin (Miller 2007; Bambi et al. 2021).

Gravitational-wave signals offer an independent avenue. During the inspiral, spin enters the waveform phase evolution through the effective spin parameter χ_{eff} and, if misaligned, through spin-orbit precession; in the merger-ringdown, the quasi-normal mode frequencies depend only on the remnant's mass and spin. LIGO/Virgo have already used these effects to bound spins of stellar-mass black holes (Abbott et al. 2016; Abbott et al. 2019), and LISA will extend the method to MBH binaries across much of cosmic history (Klein et al. 2016).

1.5 Feedback

As matter spirals inwards via an accretion disk, gravitational potential energy is converted into thermal and kinetic energy. A significant fraction of this energy is radiated away, with the remainder contributing to the black hole's mass and angular momentum (Shakura & Sunyaev 1973; Thorne 1974). The radiative efficiency of this process depends strongly on spin: for a non-rotating Schwarzschild black hole, the maximum efficiency is $\sim 6\%$, while for a maximally spinning Kerr black hole, it can reach up to $\sim 42\%$ (Reynolds & Nowak 2003). Thus, spin not only influences how black holes gain energy, but also how much energy is retained versus radiated.

Importantly, spinning black holes can also transfer energy back into their surroundings, profoundly influencing their galactic environment. One of the most powerful mechanisms for this is the Blandford-Znajek process, in which magnetic fields anchored in the surrounding accretion disk tap into the rotational energy of the black hole. These magnetic field lines, threading the ergosphere, can extract angular momentum and launch relativistic jets from the polar regions of the black hole (Blandford & Znajek 1977; Komissarov 2001). Numerical GRMHD simulations have shown that jet power can exceed the accretion power under certain conditions, especially in so-called MADs, where magnetic flux accumulation near the event horizon suppresses accretion while driving powerful outflows (Narayan et al. 2003; Tchekhovskoy et al. 2011; McKinney et al. 2012).

These jets can transport energy on kiloparsec or even megaparsec scales, inflating radio lobes and cavities in the surrounding interstellar or intracluster medium. Such radio-mode feedback is most prevalent in massive galaxies and galaxy clusters, where it heats or displaces cold gas, thereby regulating star formation and contributing to the quenching of galaxies (McNamara & Nulsen 2007; Fabian 2012; Gaspari et al. 2020).

In addition to jets, black holes also drive radiative or quasar-mode feedback, especially during periods of rapid accretion near the Eddington limit. In this regime, intense radiation pressure from the accretion disk can drive wide-angle winds and outflows that interact with the host galaxy's gas, potentially expelling it and limiting further star formation (Springel et al. 2005; Hopkins et al. 2006; Costa et al. 2014). These winds can be highly ionised (e.g., ultra-fast outflows seen in

X-ray spectra) or molecular in nature, affecting both the immediate nuclear region and the galaxy on kiloparsec scales (Feruglio et al. 2010; Tombesi et al. 2015; Fiore et al. 2017).

Together, these feedback processes are essential for explaining the observed co-evolution of black holes and their host galaxies. They provide a natural explanation for the empirical scaling relations between black hole mass and galaxy properties (e.g., $M_{BH} - \sigma$ relation; Ferrarese & Merritt 2000; Gebhardt et al. 2000) and are routinely implemented in cosmological simulations of galaxy formation(Springel et al. 2005; Weinberger et al. 2017; Dubois et al. 2021).

1.6 Black Holes in Simulations

Cosmological simulations cannot resolve the physical scales on which MBHs actually form and accrete (sub-parsec discs inside kpc-scale gravitational softenings). They therefore introduce black holes through sub-grid seeding prescriptions: either by placing a seed of fixed mass in every halo above a threshold halo mass or velocity (e.g. Illustris / TNG, EAGLE), or by spawning seeds when gas satisfies low-metallicity, dense, moderately warm conditions suggestive of direct-collapse sites (e.g. Romulus) (e.g. Springel et al. 2005; Schaye et al. 2015; Tremmel et al. 2017). Seed masses (typically $10^4 - 10^6 M_\odot$) and criteria are tuned so that subsequent growth can reproduce observed black hole-galaxy scaling relations at low redshift.

Because the Bondi radius and disc angular momentum transport are unresolved, accretion rates are estimated with modified Bondi-Hoyle formulas (often boosted to mimic unresolved cold clumping) or angular-momentum/torque-limited prescriptions; an Eddington cap is usually imposed, sometimes with allowance for mild super-Eddington episodes (Booth & Schaye 2009; Rosas-Guevara et al. 2015; Anglés-Alcázar et al. 2017). A fixed radiative efficiency (commonly $\varepsilon_r \sim 0.1$) is often assumed. A fraction of the inferred accretion energy is returned as feedback, either in a single thermal mode or in dual quasar/low-state (kinetic or jet) modes to regulate star formation and gas supply (Sijacki et al. 2007; Weinberger et al. 2017; Davé et al. 2019).

Limited force resolution also affects black hole dynamics. Many large-volume simulations forcibly reposition black hole particles to the local gravitational potential minimum at each timestep, in order to approximate the effects of rapid sinking toward galaxy centres, while others apply a sub-grid dynamical friction force so that wandering and delayed pairing occur more naturally (Tremmel et al. 2015). Mergers are declared when two black holes become closely separated (within a few softenings) and bound; the unresolved hardening and gravitational-wave inspiral phases are collapsed into an instantaneous coalescence. Most simulations do not evolve black hole spin; when spin is considered it is typically in post-processing or simplified alignment treatments (Dubois et al. 2014; Ricarte et al. 2023).

Key limitations common to all approaches include unresolved multiphase ISM structure (driving uncertain boost factors), degeneracies between seeding, accretion, and feedback efficiencies, omission or simplification of magnetic fields and radiation transport (limiting jet/spin modelling), and unresolved binary hardening timescales that influence predicted gravitational-wave event rates. Despite these, contemporary suites broadly match global galaxy statistics and black hole scaling relations, providing a physically informed–though approximate–context for more detailed, higher-resolution or post-processed studies of black hole growth and spin.

1.7 Black Hole Spin in Simulations

Large scale cosmological simulations cannot resolve the sub-parsec accretion flows that set black hole spin. As a result, two complementary families of approaches have emerged. In the first, post-processing methods reconstruct spins from simulated back hole growth histories using analytic prescriptions for thin-disk spin-up, hot/MAD spin-down, alignment torques, and merger remnant spins from numerical relativity fits. In the second, on-the-fly models evolve spin magnitude and direction during the simulation, allowing spin to feed back into radiative and jet efficiencies and, in some cases, to couple to subgrid alignment torques.

Early post-processed studies demonstrated that coherent, prolonged thin-disk accretion can rapidly spin up black holes, while chaotic or intermittent fuelling drives spins to moderate or low values; they also provided population predictions for remnant spins after mergers (e.g. Dubois et al. 2014). More recent post-processing works tailored to modern large-volume runs refined these ideas, showing how assumptions about accretion coherence, time averaging, and merger alignment shape the high-spin tail and its mass/redshift dependence (e.g. Bustamante & Springel 2019). Another line of work calibrated spin-up/down laws to GRMHD simulations and applied them to cosmological black hole histories, separating hydrodynamic spin-up from electromagnetic spin-down and distinguishing hot versus thin-disk regimes (e.g. Ricarte et al. 2025). Together, these studies established a practical toolkit for reconstructing spins from existing simulations, at the cost of neglecting real-time spin-feedback couplings.

Live, on-the-fly implementations co-evolve black hole spin with other properties like black hole mass, often including spin-dependent radiative/jet efficiencies and approximate alignment torques. Recent examples in cosmological volumes report systematically higher spins at early times, faster spin-up of low-mass black holes, and clearer links between accretion state and spin demographics (e.g. Sala et al. 2024; Beckmann et al. 2025). These on-the-fly models therefore yield self-consistent predictions (i.e. feedback reacts to spin), whereas post-processing uses fixed efficiencies from the parent run; the latter is cheaper and ideal for sweeping parameter space.

Spin-galaxy alignment offers a further point of comparison. Post-processed analyses relate black hole spins to large-scale gas angular momentum and merger geometry, predicting predominantly aligned configurations when coherent fuelling dominates (e.g. Beckmann et al. 2024). Live-spin studies, which can track real-time torques and disc reorientations, similarly find strong alignment between black hole spin and host angular momentum in settled, star-forming systems, with misalignments emerging during disturbed phases (e.g. Peirani et al. 2024). These complementary results indicate that the central galactic gas provides a robust reference axis for spin evolution on kiloparsec scales, while the duty cycle of coherent versus chaotic episodes controls the residual scatter.

Within this landscape, the present work contributes in two ways. First, it follows the post–processing paradigm but employs GRMHD-informed, spin-dependent efficiencies and distinct prescriptions for hot versus thin-disk accretion. Second–and most importantly–it explicitly treats off-centre accretion by wandering black holes through three contrasting prescriptions, quantifying how displaced growth episodes alter spin magnitudes, merger alignments, and observable spin-mass-Eddington correlations. Applying this framework to Romulus is particularly powerful: Romulus is one of the few large-volume simulations that simultaneously resolves dwarf galaxies and low-mass hosts at early times, predicts MBH populations emerging in these environments, and models black hole orbital evolution with a physically motivated dynamical-friction scheme (rather than ad-hoc recentering), yielding realistic wandering and merger timescales (Tremmel et al. 2015, 2017). The combination of a new off-center accretion treatment with a simulation that captures long-lived displacement and accurate pairing delays allows us to test competing ideas about off-center fuelling and its impact on spin in a way that has not been feasible in prior analyses. This explicit treatment of angular momentum acquisition during wandering phases highlights a facet of black hole growth that is often simplified or ignored, and it provides a bridge between post-processed reconstructions and live-spin models that will be essential for interpreting forthcoming spin measurements.

In this thesis we seek to understand how MBHs acquire and retain spin across cosmic time, and how that spin encodes their growth histories. To that end, in the following chapters, we develop a post-processing framework that combines physically motivated prescriptions for accretion-driven spin evolution and merger remnants, and, crucially, explicitly treats angular-momentum acquisition during off-centre phases. We apply this model to the Romulus25 simulation to (i) evolve spin magnitude and direction through thin-disk and hot regimes with spin-dependent radiative/jet efficiencies, (ii) test three contrasting off-centre accretion prescriptions alongside two seed masses to bracket unresolved physics, (iii) quantify spin-mass-redshift trends and spin-orbit/host alignments, and (iv) assess observational selection effects by imposing luminosity cuts relevant to current spin measurements. Together these objectives aim to place simulation-anchored, testable constraints on the coherence of black-hole growth–particularly while wandering–and to provide clear targets for future, fully self-consistent magneto hydrodynamic simulations and deeper observational surveys.

Chapter 2

Romulus25 Simulation

The Romulus25 simulation is a uniform, periodic, cosmological volume with a comoving side length of 25 cMpc, evolved from $z = 99$ to $z = 0$ with the ChaNGa TreeSPH code, producing thousands of galaxies from dwarfs to systems exceeding the Milky Way in mass; gravity is spline-softened using a gravitational softening length $\varepsilon_g = 350$pc (fully Newtonian beyond ~ 700pc) and hydrodynamics are solved with 32 neighbour SPH smoothing using equal-mass gas and star particles of $2.12 \times 10^5 \, M_\odot$ and dark matter particles of $3.39 \times 10^5 \, M_\odot$; the adopted flat ΛCDM cosmology (Planck Collaboration et al. 2014) specifies $\Omega_m = 0.3086$, $\Omega_\Lambda = 0.6914$, $h = 0.6777$, and $\sigma_8 = 0.8288$. All numerical choices and subgrid prescriptions described in this section are native to Romulus25; I analyse the resulting snapshots/derived database and did not run or modify the cosmological simulation.

All per-snapshot galaxy and SMBH properties (positions, masses, angular momentum proxies, accretion rates, merger flags) are provided in a TANGOS database (Python/SQLAlchemy interface) which pre-computes and stores values in an easily queryable SQLite schema, streamlining cross-snapshot linkage and enabling efficient construction of time series for subsequent analysis.

2.1 Gravity and Force Softening

Gravity is computed with the ChaNGa Barnes-Hut tree, using an opening criterion tuned to balance accuracy and performance in clustered regions; long-range forces are evaluated hierarchically while short-range interactions are directly summed. A spline (cubic) gravitational softening with softening length $\varepsilon_g = 350$pc is applied to dark matter, star, gas, and black hole particles. Forces are fully Newtonian for pair separations $\gtrsim 2\,\varepsilon_g = 700$pc.

The chosen softening suppresses artificial two-body scattering (particularly between massive dark matter particles and lighter stellar/gas particles), stabilises galactic potentials, and yields converged circular velocity profiles outside a few softening lengths. It also ensures that black hole orbital decay (handled via a subgrid dynamical friction model) is not dominated by discreteness noise at small radii. Because softening lengths are greater than the scale of accretion disks for any SMBH, sub-parsec disc and spin physics are unresolved and are treated analytically elsewhere; however, kiloparsec-scale potential gradients and angular momentum on which the spin post-processing relies are well resolved above the Newtonian transition radius.

2.2 Gas Cooling, Star Formation, and Stellar Feedback

Romulus25 adopts the gas-cooling implementation of Guedes et al. (2011), incorporating Compton cooling off the cosmic microwave background, primordial (H/He) atomic cooling, and metallicity-dependent low-temperature radiative cooling following Mashchenko et al. (2006). Molecular (H_2) cooling is not modelled, because the simulation's spatial and mass resolution do not capture the densities, shielding, and small-scale structure of individual star-forming clouds where H_2 would form and dominate the cooling budget.

Gas becomes eligible for star formation once it is cool and dense, specifically when its temperature falls below 10^4K and its number density exceeds $n_{gas} > 0.2\mathrm{m_p}\,\mathrm{cm}^{-3}$ (with m_p the proton mass). Star formation is modelled stochastically:

during a timestep Δt, an eligible gas particle spawns a star particle (target mass m_*) with probability

$$p = \frac{m_{gas}}{m_*}\left(1 - e^{c^* \frac{\Delta t}{t_{form}}}\right)$$

where $c^* = 0.15$ is the star formation efficiency, t_{form} is the local dynamical timescale, and $\Delta t = 10^6$yr is the star formation timescale. Newly created star particles retain 30% of the parent gas particle's mass (consistent with a Kroupa initial mass function; Kroupa 2001), with the remaining mass left in the gas reservoir or redistributed per the implementation. This probabilistic scheme reproduces a Schmidt-Kennicutt-like relation in aggregate while preserving mass and energy resolution constraints.

Feedback from massive stars is implemented with a blastwave (delayed cooling) scheme: each Type II SN deposits 10^{51} erg with an effective coupling efficiency $\varepsilon_{SN} = 0.75$, temporarily disabling radiative losses in the affected gas to ensure efficient energy transfer. Type Ia SNe and stellar winds contribute additional thermal energy and metals on longer timescales, and so cooling is not disabled in this case. This feedback regulates central gas densities, drives galactic fountains, and self-consistently produces a multiphase ISM without enforcing a subgrid two-phase equation of state.

Together, the metallicity-dependent cooling, stochastic star formation, and energetic stellar feedback establish a cycling cold-hot ISM and circumgalactic medium while preserving coherent kiloparsec-scale rotation in settled discs.

2.3 Black Hole Seeding

SMBHs are introduced via a gas property seeding prescription rather than a halo mass threshold (Tremmel et al. 2017). A gas particle is eligible to spawn a black hole if simultaneously: its metallicity satisfies $Z < 3 \times 10^{-4}$, its density exceeds $n_H > 15 n_{th}$ (with the star formation threshold $n_{th} = 0.2$ cm^{-3}), and its temperature lies within 9500 K $< T <$ 10,000 K, selecting dense, moderately warm, low-metallicity gas consistent with a heavy seed formation channel. When these criteria are met a seed of mass $10^6 M_\odot$ is created (transferring the requisite mass from the parent gas particle and, if necessary, neighbouring gas to conserve mass). This value was adopted in Romulus25 because it is large enough to exceed the gas/star particle mass by a comfortable factor, reducing numerical wandering and ensuring the subgrid dynamical-friction model operates in the appropriate mass regime, and is consistent with "heavy-seed" (direct-collapse) scenarios proposed for early SMBH formation.

This method yields an emergent occupation fraction that reflects local thermodynamic conditions rather than imposing one through a deterministic halo cut, allowing early formation in suitable pockets of pre-enriched structure while avoiding ubiquitous over-seeding in low-mass haloes. The adopted approach thus seeds black holes preferentially in dense, still chemically primitive regions likely to host rapid initial growth, while remaining agnostic about subsequent halo assembly pathways.

2.4 Black Hole Accretion and Feedback

Accretion onto SMBHs follows the rotation-aware, modified Bondi prescription of Tremmel et al. (2017). The instantaneous accretion rate is

$$\dot{M} = \alpha \times \begin{cases} \frac{\pi G^2 M^2 \rho}{\left(v_{bulk}^2 + c_s^2\right)^{\frac{3}{2}}}, & \text{if } v_{bulk} > v_\theta \\ \frac{\pi G^2 M^2 \rho c_s}{\left(v_\theta^2 + c_s^2\right)^2}, & \text{if } v_{bulk} < v_\theta \end{cases} \tag{2.1}$$

where ρ and c_s are the local gas density and sound speed. If rotational support is weak ($v_\theta < v_{bulk}$) then the bulk relative velocity v_{bulk} is used; if rotation dominates ($v_\theta > v_{bulk}$) then the rotational relative velocity v_θ is used, suppressing accretion relative to a naive Bondi estimate. The density boost factor α captures unresolved multiphase structure: $\alpha = \left(\frac{n}{n_{th}}\right)^\beta$ for $n > n_{th}$ with $\beta = 2$, and $\alpha = 1$ otherwise, where n_{th} is the star formation threshold density.

An Eddington limit is enforced by comparing the implied luminosity $L = \varepsilon_r \dot{M} c^2$ (with radiative efficiency $\varepsilon_r = 0.1$) to $L_{Edd} = 1.26 \times 10^{38} \left(\frac{M_{BH}}{M_\odot}\right) \mathrm{erg\,s^{-1}}$; if $L > L_{Edd}$ the accretion rate is reduced to $\dot{M}_{Edd} = \frac{L_{Edd}}{\varepsilon_r c^2}$. Feedback couples a fixed fraction $\varepsilon_{couple} = 0.02$ of the radiative output thermally to neighbouring gas particles, depositing $\dot{E} = \varepsilon_{couple} \varepsilon_r \dot{M} c^2$ in a kernel-weighted fashion. This continuous thermal input regulates central gas densities and moderates subsequent accretion.

Operationally, the prescription is piecewise: it switches from a bulk-dominated form to a rotation-limited form at $v_\theta = v_{bulk}$. As a result, the accretion rate is not strictly continuous at this boundary. This piecewise implementation nevertheless captures the intended suppression of inflow when rotational support dominates, while retaining the density-boost and fixed-efficiency thermal feedback terms.

2.5 Black Hole Dynamics

MBH orbits are evolved self-consistently; no artificial repositioning to galactic centres is applied. Instead, unresolved deceleration from stars and dark matter is modelled with a sub-grid Chandrasekhar dynamical friction (DF) prescription (Tremmel et al. 2015, 2017). For a black hole of mass M_{BH} moving at speed v relative to the local collisionless background of (kernel-weighted) density ρ, the DF acceleration is

$$\mathbf{a}_{DF} = -4\pi G^2 M_{BH} \rho(< v) \ln\Lambda \frac{\mathbf{v}}{v^3} \tag{2.2}$$

with Coulomb logarithm $\ln\Lambda$ limited by the gravitational softening at small scales and the local maximum impact parameter at large scales. Collisionless (stellar + dark matter) particles dominate over gas particles in this calculation (Chen et al. 2022). The required local moments (ρ, $\mathbf{v}$) are estimated from a neighbour list (typical size comparable to the SPH neighbour number), and the DF term is added explicitly to the black hole's acceleration each timestep.

This treatment allows black hole to experience realistic orbital decay timescales, including long-lived off-centre ("wandering") phases in low-density or spheroidal hosts, while preventing spurious numerical scattering by individual massive particles. Resulting displacement distributions underpin later analyses that differentiate central and off-centre accretion behaviour. Merger candidates (Section 2.6) emerge naturally as orbital energy and angular momentum are dissipated by DF and gas dynamical processes rather than by enforced positional resets.

2.6 Black Hole Mergers

A merger is registered when two black holes satisfy both a spatial and a dynamical binding criterion at a given timestep. First, their separation must fall below a fixed proximity threshold (twice the gravitational softening length, $r < 2\varepsilon$). Second, they must be gravitationally bound according to

$$\frac{1}{2}\Delta\mathbf{v} < \Delta\mathbf{a} \cdot \Delta\mathbf{r} \tag{2.3}$$

where $\Delta\mathbf{r}$, $\Delta\mathbf{a}$, and $\Delta\mathbf{v}$ are the relative position, acceleration and velocity vectors. When both conditions are met the pair is replaced by a single remnant black hole whose mass is the sum of the progenitor masses (no explicit gravitational-wave mass loss is applied at the simulation level).

To prevent spurious early 'seed clustering' events, mergers occurring within a short time window (10 Myr) of either progenitor's formation are often flagged and excluded in post-processing analyses of the merger rate. No artificial delay or hardening model is imposed once the criteria are satisfied, so coalescence effectively occurs at the resolution limit. Consequently, sub-kiloparsec inspiral and gravitational-wave driven final hardening are unresolved and their associated timescales are not explicitly modelled. The resulting catalogue provides merger times, component masses, and pre-merger relative kinematics.

2.7 Halo and Galaxy Identification

Halos are identified in each snapshot with the AHF (AMIGA Halo Finder) spherical-overdensity algorithm (Knollmann & Knebe 2009). Candidate density peaks are refined adaptively; unbound particles are iteratively removed, yielding virial masses M_{vir} and radii R_{vir} defined where the mean enclosed density $\rho(< R_{vir}) = \Delta_c(z)\rho_{crit}(z)$, where $\rho_{crit}(z) = \frac{3H^2(z)}{8\pi G}$ is the critical density of the Universe, $H(z)$ being the Hubble constant at a given redshift, and with the redshift-dependent overdensity $\Delta_c(z)$ computed following the top-hat collapse formalism (e.g. Bryan & Norman 1998). Subhalos are retained as self-bound structures within host virial radii.

Galaxies are associated with the stellar (and cold gas) components of halos; bulk galaxy properties–stellar mass, star formation rate, gas mass, and mass-weighted centre (used for black hole offset measurements)–are derived from bound particles within the virial radius. Each black hole is linked to its host galaxy and halo via particle IDs; mergers and accretion histories are thus contextualised by evolving host properties. This catalogue provides the structural and kinematic framework for analyses of black hole growth and environmental dependence elsewhere in the thesis.

Chapter 3

Black Hole Spin Model

Here I develop a hybrid post-processing model to follow the spin evolution of MBHs in the Romulus simulations by combining complementary elements from two published frameworks. From Dubois et al. (2014) I adopt the prescription for tracking the spin direction of accreting black holes and the numerical-relativity-based remnant spin (magnitude and direction) estimator applied at black hole mergers. From Ricarte et al. (2023) I incorporate the spin-up prescription that links accretion mode to changes in spin magnitude, distinguishing thin-disk (coherent) and hot (chaotic) fuelling based on the Eddington ratio. The resulting model updates each black hole's spin vector through time using the simulation's recorded accretion history, off-centre state, and merger events. In this chapter I describe the accretion-mode classification and spin-up rules (section 3.1) following Ricarte et al. (2023), the spin reorientation (section 3.2) and merger remnant treatment (section 3.3 following Dubois et al. (2014), and the modifications introduced for this work (off-centre regime, and varying black hole seed-mass) (section 3.4).

3.1 Accretion-Driven Spin-Up Prescription

In this section I outline how black-hole spin magnitudes are updated in response to gas accretion. Following the mode-dependent formulation of Ricarte et al. (2023), each timestep is classified as thin-disk or hot accretion according to an Eddington-ratio threshold; only thin-disk episodes drive significant spin-up, capped at the Thorne limit. The subsections below detail the mode-classification procedure, the spin-up equations, and their implementation in the Romulus post-processing pipeline.

3.1.1 Accretion Mode Classification

In each post-processing snapshot, we decide whether the current accretion episode can spin-up the black hole. The simulation outputs a 10 Myr time-averaged accretion rate $\dot{M}$, so we only need to compare it with the appropriate Eddington limit.

1. **Eddington accretion rate with spin-dependant radiative efficiency**
   The Eddington mass-accretion rate is

$$\dot{M}_{Edd} = \frac{4\pi G M_{BH} m_p}{\eta(a)\sigma_T c(1-\eta(a))} \tag{3.1}$$

   with G the gravitational constant, m_p the proton mass, σ_T the Thomson cross-section, c the speed of light, and η the radiative efficiency. The radiative efficiency η depends on the current spin magnitude because spin changes the location of the ISCO. A rotating black hole drags spacetime so gas on prograde orbits can remain stably bound deeper in the potential well–i.e. its ISCO moves inward from $6GMc^2$ for a Schwarzschild hole to $\sim 1.24GMc^2$ at $a \sim 0.998$. Falling farther releases more gravitational binding energy per unit rest mass, increasing the radiative efficiency. For a prograde thin disk the dimensionless ISCO radius (in units of $\frac{GM}{c^2}$) is

$$r_{ISCO}(a) = 3 + Z_2 - \mathrm{sign}(a)\sqrt{(3-Z_1)(3+Z_1+2Z_2)} \tag{3.2}$$

   where Z_1 and Z_2 are auxiliary variables, found using

$$Z_1 = 1 + (1-a^2)^{\frac{1}{3}}[(1+a)^{\frac{1}{3}} + (1-a)^{\frac{1}{3}}] \tag{3.3}$$

$$Z_2 = \sqrt{3a^2 + Z_1} \tag{3.4}$$

The spin-dependant radiative efficiency is then

$$\eta(a) = 1 - \sqrt{1 - \frac{2}{3r_{ISCO}(a)}} \tag{3.5}$$

This prescription matches the form adopted by Ricarte et al. (2023) and reduces to $\eta \simeq 0.057$ for a Schwarzschild black hole, and $\eta \simeq 0.32$ for a Kerr black hole with $a = 0.9$.

2. **Instantaneous Eddington ratio**
   The Eddington ratio is calculated as

$$f_{Edd} = \frac{\dot{M}}{\dot{M}_{Edd}} \tag{3.6}$$

3. **Thin-disk vs. hot accretion criterion**
   We adopt the critical threshold from Ricarte et al. (2023):

$$f_{crit} = 0.03$$

For $f_{Edd} > 1$ (super-Eddington accretion), the episode is still treated as

Accretion Mode	Criterion	Physical Interpretation
Thin-disk (coherent)	$f_{Edd} \geq f_{crit}$	Radiatively efficient, geometrically thin; contributes coherent angular momentum and spins up the hole.
Hot (chaotic)	$f_{Edd} < f_{crit}$	Radiatively inefficient, geometrically thick; adds mass with negligible net spin-up.

Table 3.1: Summary of the two, eddington ratio-dependent, accretion modes adopted in this model

thin-disk for spin purposes. Although super-Eddignton accretion is not seen in the Romulus simulations, if black hole seed masses are altered in post-processing (discussed in section 3.4) super-Eddington rates can be achieved.

3.1.2 Thin-Disk Spin-Up Calculations

When a timestep is flagged thin-disk, the model follows Ricarte et al. (2023), who write the dimensionless spin-evolution equation

$$s = \frac{da}{dt}\frac{M}{\dot{M}} \qquad \text{with} \qquad s = s_{HD} + s_{EM} \tag{3.7}$$

s_{HD} is the hydrodynamic spin-up term from coherent accretion, and s_{EM} is the electromagnetic spin-down term that accounts for angular-momentum loss through jets.

1. **Hydrodynamic term s_{HD}**
   Using the relativistic thin-disk solution (Bardeen 1970), one obtains

   $$s_{HD} = \frac{l_{ISCO}(a)}{M_{BH}c} - 2ae_{ISCO}(a) \tag{3.8}$$

   where $e_{ISCO}(a)$ and $l_{ISCO}(a)$ are the specific energy and specific angular momentum at the ISCO for spin a. Since e_{ISCO} and l_{ISCO} are not output by the simulation, s_HD is instead modelled, based on a general relativistic radiative magnetohydrodynamic (GRRMHD) simulations of radiative MAD accretion disks around black holes, as

   $$s_{HD} = \frac{s_{thin} + s_{min}\xi}{1+\xi} \qquad \text{with} \qquad \xi = 0.017 f_{Edd} \tag{3.9}$$

   s_{thin} represents the analytic solution for a thin disk, which models approach as $f_{Edd} \to 0$.

   $$s_{thin}(a) = l_{thin}(a) - 2ae_{thin}(a) \tag{3.10}$$

   where $e_{thin}(a)$ and $l_{thin}(a)$ are the specific energy and specific angular momentum at the marginally stable (ms) orbit of a thin-disk (calculated as for the ISCO radius).

   $$e_{thin} = \left(1 - \frac{2}{3r_{ms}}\right)^{\frac{1}{2}} \tag{3.11}$$

   $$l_{thin} = \frac{2}{3\sqrt{3}}\left[1 + 2\left(3r_{ms} - 2\right)^{\frac{1}{2}}\right] \tag{3.12}$$

s_{min} represents the solution for a non-radiative system, which the model approaches as $f_{Edd} \to \infty$.

$$s_{min} = 0.086 - 1.94a \tag{3.13}$$

2. **Electromagnetic term** s_{EM}
   The electromagnetic term s_{EM} represents spin-down by jet launching: large-scale magnetic fields anchored in the accretion flow thread the black hole horizon and, via the Blandford–Znajek process, extract rotational energy and angular momentum. This outward electromagnetic torque saps away the black hole's spin, yielding a negative contribution to s. Ricarte et al. (2023) adopt a phenomenological Blandford-Znajek spin-down rate

$$s_{EM} = -\mathrm{sign}(a)\eta_{EM}\left(\frac{1}{k\Omega_H} - 2a\right) \tag{3.14}$$

where the k parameter is the ratio of the angular frequency of the field lines relative to that of the black hole. Based on other simulations (McKinney et al. 2012; Penna et al. 2013; Chael et al. 2023), it is set as

$$k = \begin{cases} 0.23, & a < 0 \\ (min)(0.1 + 0.5a, 0.35), & a > 0 \end{cases} \tag{3.15}$$

Ω_H is the angular velocity of the horizon

$$\Omega_H = \frac{|a|}{2r_H} \tag{3.16}$$

with the radius of the horizon given by

$$r_H = 1 + \sqrt{1 - a^2} \tag{3.17}$$

and η_{EM} is the electromagnetic jet efficiency

$$\eta_{EM} = \frac{\kappa}{4\pi}\phi^2\Omega_H^2\left[1 + 1.38\Omega_H^2 - 9.2\Omega_H^4\right] \tag{3.18}$$

with the dimensionless magnetisation parameter modelled based on the previously mentioned GRRMHD simulations

$$\phi = \phi_{MAD}(a)\frac{\left(\frac{f_{Edd}}{f_c}\right)^{\alpha}}{1+\left(\frac{f_{Edd}}{f_c}\right)^{\alpha}} \tag{3.19}$$

where f_c is the Eddington ratio at the midpoint of the transition from non-magnetised to highly magnetised and α is a measure of how quickly the magnetisation parameter evolves around f_c. As in Ricarte et al. (2023), they are set to $f_c = 1.88$ and $\alpha = 1.29$ ϕ_{MAD} is the spin dependant magnetisation parameter of systems with a MAD. Here, it is modelled based on GRMHD simulations, run by Narayan et al. (2022), of MADs and their jets.

$$\phi_{MAD}(a) = 52.6 + 34a - 14.9a^2 - 20.2a^3 \tag{3.20}$$

Misalignment of the accretion flow is handled separately, in section 3.2. Here, we assume the thin-disk is aligned at the ISCO.

3.1.3 Hot Accretion Spin-Up Calculations

When the Eddington ratio falls below the thin-disk threshold ($f_{Edd} < f_{crit} = 0.03$), the accretion flow is presumed to be a geometrically thick, radiatively inefficient, hot flow. Ricarte et al. (2023) calibrate its impact on spin using GRMHD simulations of MADs from Narayan et al. (2022). In this regime the specific angular momentum delivered to the black hole is greatly reduced.

Based on the simulations from Narayan et al. (2022), the spin-up parameter is modelled as

$$s_{MAD}(a) = 0.45 - 12.53a - 7.80a^2 + 9.44a^3 + 5.71a^4 - 4.03a^5 \tag{3.21}$$

3.1.4 Time Integration

The spin-evolution equations give an ordinary differential system for black-hole mass and spin magnitude

$$\frac{dM}{dt} = \dot{M}, \quad \frac{da}{dt} = s(a, f_{Edd})\frac{\dot{M}}{M} \tag{3.22}$$

1. **Global and sub-timesteps**
   We integrate from snapshot t_i to $t_{i+1} = t_i + \Delta T$ with

   $$\Delta T = 0.1\text{Gyr}.$$

   For improved accuracy, we divide this main step into

   $$N_{sub} = 1000 \implies \delta t = \frac{\Delta T}{N_{sub}} = 10^5\text{yr}.$$

   Within a main step, $\dot{M}$ is held constant as the simulation only outputs it every ΔT, while M and a evolve every sub step.

2. **Fourth-order Runge-Kutta (RK4) scheme**
   For every sub step $n \to n+1$ we evolve the state vector $\mathbf{y} = (M, a)$:

   $$\mathbf{y}_{n+1} = \mathbf{y}_n + \frac{\delta t}{6}(k_1 + 2k_2 + 2k_3 + k_4) \tag{3.23}$$

   where

   $$\begin{aligned} k_1 = \mathbf{F}(\mathbf{y}_n), \quad k_2 = \mathbf{F}(\mathbf{y}_n + \tfrac{1}{2}k_1), \\ k_3 = \mathbf{F}(\mathbf{y}_n + \tfrac{1}{2}k_2), \quad k_4 = \mathbf{F}(\mathbf{y}_n + k_3) \end{aligned} \tag{3.24}$$

   and $\mathbf{F} = (\dot{M}, s\frac{\dot{M}}{M})$ from equation 3.22.

3. **Mass interpolation inside RK4**
   Because $\dot{M}$ is fixed over ΔT, the mass inflow is linear in time over the same:

   $$M(t) = M_0 + \dot{M}(t - t_0) \tag{3.25}$$

By interpolating the mass, the mass-dependent Eddington ratio can be calculated for each intermediate step within the RK4. This allows the black hole to possibly move from one accretion regime to another, even within a sub step.

4. **Spin cap at the Thorne limit**
   In a perfectly coherent, radiatively efficient thin-disk, continued prograde accretion would naively drive $a \to 1$. In practice, capture of retrograde photons (and other radiative angular-momentum losses) from the inner disk exerts a counter-torque that sets an equilibrium (the "Thorne limit") at $a_T = 0.998$ (Bardeen 1970; Thorne 1974). Above this, additional accreted rest mass increases M faster than J, so a cannot grow further. Our electromagnetic (jet) spin-down term can also push the equilibrium to slightly lower values, but we still impose a hard cap at the Thorne limit to avoid unphysical super-Thorne spins.

   After each main step, if $a \leq a_T$, it is accepted, but if $a > a_T$, it is set to a_T. The excess angular momentum is interpreted to be removed by photon capture/electromagnetic torques and is not carried forward.

At the end of the 0.1 Gyr interval the updated (M,a) serve as initial conditions for the next main step.

3.2 Spin Orientation Tracking and Re-Alignment Prescription

We evolve the spin direction (unit vector $\hat{\mathbf{s}} \equiv \frac{\mathbf{J}_{BH}}{|\mathbf{J}_{BH}|}$) following the prescription of Dubois et al. (2014), which couples the black hole spin direction to the angular momentum of nearby gas. This section describes how the reference gas angular momentum is measured for central black holes, the alignment/counter-alignment criterion, and how $\hat{s}$ is rotated toward (or away from) the gas.

3.2.1 Sampling the Reference Gas Angular Momentum

The simulation outputs the angular momentum radial profile and the angular momentum phi and theta radial profiles for each galaxy at each main timestep. In post-processing we did not have access to raw neighbour lists around each black hole at arbitrary sub-timesteps, so we could not recompute the local gas angular momentum as in Dubois et al. (2014). Instead, we use the angular momentum of the central 1.4 kpc (four gravitational softening lengths; the same aperture used in the simulation for central gas density) of the host galaxy. The galaxy's angular momentum is vectorially integrated within the 1.4 kpc aperture, and is normalised to obtain the direction of the angular momentum.

$$\hat{\mathbf{l}} \equiv \frac{\mathbf{J}_{cen}}{|\mathbf{J}_{cen}|}$$

The disk around the black hole is assumed to be aligned with this gas angular momentum.

3.2.2 Misalignment and the Alignment/Counter-Alignment Criterion

The misalignment angle is defined as

$$\cos\theta = \hat{\mathbf{s}} \cdot \hat{\mathbf{l}} \tag{3.26}$$

Following King et al. (2005)/Dubois et al. (2014), if the central disk angular momentum J_d are partially or entirely anti-aligned (i.e. $\cos\theta < 0$) and J_d is insufficient to flip the black hole from its degree of misalignment (i.e. $\cos\theta < -\frac{J_d}{2J_{BH}}$), accretion proceeds in a counter-aligned configuration and the black hole's angular momentum anti-aligns with the disk. Otherwise, the black hole always aligns with the disk.

In order to assess the criterion, the magnitude of the disk angular momentum relative to that of the black hole is required. Since we cannot resolve the accretion disk, we assume, as in Dubois et al. (2014), the tilted thin disk solution (Shakura

& Sunyaev 1973). This allows the relative magnitudes to be found as

$$\frac{J_d}{2J_{BH}} \sim 6.8 \times 10^{-2} \left(\frac{f_{Edd}}{\varepsilon_{r,01}}\right)^{\frac{1}{8}} M_{BH,8}^{\frac{23}{16}} a^{\frac{3}{16}} \alpha_{t,01}^{-\frac{7}{4}} \left(\frac{\nu_2/\nu_1}{85}\right)^{-\frac{19}{16}} \tag{3.27}$$

where $\varepsilon_{r,01}$ is the radiative efficiency normalized to 0.1, ν_1 and ν_2 are the viscosities responsible for driving accretion/transferring angular momentum and for disk warp propagation, respectively, and $\alpha_{t,01}$ is a viscosity parameter normalised to 0.1. The radiative efficiency is calculated as normal, based on the current black hole spin. The normalised viscosity parameter is set to 1, as a typical value from observations (King et al. 2007). This allows the ratio of the viscosities to be calculated

$$\frac{\nu_2}{\nu_1} = \frac{2\alpha_t^2\,(1+7\alpha_t)}{4+\alpha_t^2} \sim 85 \tag{3.28}$$

After each main timestep, the angular momentum ratio between the disk and the black hole is calculated, the misalignment of the disc's current angular momentum and the black hole's previous spin direction is found, and the alignment criterion is evaluated. If $\cos\theta < -\frac{J_d}{2J_{BH}}$, the new black hole spin direction is set to $-\hat{\mathbf{l}}$, otherwise it is set to $\hat{\mathbf{l}}$.

3.3 Merger Handling and Remnant Spin Prescription

We compute remnant (post–coalescence) spin magnitude and direction using the vector fitting formalism implemented in Dubois et al. (2014), which is based on the numerical–relativity fits of Rezzolla et al. (2008) (and subsequent refinements).

3.3.1 Merger Event Preparation

When a black hole binary merger event is identified in the simulation, the merger time is obtained form the simulation merger database. At the merger time, we use

- component black hole masses: m_1 and m_2, chosen so $m_1 \geq m_2$, integrated until the most recent sub timestep, as described in section 3.1.4

- component spin magnitudes: a_1 and a_2, integrated until the most recent sub timestep, as described in section 3.1.4
- component spin unit vectors: $\hat{\mathbf{s}}_1$ and $\hat{\mathbf{s}}_2$, integrated until the most recent main timestep, as described in section 3.2

From the simulation outputs, the position ($\mathbf{r}_1$ and $\mathbf{r}_2$) and velocity ($\mathbf{v}_1$ and $\mathbf{v}_2$) vectors of the component black holes are taken at the most recent main timestep for which they are all defined.

3.3.2 Calculating Orbital Angular Momentum

The orbital angular momentum of the black hole binary is a key factor in determining the remnant spin. Because the simulation does not model the final inspiral of black hole binaries, the orbital angular momentum must be estimated. Rezzolla et al. (2008) define a normalised orbital angular momentum parameter

$$\boldsymbol{\ell} = \frac{\mathbf{l}}{m_1 m_2} \tag{3.29}$$

where $\mathbf{l}$ is the orbital angular momentum of the binary just prior to coalescence, i.e. the difference between the orbital angular momentum of the binary when widely separated ($\mathbf{L}$) and the angular momentum extracted via gravitational waves during the inspiral. Rezzolla et al. (2008) find an analytical expression for the magnitude ℓ.

$$\begin{aligned} \ell = & \frac{s_4}{(1+q^2)^2}\left(a_1^2 + a_2^2 q^4 + 2\mathbf{a_1}\cdot\mathbf{a_2} q^2\right) \\ & + \frac{s_5\mu + t_0 + 2}{1+q^2}\left(a_1\cos\phi_1 + a_2 q^2\cos\phi_2\right) \\ & + 2\sqrt{3} + t_2\mu + t_3\mu^2 \end{aligned} \tag{3.30}$$

where ϕ_i is the angle between $\mathbf{a}_i$ and $\boldsymbol{\ell}$ for $i = 1,2$, q is the mass ratio of the component black holes: $q = \frac{m_2}{m_1}$, and $\mu = \frac{q}{(1+q)^2}$. The numerical constants are $s_4 = -0.129$, $s_5 = -0.384$, $t_0 = -2.686$, $t_2 = -3.454$, and $t_3 = 2.353$.

To determine the direction of $\boldsymbol{\ell}$, we assume it is aligned with $\mathbf{L} = \mathbf{L}_1 + \mathbf{L}_2$, i.e. that the gravitational waves extract angular momentum isotropically. To calculate

$\mathbf{L}_i$, for $i = 1,2$, we first find the centre of mass position and velocity.

$$\mathbf{r}_{com} = \frac{m_1\mathbf{r}_1 + m_2\mathbf{r}_2}{m_1 + m_2} \tag{3.31}$$

$$\mathbf{v}_{com} = \frac{m_1\mathbf{v}_1 + m_2\mathbf{v}_2}{m_1 + m_2} \tag{3.32}$$

and use these to calculate $\mathbf{L}_i$

$$\mathbf{L}_i = m_i\left(\mathbf{r}_i - \mathbf{r}_{com}\right) \times \left(\mathbf{v}_i - \mathbf{v}_{com}\right) \tag{3.33}$$

After summing $\mathbf{L}_1$ and $\mathbf{L}_2$, $\mathbf{L}$ is finally normalised, to find the direction of $\boldsymbol{\ell}$

$$\hat{\boldsymbol{\ell}} \equiv \frac{\mathbf{L}}{|\mathbf{L}|} \tag{3.34}$$

3.3.3 Remnant Spin Calculation

Finally, the remnant spin vector is calculated using the analytical fit of Rezzolla et al. (2008) based on relativistic numerical simulations of black hole binaries.

$$\mathbf{a} = \frac{1}{1+q^2}\left(\mathbf{a}_1 + \mathbf{a}_2 q^2 + \mathbf{l}q\right) \tag{3.35}$$

The magnitude and direction of the spins used to calculate the remnant spin, as well as the remnant spin magnitude and direction are recorded in a merger database. The remnant spin magnitude is used as the next initial spin for spin up via accretion, and is integrated as in section 3.1.4 until the next main timestep. The remnant spin direction is used as the previous spin direction when evolving the direction at the next main time step as in section 3.2.

3.4 Model Variations

In the previous sections we discussed how a black hole's spin direction follows the angular momentum of gas in the central region of its host galaxy when the black hole is also central. Here central is defined as the inner sphere of radius $r_{cen} = 4\varepsilon_g = 1.4$kpc, chosen to be large enough to average over particle noise yet small enough to trace the nucleus where coherent inflow originates. However, in Romulus many MBHs spend substantial time off-centre. In those phases the local accretion geometry is poorly resolved and the central gas no longer provides an obvious proxy for the spin-alignment axis. To bracket this uncertainty we introduce a set of alternative prescriptions for how off-centre black holes accrete and align, and we also explore how the adopted seed mass sets the early Eddington history and thus spin build-up.

In addition to the baseline spin-evolution machinery, we explore model variations that represent key unresolved physics in how MBHs grow while they are not central to their host galaxy, and how their initial seed masses set early accretion and spin trajectories. Specifically, we implement three alternative off-centre accretion prescriptions–an assumed central model, a chaotic accretion model and a gas adhesion model. Each prescription is evaluated for two seed masses ($M_{seed} = 10^5 M_\odot$ and $M_{seed} = 10^6 M_\odot$) to test sensitivity to the observationally unconstrained initial conditions and the resulting differences in early Eddington ratios and spin build-up. The following subsections detail the implementation of the three off-centre prescriptions, the rationale and procedure for the seed mass experiment, and finally a consolidated summary mapping these choices to the six labelled models used in the comparison and results.

3.4.1 Assumed Central Off-Centre Accretion Prescription

The assumed central prescription is a deliberately simple baseline used to validate the rest of the spin–evolution pipeline (accretion mode classification, spin magnitude updates, orientation tracking, merger handling) without introducing additional complexity from modelling the detailed dynamics of wandering black holes. It serves as a control case against which the more physically motivated prescriptions (chaotic accretion and gas adhesion) can be compared.

Under this prescription, even when a black hole is spatially offset from the potential minimum of its host, it is treated as if it were still located at the centre for the purpose of selecting the gas reservoir from which it accretes. This defines the angular momentum vector that sets spin alignment. Thus, displacement does not alter either the magnitude or the coherence of the accreted angular momentum in this model.

When a black hole is identified as > 1.4kpc from the galactic centre (i.e. not within the central gas reservoir used in section 3.2), the procedure is as follows:

1. **Centre determination**: For every main timestep we identify the host galaxy centre (as in section 3.2) and use the central 1.4 kpc angular momentum vector $\hat{\mathbf{l}}_{cen}$ regardless of the black hole's actual position.

2. **Accretion Mode**: The Eddington ratio is computed normally from the simulation output $\dot{M}$; no suppression or modification is applied during off-centre intervals.

3. **Spin magnitude updates**: Accretion episodes apply the full spin-up prescriptions detailed in 3.1.

4. **Spin direction updates**: The direction of black hole spin is found as described in section 3.2, using the previously found $\hat{\mathbf{l}}_{cen}$ as the direction of the surrounding gas angular momentum.

This prescription was easy to implement as no additional state needs to be tracked. The existing central angular momentum calculation suffices; wandering status is effectively ignored in the spin and accretion update routines. However, it neglects angular momentum decorrelation. Spatial offsets that could misalign the small-scale disk from the global galactic plane are not captured.

3.4.2 Chaotic Off-Centre Accretion Prescription

The chaotic off-centre prescription is designed to bracket the minimum coherent spin growth plausibly associated with wandering phases. It tests how strongly the spin evolution and alignment statistics depend on assuming that displaced black holes accrete gas with randomly oriented angular momentum parcels that cancel net spin-up, while still allowing mass build-up. Comparing this to the assumed central and gas adhesion regimes isolates the impact of suppressing coherent angular-momentum delivery during off-centre episodes.

Under this prescription, when the black hole is off-centre, the inflowing material arrives in short, randomly oriented packets. Mass does accumulate at the simulation-provided rate $\dot{M}$, but the vector sum of specific angular momentum additions is taken to be zero. Consequently, the hydrodynamic spin-up term vanishes but spin can still decrease through electromagnetic jet torques, in high eddington ratio episodes, or through MAD spin-down, in low eddington ratio episodes.

When a black hole is identified as > 1.4kpc from the galactic centre, the procedure is as follows:

1. **Accretion Mode**: Mode classification (thin-disk vs hot) still uses the spin-dependent radiative efficiency and the supplied mass accretion rate; wandering does not suppress the mass inflow rate (so the mass history is identical across prescriptions for identical $\dot{M}$.

2. **Spin magnitude updates**: In the thin-disk regime, the evaluation of s_{HD} is skipped, only s_{EM} is computed. In the hot accretion regime, s_{MAD} is computed as normal. The spin is then updated as described in section 3.1.4 using the spin-up parameter calculated here.

3. **Spin direction updates**: Alignment torque toward the central angular momentum vector is disabled while off-centre, i.e. the direction remains unchanged from the previous step

While more physically realistic than the assumed central prescription, this prescription does have some physical limitations. Real off-centre accretion may retain some partial coherence; setting $s_{HD} = 0$ is an extreme (lower-limit) scenario. True chaotic feeding could also induce small random tilts; diffusion could approximate this but is not included in the these runs. Finally, any gas captured while wandering is not retained for future coherent spin-up, however a prescription including this is discussed in section 3.4.3.

3.4.3 Gas Adhesion Off-Centre Accretion Prescription

The gas adhesion prescription explores an intermediate, more physically motivated scenario between the optimistic assumed central model and the pessimistic chaotic model. It assumes that when a black hole leaves the galactic centre it retains (or remains coupled to) a coherent mini-disc/bound gas reservoir whose angular momentum direction reflects that of the last time it was central. While displaced, the black hole continues to accrete (at the simulation-provided rate) from this "adhesive" reservoir, so coherent spin-up can continue even off-centre. This probes whether preserving a memory of the central disk can sustain higher spins during wandering phases.

Upon crossing outward past the displacement threshold (becoming off-centre), the black hole captures a gas reservoir whose angular momentum unit vector is the central 1.4 kpc vector $\hat{\mathbf{l}}_{cen}$ at that moment. That vector is then frozen as the adhesion vector $\hat{\mathbf{l}}_{adh}$ for the entire off-centre interval. All alignment operations while off-centre reference $\hat{\mathbf{l}}_{adh}$ rather than the instantaneous central $\hat{\mathbf{l}}_{cen}$. When the black hole re-enters the centre of the galaxy, $\hat{\mathbf{l}}_{adh}$ is refreshed to the current $\hat{\mathbf{l}}_{cen}$.

When a black hole is identified as > 1.4kpc from the galactic centre, the procedure is as follows:

1. **Gas reservoir determination**: The most recent timestep in which the black hole was central is identified and the central 1.4 kpc angular momentum vector of the host galaxy at that time is calculated, as described in section 3.2, to find $\hat{\mathbf{l}}_{adh}$.

2. **Accretion Mode**: Mode classification (thin-disk vs hot) still uses the spin-dependent radiative efficiency and the supplied mass accretion rate; wandering does not suppress the mass inflow rate (so the mass history is identical across prescriptions for identical $\dot{M}$.

3. **Spin magnitude updates**: Accretion episodes apply the full spin-up prescriptions detailed in 3.1.

4. **Spin direction updates**: The direction of black hole spin is found as described in section 3.2, using the previously found $\hat{\mathbf{l}}_{adh}$ as the direction of the surrounding gas angular momentum.

This prescription is more physically motivated than the assumed central prescription as the off-centre status is tracked and spin orientation is not continuously updated from the evolving central gas reservoir while the black hole is off-centre. It is also less extreme than the chaotic prescription as coherent spin-up can continue away from the centre of the galaxy if the accretion rates are high enough. However, a real bound gas reservoir would likely be drained over time through accretion or be warped by surroundings, but in this model the reservoirs supply infinite coherent angular momentum for the duration of the displacement.

3.4.4 Seed Mass Experiment

The Romulus simulation adopts a default black hole seed mass of $10^6 M_\odot$. To test the sensitivity of early spin evolution to the initial mass scale, we re-run each off-centre accretion prescription with an artificially lowered seed mass of $10^5 M_\odot$. Holding all other inputs fixed isolates how a smaller initial mass (at fixed accretion history) modifies the Eddington ratio and hence the time spent in the thin-disk (coherent) regime that drives spin-up.

For every black hole and timestep we use the same simulation-output 10 Myr averaged mass accretion rate $\dot{M}$ in both seed experiments. Only the starting mass differs. With:

$$f_{Edd} = \frac{\dot{M}}{\dot{M}_{Edd}(M,a)} \propto \frac{\eta(a)}{M} \tag{3.36}$$

the lower-mass seeds have systematically higher f_{Edd} at early times. Because the spin-dependent radiative efficiency $\eta(a)$ is applied self-consistently, any accelerated spin-up in the low-seed case (raising η) slightly reduces subsequent $\dot{M}_{Edd}$, reinforcing high f_{Edd} and prolonging the thin-disk phase–a positive feedback on early coherent growth.

The seed mass experiment is implemented by simply changing the initial mass of the black holes and integrating the mass as normal (described in section 3.1.4). All other procedures remain the same.

Lower seed masses act as a control knob that rescales the Eddington ratio for an unchanged $\dot{M}$. This mimics a scenario where the black holes were seeded in an environment which is more gas-rich than the simulation's default, allowing us to probe whether early sustained near-Eddington, coherent accretion can leave a long-lived spin imprint. As black holes grow, the fractional difference between the two seed choices becomes negligible (once $M \gg 10^6 M_\odot$), so the seed experiment primarily constrains early-epoch spin divergence and its persistence to $z \sim 0$.

While changing the seed mass in post-processing offers interesting insights, it is not without limitations. Most notably, using identical $\dot{M}$ histories for black holes with different total masses (due to different seed masses) ignores possible physical coupling between seed mass and gas inflow. This creates physically unrealistic black hole accretion rates.

3.4.5 Model Variations Summary

Table 3.2 summarises the six model variants that we ran, formed by the combination of the three off-centre accretion prescriptions with the two seed masses, all described in the preceding subsections.

Model Name	**Off-Centre Model**	**Black Hole Seed Mass**
10^5&Central	Assumed Central	$10^5 M_\odot$
10^6&Central	Assumed Central	$10^6 M_\odot$
10^5&Chaotic	Chaotic Accretion	$10^5 M_\odot$
10^6&Chaotic	Chaotic Accretion	$10^6 M_\odot$
10^5&Adhesion	Gas Adhesion	$10^5 M_\odot$
10^6&Adhesion	Gas Adhesion	$10^6 M_\odot$

Table 3.2: *Summary of the six black hole spin evolution models explored in this study.*

Chapter 4

Results

4.1 Model Comparison - Off-Centre Accretion Prescriptions and Seed Masses

To examine the effects of different off-centre regimes and seed masses on final black hole spin, I first compare the spin distributions at redshift zero for each model. The results are binned by black hole mass to account for expected mass-dependent spin evolution due to differing accretion histories and merger rates.

Black holes are grouped into three mass bins depending on their seed mass: $10^5 < \frac{M_{BH}}{M_\odot} < 10^6$, $10^6 < \frac{M_{BH}}{M_\odot} < 10^7$, $\frac{M_{BH}}{M_\odot} > 10^7$ for models using $M_{seed} = 10^5 M_\odot$; and $10^6 < \frac{M_{BH}}{M_\odot} < 10^7$, $10^7 < \frac{M_{BH}}{M_\odot} < 10^8$, $\frac{M_{BH}}{M_\odot} > 10^8$ for models using $M_{seed} = 10^6 M_\odot$.

These bins were chosen to (i) separate seed-scale black holes from those that have undergone significant growth, and (ii) isolate the high-mass end where merger-driven spin evolution may dominate.

Figure 4.1 shows the normalised spin distributions at z=0 for each model, with panels grouped by mass bin. Each curve represents a different off-centring regime and seed mass combination. Across all models, spin values range from 0 to 1, with distinct patterns emerging as a function of both seed mass and off-centre treatment.

Across all six models, the redshift $z = 0$ spin distributions are heavily skewed toward low spin magnitudes, with the majority of black holes exhibiting spins

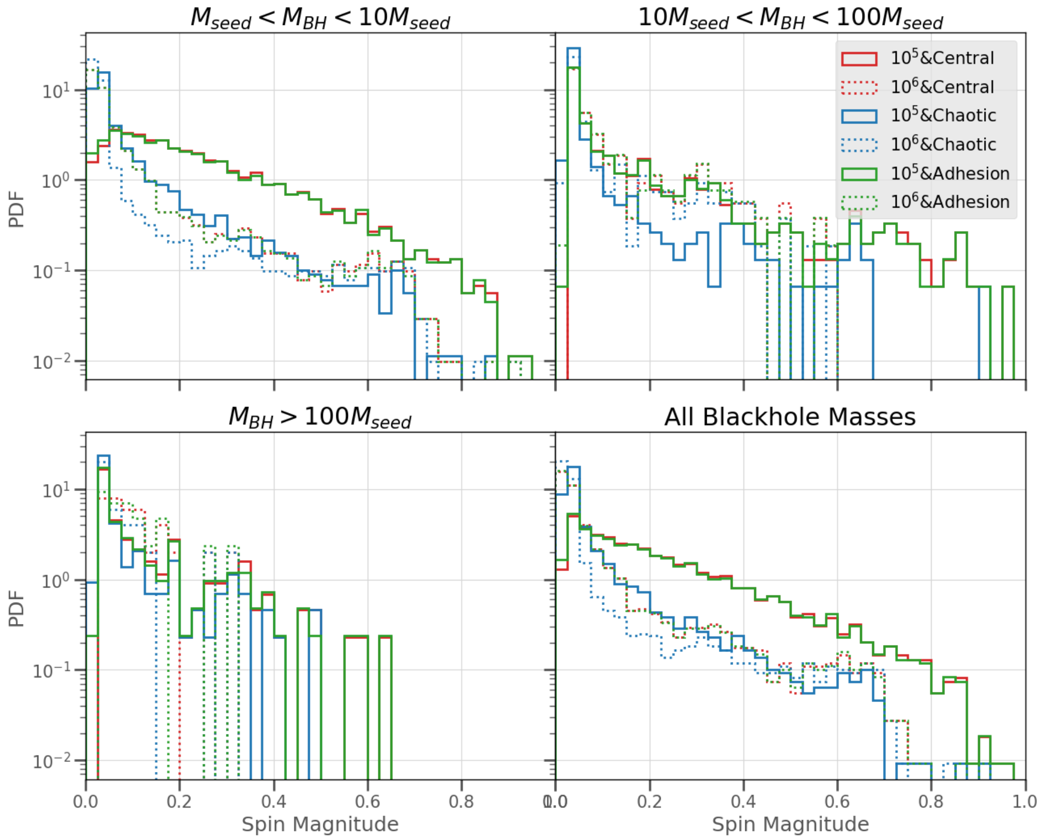

Figure 4.1: *Spin magnitude distributions of MBHs at redshift z=0 for all six model variations. The four panels plot different mass bins: top left: $M_{seed} < M_{BH} < 10M_{seed}$, top right: $10M_{seed} < M_{BH} < 100M_{seed}$, bottom left: $M_{BH} > 100M_{seed}$, bottom right: all masses. Models are differentiated by using different colours for each off-centre accretion prescription (Assumed central in red, Chaotic in blue, and Gas Adhesion in green) and using a different linestyle for each seed mass (solid for $10^5 M_\odot$ and dotted for $10^6 M_\odot$. All six models are dominated by low-spin values. Differences in spin distribution shape between regimes and seed masses highlight the sensitivity of spin evolution to seeding and dynamical treatment. Notably, the Chaotic Accretion prescription exhibits a more sharply peaked distribution at very low spins ($a < 0.05$), especially in the two higher mass bins.*

$a < 0.1$. This behaviour is most pronounced in the intermediate and high mass bins, which require higher accretion rates compared their lower mass counterparts to achieve the same eddington ratio necessary for coherent growth which efficiently spins up black holes. Clear differences emerge between the three off-centring regimes, reflecting the impact of how gas accretion is treated when black holes

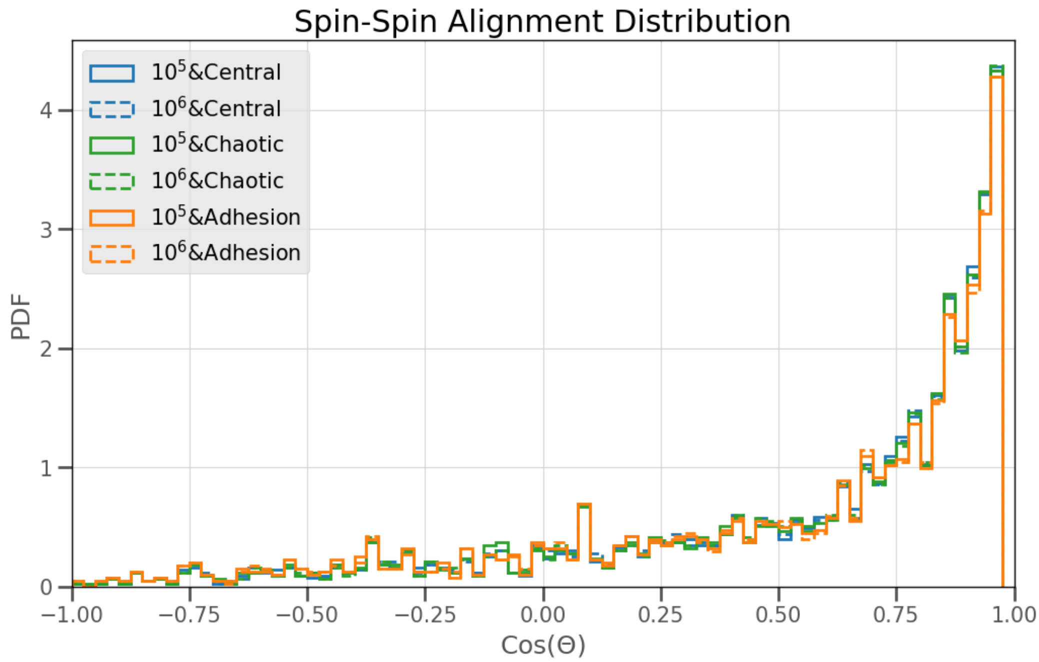

Figure 4.2: *Distribution of cosine of the angle between spin vectors of merging black hole pairs for all six spin evolution models, at the moment just prior to merger. Models using a black hole seed mass of* $10^5 M_\odot$ *are shown using solid lines, and models using a seed mass of* $10^6 M_\odot$ *are shown using dashed lines. Colour indicates off-centre regime: assumed central (blue), chaotic accretion (green), and gas adhesion (orange). All models show a strong peak near* $cos(\theta) = 1$*, indicating preferential alignment of spins, with little variation across regimes or seed mass assumptions.*

are displaced from the centres of their host galaxies.

The Chaotic Accretion prescription yields the narrowest spin distributions, sharply peaked at near-zero values. In contrast, the Assumed Central prescription and the Gas Adhesion prescription allow for more sustained spin-up. This results in broader spin distributions and a small tail of moderate-spin black holes with $a \approx 0.3 - 0.5$, especially at lower masses where accretion is more influential.

Seed mass also modulates the outcome. Models seeded with $10^5 M_\odot$ black holes (top row) tend to show slightly broader spin distributions than their $10^6 M_\odot$ counterparts (bottom row), particularly in the lowest mass bin. For a fixed accreted mass ΔM, a smaller black hole experiences a much larger fractional growth ($\Delta \frac{M}{M_{BH}}$), so the angular momentum of the newly accreted gas has a correspondingly bigger

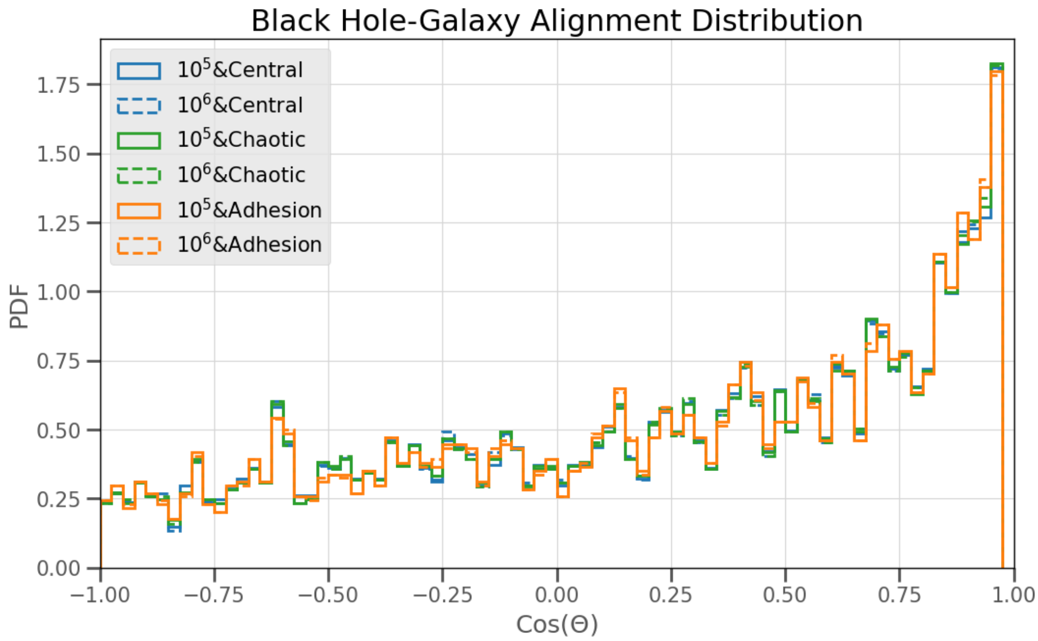

Figure 4.3: *Distribution of the cosine of the angle between each merging black hole's spin vector and the angular momentum of the centre of its host galaxy, measured just prior to merger for all models. Models using a black hole seed mass of $10^5 M_\odot$ are shown using solid lines, and models using a seed mass of $10^6 M_\odot$ are shown using dashed lines. Colour indicates off-centre regime: assumed central (blue), chaotic accretion (green), and gas adhesion (orange). All models show a clear tendency toward alignment ($cos(\theta) = 1$), with minimal variation between models.*

leverage on the spin. Once $M_{BH} \gg \Delta M$, additional episodes contribute less to the existing angular momentum budget and a evolves slowly. This highlights that the earliest growth phase (before $\sim 10^6 M_\odot$) is crucial for setting the eventual spin distribution so ignoring this initial build-up risks missing key spin evolution.

Together, these results highlight how both seeding strategy and assumptions about accretion physics significantly shape the final spin distribution of MBHs.

Figures 4.2 and 4.3 show the distributions of spin alignments for merging black holes across all six models, comparing both spin-spin alignment (between the two black holes) and black hole-galaxy alignment (between each black hole' spin and the angular momentum of its host galaxy's centre). In both cases, the distributions are highly consistent across all models, regardless of off-centre regime or seed mass.

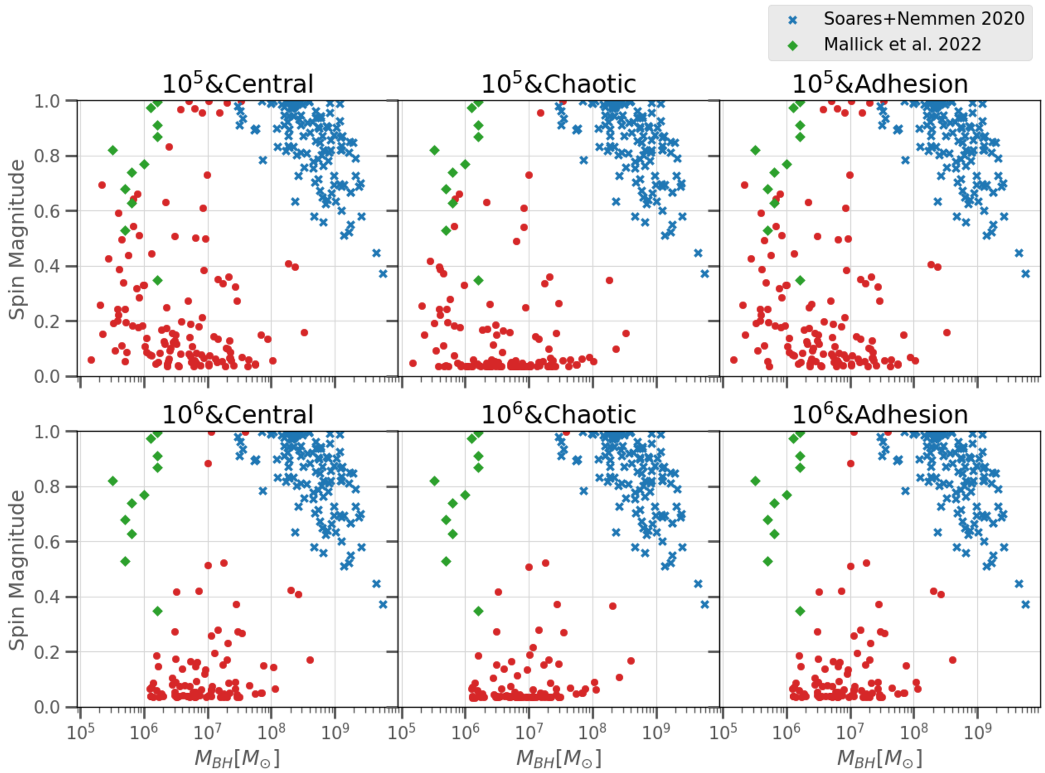

Figure 4.4: *Spin magnitude versus black hole mass (M_{BH}) at $z = 1$ for all six models, showing only black holes with bolometric luminosities exceeding* $10^{42}\mathrm{erg\,s^{-1}}$, *representing an approximate observational detection threshold. Simulated black holes are shown in red, while observational data from Soares & Nemmen (2020) (blue crosses) and Mallick et al. (2022) (green diamonds) are over plotted for comparison. The differences across models reflect the influence of varying seed masses and off-centre growth prescriptions on the spin–mass distribution of observable black holes.*

The spin-spin alignment distribution shows a strong peak near $cos(\theta) = 1$, indicating that black holes typically merge with closely aligned spins. This likely reflects the fact that merging black holes are embedded in a shared gaseous environment in the galactic centre, where both rapidly align their spins with the same angular momentum direction - through both direct accretion and torques from surrounding gas.

Similarly, the black hole-galaxy alignment distribution shows a preference for spin vectors to align with the angular momentum of the galaxy centre. While the alignment is slightly weaker than the spin-spin case, it remains consistent across models. This implies that black holes maintain strong coupling to their local galactic environment up to the point of merger, and that even in cases of more significant wandering, alignment is generally re-established before coalescence.

Figure 4.4 shows the spin-mass distribution, at redshift $z = 1$, of black holes with bolometric luminosities exceeding $10^{42}\mathrm{erg\,s^{-1}}$, overlaid with observational spin measurements from Soares & Nemmen (2020) and Mallick et al. (2022). These observations cover redshifts up to $z \sim 2.6$. Across all six models, the simulated population of active black holes occupies a distinct region of the spin-mass parameter space compared to observational data. Most notably, the simulations tend to produce predominantly low-spin ($a < 0.3$) active black holes across the full mass range. This is in contrast to the observed systems–particularly those compiled by Soares & Nemmen (2020)–which exhibit systematically higher spins, often exceeding $a > 0.6$, and reach a higher mass.

A notable trend is that the lower seed mass models (top row) produce more high-spin black holes than their higher seed mass counterparts (bottom row). This suggests that lower initial black hole masses may allow for a greater degree of spin-up through accreting at higher eddington ratios.

Among the models, both the Assumed Central prescription and the Gas Adhesion prescription do produce a small number of high-spin black holes ($a > 0.6$), although these remain relatively rare. Even in these cases, the high-spin population is significantly underrepresented compared to observations. The Chaotic Accretion prescription, as expected, shows the lowest spin values overall, consistent with repeated accretion of misaligned material leading to spin-down.

The general failure of all models to reproduce the observed population of high-spin, high-mass black holes may point to missing physical processes in the simulation framework–such as sustained coherent accretion, or late mergers–or to observational selection effects that preferentially detect high-spin systems via X-ray reflection or continuum fitting techniques (Reynolds 2019).

In the following sections, we present results primarily from the 10^5&Adhesion model, which adopts the gas adhesion regime and uses a black hole seed mass of $10^5 M_\odot$. This model was selected for detailed analysis based on its physical

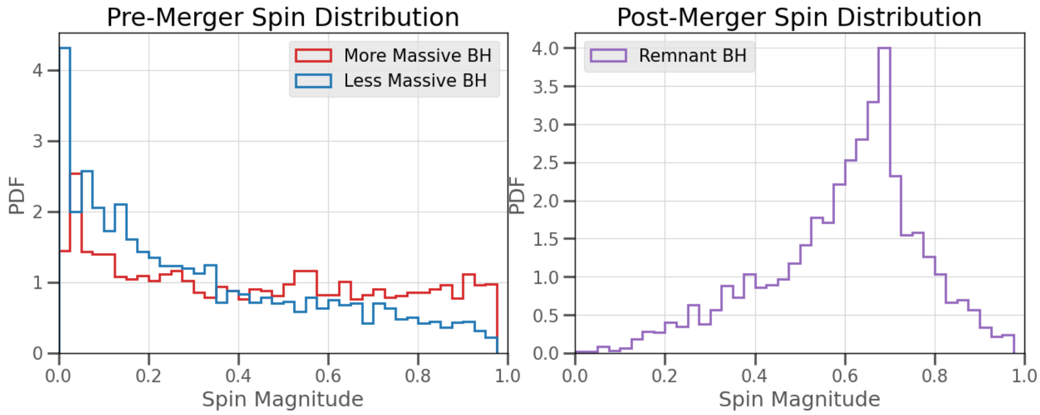

Figure 4.5: *Spin magnitude distributions of merging black holes in the* 10^5*&Adhesion model, shown before and after merger. Left: Spin magnitudes of the more massive (red) and less massive (blue) progenitor black holes immediately prior to merging. Both populations are dominated by low spins, particularly the less massive black holes. Right: Spin distribution of the remnant black holes immediately after the merger. A clear peak emerges at moderate spin values (*$a \sim 0.7$*). These results highlight the role of mergers in building up the spin of MBHs in this model.*

plausibility and ability to replicate current measurements of spin. While all six models explored in this study offer valuable insights, we focus on the 10^5&Adhesion model to provide a clearer and more concise interpretation of black hole spin and feedback evolution.

4.2 Black Hole Mergers

Figure 4.5 shows the distribution of spin magnitudes for merging black holes in the 10^5&Adhesion model, comparing their spin immediately before the merger (left panel) to that of the merger remnant (right panel). The left panel distinguishes between the more massive and less massive black holes involved in each merger, while the right panel displays the resulting post-merger spin distribution.

The left panel shows that prior to merger, the spin magnitudes of both more massive black holes is relatively flat across the entire spin range whereas less massive black holes in the 10^5&Adhesion model show a peak near zero. The less massive black holes showing a stronger preference for low spins likely reflects

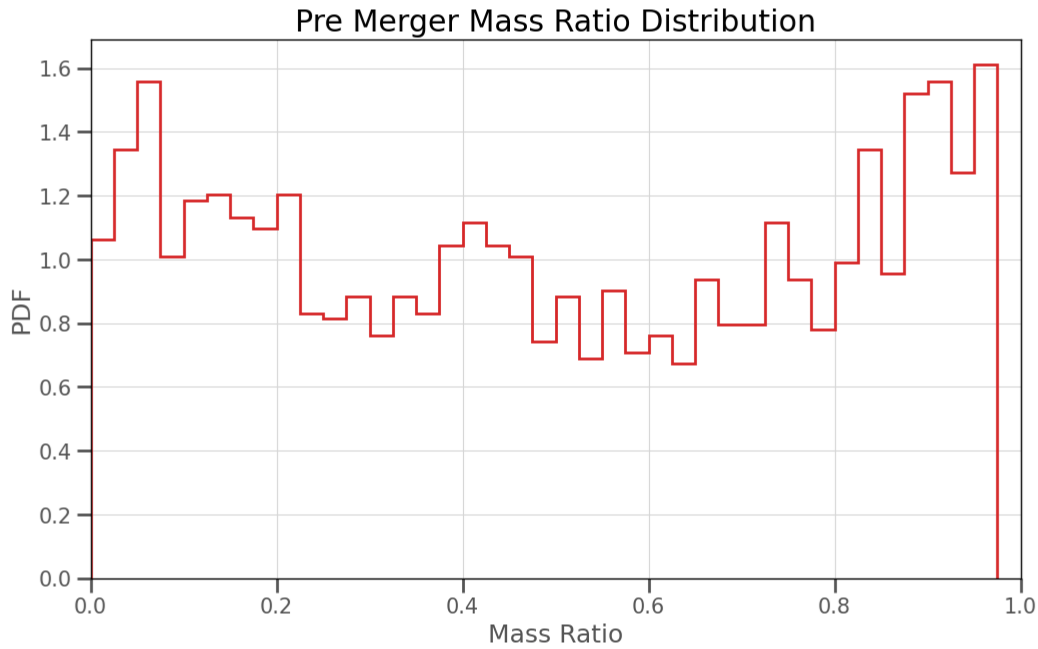

Figure 4.6: *Distribution of black hole merger mass ratios in the* 10^5*&Adhesion model. The mass ratio is defined as* $q = \frac{M_{less}}{M_{more}}$*, where* M_{less} *and* M_{more} *are the masses of the less and more massive black holes, respectively, just prior to merging. A value of* $q = 1$ *corresponds to an equal-mass merger, while lower values indicate more unequal mass ratios.*

shorter accretion histories and fewer coherent episodes of angular momentum buildup.

Following the mergers, however, the right panel reveals that the remnant black holes typically exhibit moderate spin values, peaking around $a \sim 0.7$. This aligns with expectations from numerical relativity, where non-spinning or low-spin black holes merging with mass ratios $\gtrsim 0.3$ tend to produce remnants with spins in the $a \sim 0.6-0.8$ range (Rezzolla et al. 2008; Buonanno et al. 2008). The emergence of this peak in the post-merger distribution demonstrates how mergers act as a significant channel for spin-up, even when the progenitor black holes are slowly rotating.

This transition highlights the dual role of accretion and mergers in shaping the spin evolution of MBHs: while gas accretion alone may not lead to high spins in chaotic environments, mergers can reset and boost spins, contributing to the

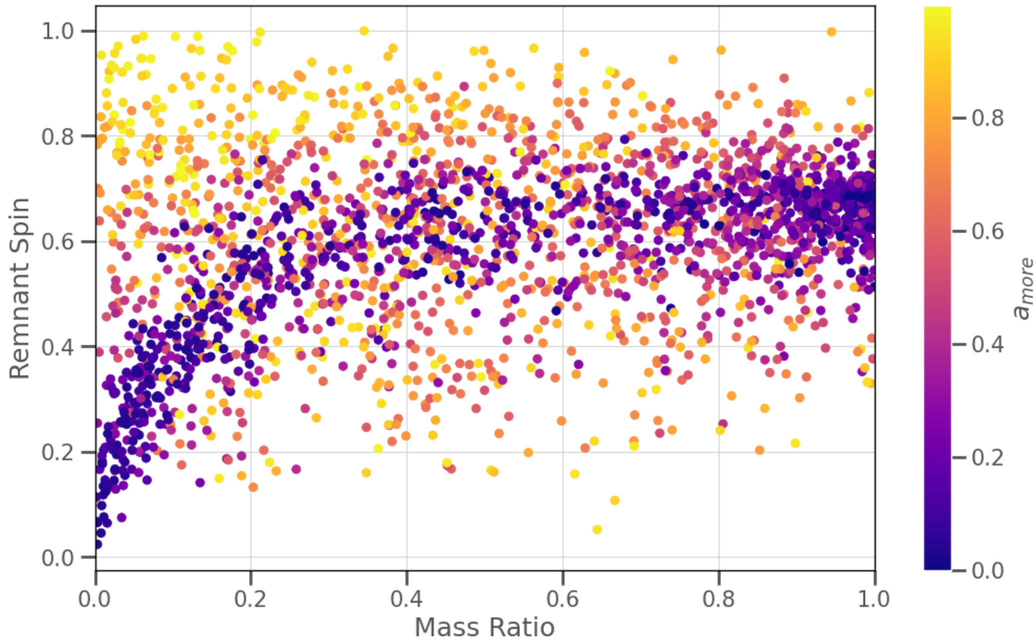

Figure 4.7: *Remnant black hole spin as a function of merger mass ratio for the* 10^5*&Adhesion model. Each point represents a merger event, with the x-axis showing the mass ratio* $q = \frac{M_{less}}{M_{more}}$, *and the y-axis showing the dimensionless spin of the resulting remnant black hole. Points are colour-coded by the spin* a_{more} *of the more massive progenitor black hole at the time of the merger, as indicated by the colourbar. A positive correlation is evident between mass ratio and remnant spin. The influence of the more massive black hole's spin is also visible, especially at low mass ratios where its contribution dominates the final spin.*

moderate-spin population observed at later stages. These moderate remnant spins may also serve as seeds for further spin-up via coherent accretion, depending on subsequent gas availability and inflow structure.

Figure 4.6 shows the distribution of black hole merger mass ratios in the 10^5&Adhesion model. The mass ratio is defined as $q = \frac{M_{less}}{M_{more}}$, where M_{less} and M_{more} are the masses of the less and more massive black holes in each merger event. A mass ratio of 1 corresponds to an equal-mass merger, while lower values indicate more unequal mass mergers.

The distribution is relatively broad, with a notable rise toward both low ($q < 0.1$) and high ($q > 0.9$) mass ratios. This indicates that the 10^5&Adhesion model includes both extreme ratio and near-equal mass mergers. The central dip suggests that intermediate mass ratio mergers ($q \sim 0.4 - 0.7$) are less common. The significant fraction of near-equal mass mergers is particularly relevant, as such events are known to produce higher post-merger spins when the pre-merger spins are low or aligned (e.g., Rezzolla et al. 2008). Conversely, low-q mergers are generally less efficient at spinning up the remnant black hole (e.g., Lousto & Zlochower 2011). These trends are consistent with the post-merger spin distribution seen in Figure 4.5, where remnants peak at moderate spin values around $a \sim 0.7$.

Figure 4.7 shows the relationship between black hole merger mass ratio and the resulting remnant spin in the 10^5&Adhesion model, with point colour indicating the dimensionless spin a_{more} of the more massive progenitor. Each point represents an individual merger event, with the x-axis indicating the mass ratio $q = \frac{M_{less}}{M_{more}}$ (ranging from very unequal mergers at low q to equal-mass mergers at $q \sim 1$), and the y-axis showing the dimensionless spin of the remnant black hole produced by that event.

There is a clear trend of increasing remnant spin with increasing mass ratio, with particularly low-spin remnants associated with highly unequal mergers ($q \lesssim 0.1$). However, at these low mass ratios, remnant spins span a wide range with a strong dependence on the spin of the primary black hole. This is reflected in the colour gradient: for small q, the distribution of remnant spins tracks a_{more} closely. This aligns with theoretical expectations (e.g., Lousto & Zlochower 2011), as low-q mergers typically involve a smaller contribution of orbital angular momentum and result in less spin-up of the remnant. At higher q, this dependence diminishes, and remnant spins converge toward a narrower range, around $a \sim 0.7$, largely independent of the progenitor spins. This behaviour is consistent with semi-analytic prescriptions (e.g., Rezzolla et al. 2008), which predict that the remnant spin becomes increasingly insensitive to progenitor spin alignment and magnitude as the mass ratio approaches unity. The concentration of points around $a \sim 0.7$ in this high q region reflects the high frequency of near-equal mass mergers in the 10^5&Adhesion model (see Figure 4.6), and suggests that these mergers dominate the population of moderately spinning remnants.

The distribution of spin-spin alignment angles shown in Figure 4.2 demonstrates a strong preference for aligned spin orientations between the two

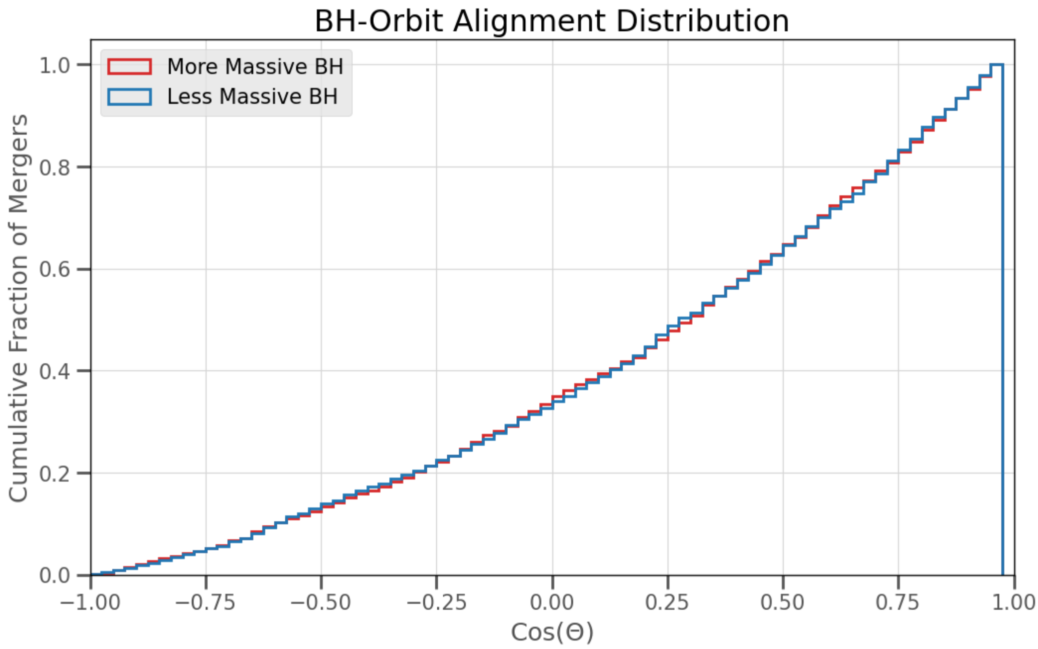

Figure 4.8: *Cumulative distribution of the cosine of the angle between each black hole's spin vector and the binary orbital angular momentum vector immediately prior to merger, for the* 10^5*&Adhesion model. The red histogram shows the more massive black hole in each merger, while the blue histogram corresponds to the less massive black hole. Both distributions show a clear preference for aligned configurations.*

progenitor black holes at the point of merger in the 10^5&Adhesion model. The figure plots the probability density function of $\cos\theta$, where θ is the angle between the spin vectors of the two black holes immediately before merger. A value of $\cos\theta = 1$ corresponds to perfectly aligned spins, $\cos\theta = -1$ to perfectly anti-aligned spins, and $\cos\theta = 0$ to orthogonal spin orientations. The sharp peak at $\cos\theta \approx 1$ indicates that the majority of black hole pairs merge with their spins closely aligned.

This result suggests that the mergers in this model predominantly occur under conditions that facilitate spin alignment, such as prolonged gas-rich inspirals in which accretion torques act to reorient black hole spins toward the orbital angular momentum vector. In such environments, the angular momentum of the circumbinary disc can couple efficiently to the black hole spin vectors, driving alignment over relatively short timescales.

This strong alignment signature has important implications for the spin magnitude and orientation of the post-merger remnant, as well as for the recoil velocity imparted via gravitational wave emission. Aligned spins tend to suppress recoil kicks and enhance remnant spin magnitudes (Campanelli et al. 2007).

Figure 4.8 presents the distribution of the cosine of the angle between each progenitor black hole's spin vector and the binary's orbital angular momentum vector, measured just prior to merger in the 10^5&Adhesion model. The distribution is shown separately for the more massive (red) and less massive (blue) members of each binary system. Values of $\cos\theta$ near 1 indicate near-perfect alignment between the spin and orbital angular momentum vectors, while values near -1 correspond to strong anti-alignment. A uniform, isotropic distribution of spin orientations would be expected to yield a flat histogram across the interval [-1,1].

Both populations exhibit a clear excess of systems with $\cos\theta > 0$, indicating a statistical preference for alignment between black hole spin vectors and the orbital angular momentum of the binary. Only $\sim 30\%$ of systems are misaligned ($\cos\theta < 0$, consistent with strong coupling to the angular momentum structure of the surrounding gas.

This preferential alignment suggests that both black holes are repeatedly torqued toward the same large-scale angular momentum, and their orbits are set by the same galactic potential, so spin-orbit coupling is regulated by these shared conditions. The mass dependence observed may be a consequence of the more massive black holes dominating the angular momentum budget of the system or experiencing more sustained, coherent accretion episodes.

4.3 Spin Magnitude

Figure 4.9 shows the median spin magnitude of black holes as a function of redshift in the 10^5&Adhesion model, separated into four distinct black hole mass bins. The shaded regions indicate the interquartile range, capturing the spread in spin magnitudes within each mass category at a given time.

Lower-mass black holes ($10^5 < \frac{M_{BH}}{M_\odot} < 10^6$) consistently maintain relatively low median spins ($a \sim 0.1 - 0.2$) throughout cosmic time. This limited spin-up

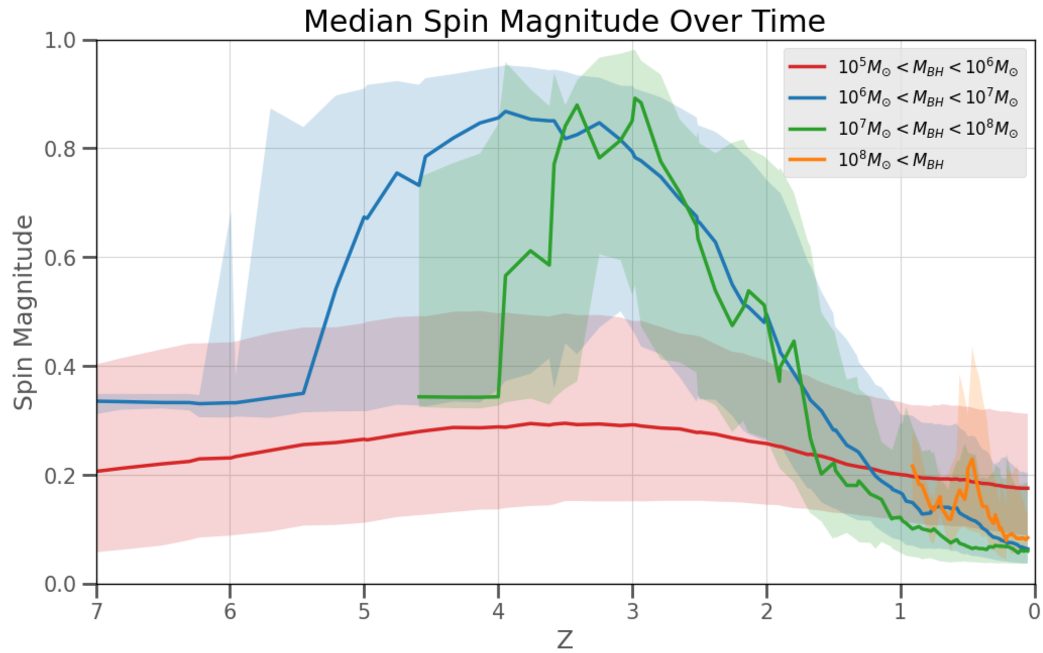

Figure 4.9: *Median black hole spin magnitude as a function of redshift for the* 10^5*&Adhesion model, separated into four black hole mass bins:* $10^5 < \frac{M_{BH}}{M_\odot} < 10^6$ *(red),* $10^6 < \frac{M_{BH}}{M_\odot} < 10^7$ *(blue),* $10^7 < \frac{M_{BH}}{M_\odot} < 10^8$ *(green), and* $\frac{M_{BH}}{M_\odot} > 10^8$ *(orange). Each solid line shows the median spin magnitude within the bin at each redshift, while the shaded regions represent the interquartile range.*

reflects their more stochastic and less prolonged accretion histories. As black holes grow and transition into higher mass bins ($10^6 < \frac{M_{BH}}{M_\odot} < 10^6$ and$10^7 < \frac{M_{BH}}{M_\odot} < 10^8$, median spins peak notably at intermediate redshifts ($z \sim 4-6$), reaching values of $a \sim 0.7-0.8$. This spin-up can be attributed to phases of rapid, coherent gas accretion, possibly corresponding to periods of active galaxy assembly and increased gas inflow at earlier epochs. Subsequently, median spins decrease significantly toward lower redshifts, converging toward low values ($a < 0.2$) in the highest mass bin ($\frac{M_{BH}}{M_\odot} > 10^8$) at $z \sim 0$. This decline results from sustained periods of lower relative accretion rates, keeping the black holes accreting in the hot accretion regime, where jets efficiently transfer rotational energy out of the black holes, at later cosmic times.

The increasing interquartile range at intermediate redshifts for higher mass bins indicates a growing diversity of accretion and merger histories during peak cosmic

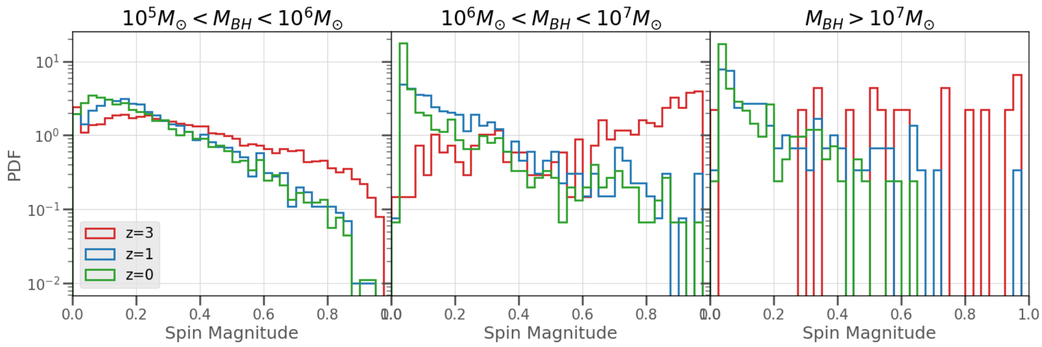

Figure 4.10: *Spin magnitude distributions for black holes in the* 10^5*&Adhesion model at three selected redshifts (*$z = 3$*,* $z = 1$*, and* $z = 0$*). Distributions are shown for three distinct mass bins:* $10^5 < \frac{M_{BH}}{M_\odot} < 10^6$ *(left),* $10^6 < \frac{M_{BH}}{M_\odot} < 10^7$ *(middle), and* $\frac{M_{BH}}{M_\odot} > 10^7$ *(right). The different redshifts are shown in red (*$z = 3$*), blue (*$z = 1$*), and green (*$z = 0$*).*

star formation periods. This spread narrows at late times as low accretion rates act to reduce spins, consistent with theoretical expectations for MBH populations (e.g., Volonteri et al. 2013; Dubois et al. 2014).

Having examined the median spin evolution over cosmic time, we now focus on the complete spin distributions at specific snapshots ($z = 3$, $z = 1$, and $z = 0$), shown in Figure 4.10. Because the highest mass bin sample at each snapshot is small, the histogram is dominated by Poisson noise. This is most evident at the intermediate and high redshift. The following statements focus on $M_{BH} < 10^7M_\odot$, where statistics are more robust.

At high redshift ($z = 3$), spins exhibit a broad distribution across mass bins, reflecting diverse accretion histories and merger events during early galaxy formation stages. Moving toward intermediate ($z = 1$) and low ($z = 0$) redshifts, we see a systematic shift toward lower spin magnitudes across all mass bins. By the present day, the majority of black holes across all masses have spins clustered near zero, with the distribution strongly peaked at very low spin values ($a \lesssim 0.1$). This reflects the increasing prevalence of lower accretion rates. Particularly notable is the narrowing of the spin distribution at $z = 0$ for the higher mass bin ($10^6 < \frac{M_{BH}}{M_\odot} > 10^7$), highlighting how even more massive black holes preferentially evolve towards very low spins.

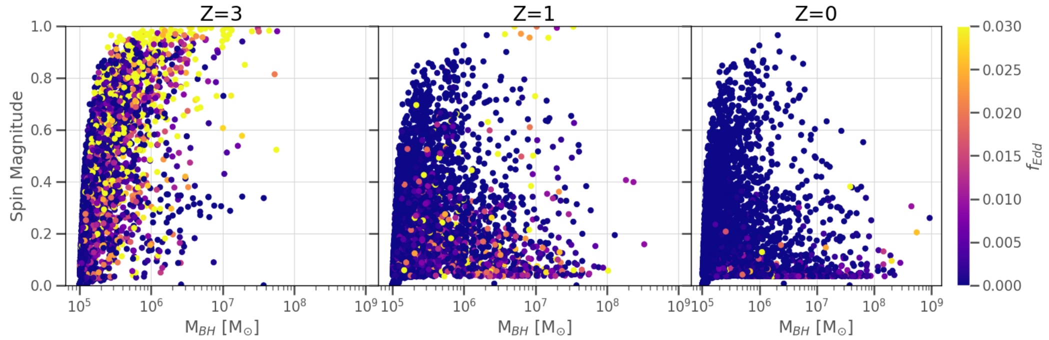

Figure 4.11: *Spin magnitude versus black hole mass in the* 10^5*&Adhesion model at three cosmic epochs* ($z = 3$, $z = 1$, $z = 0$). *Points are colour-coded by the instantaneous Eddington ratio* f_{Edd}. *Yellow points mark systems at or above the critical threshold* $f_{crit} = 0.03$, *identifying black holes currently accreting in the thin-disk regime; darker colours indicate progressively lower* f_{Edd} *values associated with hot accretion states. These panels illustrate the evolution of the black hole population from higher-*f_{Edd} *systems at early times to lower-*f_{Edd} *systems at the present epoch and how this impacts black hole spin.*

The overall evolution captured in these snapshots aligns closely with observational evidence of decreasing accretion rates and evolving accretion regimes over cosmic time (Fanidakis et al. 2011), providing insights into the mechanisms driving black hole spin evolution.

Figure 4.11 shows the distribution of spin magnitudes for the full population of black holes in the 10^5&Adhesion model as a function of their mass at three epochs ($z = 3, 1, 0$), colour-coded by instantaneous Eddington ratio (f_{Edd}). At high redshift ($z = 3$), a significant portion of the population exhibits higher spins ($a > 0.6$) and elevated accretion rates near or above the critical Eddington fraction ($f_{crit} = 0.03$), consistent with thin-disk accretion dominating at early epochs. At this stage, many black holes undergo coherent accretion, efficiently spinning them up.

By $z = 1$, the mass distribution broadens and spin shifts towards lower magnitudes, reflecting a declining fraction of black holes in the thin-disk regime as accretion rates generally decrease. Lower Eddington ratios dominate, corresponding to more frequent hot accretion states, where accretion-driven spin-up is inefficient.

By $z = 0$, the vast majority of black holes display very low instantaneous

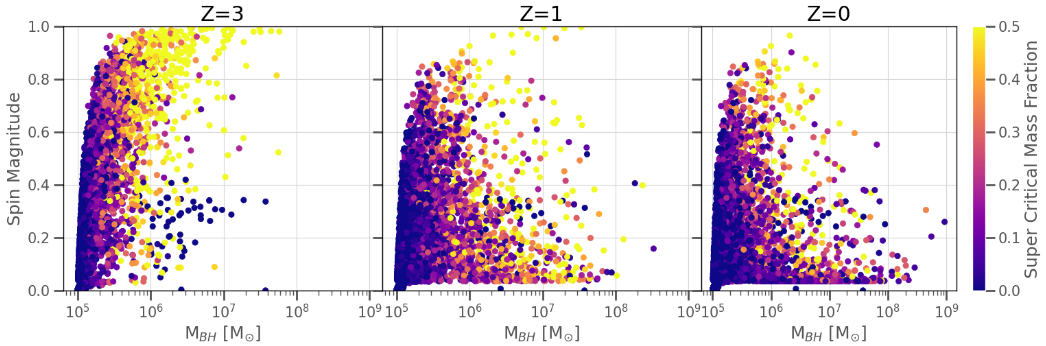

Figure 4.12: *Spin magnitude versus black hole mass in the* 10^5*&Adhesion model at three redshifts (*$z = 3$*,* $z = 1$*,* $z = 0$*), colour-coded by the super critical mass fraction accumulated up to each epoch. Yellow symbols correspond to black holes that have accumulated half or more of their accreted mass through coherent, thin-disk mode accretion. This figure links present-day spin to the levels of coherent accretion of each black hole across cosmic time.*

Eddington ratios, characteristic of hot, radiatively inefficient accretion conditions. This epoch is dominated by relatively low spins ($a \lesssim 0.4$), especially at higher masses, as chaotic accretion and spin-down mechanisms such as jet production increasingly suppress spin growth.

Overall, this figure demonstrates a clear evolutionary sequence wherein black hole spin magnitudes correlate strongly with instantaneous accretion conditions. It reinforces the picture that high spins are built and maintained during phases of near-Eddington, thin-disk accretion common at early times, whereas declining accretion rates and increasingly chaotic growth at later epochs drive the population toward low spins–especially among the most massive black holes.

Figure 4.12 demonstrates the correlation between black hole spin magnitude and the fraction of mass accumulated through coherent, thin-disk accretion (super critical mass fraction) across cosmic time. At high redshift ($z = 3$), many black holes exhibit high spin magnitudes and elevated super critical mass fractions, indicating substantial spin-up driven by coherent accretion during this period of rapid, gas-rich growth (e.g. Dubois et al. 2014; Sesana et al. 2014). Systems coloured yellow, representing those that have accumulated at least half their mass through sustained thin-disk accretion, predominantly populate the high-spin region.

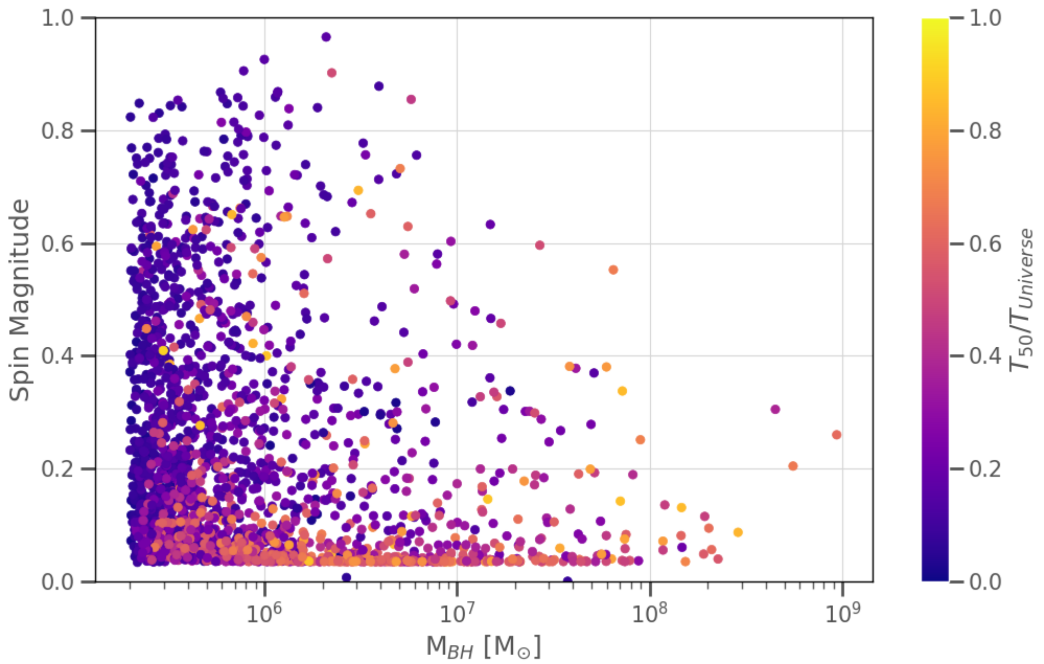

Figure 4.13: *Spin magnitude versus black hole mass in the* 10^5*&Adhesion model at* $z = 0$*. Points are colour-coded by the scaled half-mass assembly time* $\frac{T_{50}}{T_{Universe}}$*, where* T_{50} *is the cosmic age when the black hole first reached 50% of its current mass and* $T_{Universe}$ *is the age of the Universe. Low values indicate systems that assembled half their mass early in the Universe; values near 1 assembled half their mass comparatively late. This figure links present-day spin to the timing of past mass build-up.*

At intermediate redshift ($z = 1$), there is a noticeable decline in the number of systems with large super critical mass fractions. Although some black holes still exhibit high spin magnitudes, fewer maintain extensive coherent accretion histories, reflecting the transition from coherent, rapid accretion in the early Universe to increasingly chaotic and low-efficiency accretion flows at later times (e.g. Fanidakis et al. 2011; Volonteri et al. 2013).

By the present epoch ($z = 0$), black holes predominantly occupy regions of lower spin magnitude and reduced super critical mass fractions. While isolated examples of high-spin, high-coherence systems persist, they are substantially outnumbered by black holes whose growth has been dominated by chaotic accretion events. This highlights the cumulative impact of reduced coherent accretion over cosmic time, reflecting the gas-poor, merger-driven, or episodic accretion conditions characteristic of the local Universe (e.g. Ho 2008).

Overall, this figure reinforces the direct link between sustained coherent (thin-disk) accretion and black hole spin evolution, highlighting how the spin magnitude encodes vital information about a black hole's cumulative accretion history.

Figure 4.13 displays the present-day ($z = 0$) spin magnitudes and black hole masses, colour-coded by the scaled half-mass assembly time, $T_{50}/T_{Universe}$. Black holes that assembled half of their current mass earlier in cosmic history (darker points) generally populate a broader range of spins, including numerous moderate to high-spin systems. This result likely reflects more sustained coherent accretion phases during earlier, gas-rich epochs that allowed substantial spin-up. Conversely, black holes with higher values of $T_{50}/T_{Universe}$ (lighter points), indicating relatively late mass assembly, predominantly exhibit lower spin magnitudes. This behaviour is consistent with later black hole growth occurring in a less gas-rich, more merger-driven, and chaotic accretion environment, which is less conducive to sustained coherent spin-up. The distribution highlights how the timing of black hole mass assembly significantly impacts their spin evolution, providing a direct observational link between black hole spin and the formation history of MBHs.

While figure 4.13 showed how *when* a black hole assembled its mass relates to present spin, figure 4.14 adds a complementary view: how that mass was acquired–specifically, the fraction of its lifetime spent in the thin-disk (radiatively efficient) accretion regime. At the highest redshift, many black holes–especially those with masses above $M_{BH} \gtrsim 10^6 M_\odot$–have spent a large fraction of their lifetimes in thin-disk mode (super critical lifetime fraction$\gtrsim 0.33$, possibly up to ~ 1 given the young Universe) and correspondingly attain high spin magnitudes ($a \gtrsim 0.6$). This reflects the early Universe's gas-rich conditions, where prolonged, coherent accretion episodes dominate growth and efficiently spin up black holes (e.g. Dubois et al. 2014). Systems with smaller thin-disk exposure (super critical lifetime fraction$\lesssim 0.15$) populate the lower-spin tail. This association between high lifetime fractions and high spin is expected: extended coherent, radiatively efficient accretion efficiently spins up Kerr black holes (Bardeen 1970; Thorne 1974).

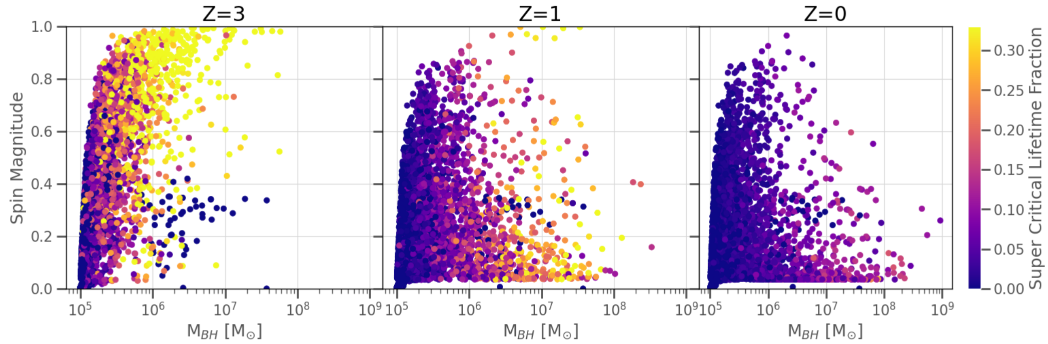

Figure 4.14: *Spin magnitude versus black hole mass in the* 10^5*&Adhesion model at three redshifts (*$z = 3$*,* $z = 1$*,* $z = 0$*), colour-coded by the super critical lifetime fraction, i.e. the fraction of cosmic time (from birth to the snapshot epoch) that each black hole has spent accreting above the critical Eddington ratio* $f_{crit} = 0.03$ *(thin-disk, radiatively efficient regime). Dark colours mark systems that have spent little of their lives in thin-disk mode (mostly hot accretion), whereas yellow points have experienced prolonged cumulative thin-disk growth–at least* $\frac{1}{3}$ *of their lifetime.*

By $z = 1$ the prevalence of long-lived thin-disk episodes declines, and the correlation between the super critical lifetime fraction and spin weakens. Although a handful of intermediate-mass black holes still reach high thin-disk fractions and spins, the bulk of the population–even for intermediate to high spins–has accumulated only sporadic thin-disk growth (purple points). Black holes with high thin-disk fractions are still seen across the spin spectrum, with many of them lying at low spin values ($a \lesssim 0.4$) This indicates that intervening chaotic or misaligned accretion episodes have begun to decay spins built up during earlier coherent phases (King & Pringle 2006).

At $z = 0$ nearly all black holes have spent a small fraction of their lifetimes in the thin-disk regime. Even those with moderate spins ($a \sim 0.3$) exhibit thin-disk fractions below 10%. The lack of sustained thin-disk growth in the local Universe is attributable to lower gas fractions, more merger-driven or feedback-regulated accretion, and longer quiescent intervals, which together inhibit coherent spin-up.

Across redshift, the super critical lifetime fraction acts as a spin memory parameter only when subsequent chaotic evolution is limited. Early-universe black holes that spend much of their short lifetimes in thin-disk mode reach and retain high spins; with time, cumulative episodes of misaligned and low-rate accretion drive convergence to low spins, decoupling present spin from past thin-disk duty cycle. This evolution explains why high spins are expected (and observed) primarily in rapidly accreting quasars at high redshift.

4.4 Observation Bias

In practice, measurements of SMBH spin rely on bright AGN whose high accretion luminosities make them accessible to X-ray reflection or continuum-fitting techniques(Reynolds 2014). In this subsection, we look at how imposing a simple luminosity cut ($L_{bol} > 10^{42}\mathrm{erg\,s^{-1}}$) skews the underlying spin-mass population. By comparing the full simulation sample to the luminosity-selected subset, we can assess the degree to which observational selection effects bias our inferences about typical black hole spin magnitudes and their growth histories.

Figure 4.15 shows the median spin magnitude as a function of redshift for luminous systems in the 10^5&Adhesion model, divided into the same mass bins used previously in Figure 4.9.

Compared to the full population shown in Figure 4.9, luminous black holes exhibit systematically higher median spins at all redshifts, especially at early times ($z > 4$). This suggests that observationally selected AGN-like black holes preferentially populate the high-spin tail of the distribution. The intermediate mass bins ($10^6 < \frac{M_{BH}}{M_\odot} < 10^8$) show a prominent peak in median spin ($a \sim 0.8$) at high redshifts ($z \sim 4-6$), reflecting epochs of intense and well-aligned gas inflow that efficiently spins up black holes. However, median spin values decrease substantially toward lower redshifts, approaching low-to-moderate spins at the present epoch, which implies a transition toward more chaotic, misaligned accretion events driving spin-down.

The narrowing of the interquartile ranges at lower redshifts indicates greater uniformity in the spin evolution of luminous black holes at late cosmic times. This reduced diversity likely reflects the convergence toward conditions which promote lower relative black hole accretion rates, consistent with theoretical expectations and observational results from local AGN populations (Marconi et al. 2004; Georgakakis et al. 2017). Overall, these results highlight the strong influence of

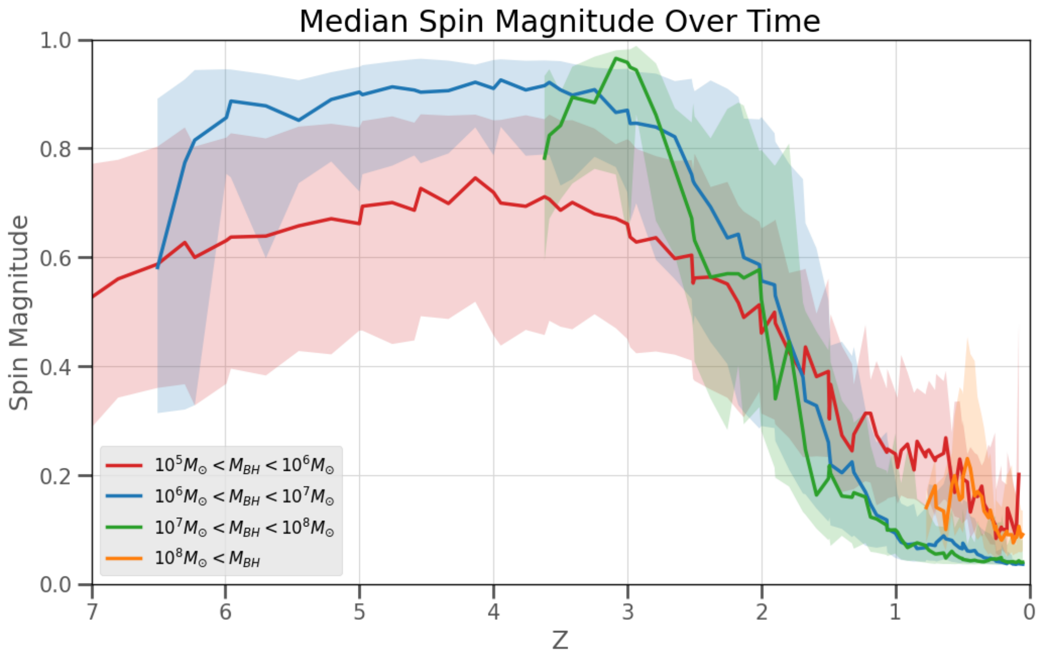

Figure 4.15: *Median spin magnitude evolution over cosmic time for luminous black holes ($L_{bol} > 10^{42}\mathrm{erg\,s^{-1}}$) in the 10^5&Adhesion model, shown across four distinct black hole mass bins: $10^5 < \frac{M_{BH}}{M_\odot} < 10^6$ (red), $10^6 < \frac{M_{BH}}{M_\odot} < 10^7$ (blue), $10^7 < \frac{M_{BH}}{M_\odot} < 10^8$ (green), and $\frac{M_{BH}}{M_\odot} > 10^8$ (orange). Solid lines indicate the median spin magnitude within each bin, while the shaded areas represent the interquartile range. Restricting the sample to observationally luminous black holes emphasizes spin trends most directly comparable with active black hole populations observed at various cosmic epochs.*

luminosity selection when comparing simulated spin distributions with observed populations of active black holes across cosmic history.

Figure 4.16 isolates the subset of black holes that would be detectable in spin measurements and shows how their spin-mass-growth correlations differ from the full population at $z = 1$.

In the top-left panel (instantaneous f_{Edd}), nearly all luminous black holes lie at moderate to high Eddington ratios ($f_{Edd} \gtrsim 0.01$), in stark contrast to the full sample (Figure 4.11), which included many low-accretion ($f_{Edd} < 10^{-3}$) systems. This selection bias toward high-$\dot{M}$ naturally favours higher spin magnitudes, since thin-disk accretion efficiently spins up black holes.

Similarly, the bottom-left panel (scaled half-mass assembly time $\frac{T_{50}}{T_{Universe}}$) shows

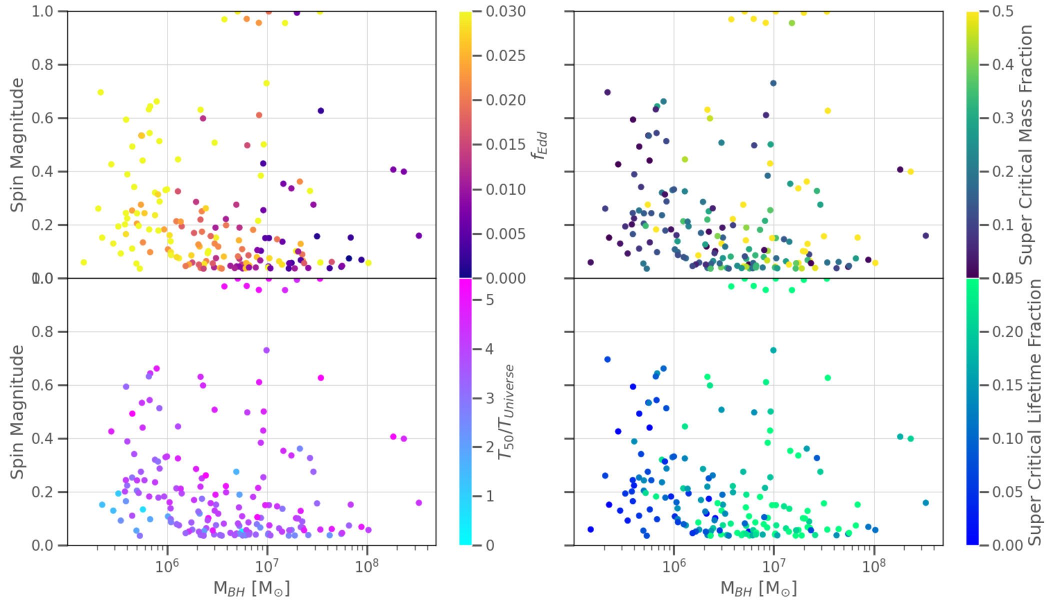

Figure 4.16: *Spin magnitude versus black hole mass at $z = 1$ for luminous black holes ($L_{bol} > 10^{42}$erg s^{-1} in the 10^5 $M_\odot$ Adhesion model. Each panel is colour-coded by a different growth-related quantity: instantaneous Eddington ratio f_{Edd} (top left), scaled half-mass assembly time $\frac{T_{50}}{T_{Universe}}$ (bottom left), super critical mass fraction (top right), and super critical lifetime fraction (bottom right). This figure allows for a direct comparison between the luminous black hole population and the, previously shown, full population, revealing how selection biases toward high-luminosity systems may skew interpretations of spin–growth correlations.*

that observed AGN are skewed toward earlier-assembling black holes–points are concentrated at lower $\frac{T_{50}}{T_{Universe}}$ compared to the full population. In the full population plot (Figure 4.13), late-assemblers spanned spins from low to moderate, but here they are under-represented.

The top-right panel (super-critical mass fraction) highlights that luminous systems span the full range of super-critical mass fractions, indicating that bright sources are not exclusively those that have grown predominantly through coherent, thin-disk accretion. While higher spins tend to be associated with larger fractions, substantial scatter persists: many low-spin luminous black holes show sizeable coherent fractions, and conversely some moderate-spin systems exhibit predominantly chaotic/hot growth histories. Thus instantaneous brightness does not uniquely trace the cumulative importance of thin-disk accretion; both coherent and incoherent growth pathways are represented within the observable sample.

Finally, the bottom-right panel (thin-disk lifetime fraction) confirms that luminous AGN have spent a substantial fraction ($\gtrsim 10\%$) of their lives in the radiatively efficient mode. In contrast, the full population at $z = 1$ (Figure 4.14) had a broad mix of low and high lifetime fractions.

Overall, compared to the unbiased, full population, the luminosity-selected sample is clearly biased toward black holes with higher instantaneous accretion rates, earlier mass assembly, and more extensive coherent accretion histories–properties that all drive higher spin magnitudes. Any observational inference about the typical spin distribution must therefore account for these possible selection effects.

Chapter 5

Discussion

5.1 Physical Interpretation and Theoretical Context

The post-processed spin histories extracted from the Romulus simulation support a picture in which coherent thin-disk accretion at early times is the dominant channel for producing high spin magnitudes, while subsequent low-Eddington, hot/chaotic phases erode the initial diversity and drives the population toward low spins by the present epoch. At $z \gtrsim 3$ the majority of black holes in the 10^5&Adhesion model experience sustained periods above the critical Eddington fraction, building both mass and spin efficiently. This is consistent with analytic expectations that prolonged prograde accretion at high f_{Edd} rapidly spins up modest-mass seeds toward the Thorne limit (e.g. Bardeen 1970; Thorne 1974). As cosmic gas fractions decline and feedback regulates inflows, the duty cycle of thin-disk episodes diminishes: by $z = 0$ most systems have spent only a few percent of their lifetimes in the coherent regime, and the cumulative action of hot-mode inefficiently spinning up black holes and jet torques produces the broad low-spin peak we observe. This transition matches semi-analytic treatments in which stochastic angular momentum flips or intermittent accretion episodes drive a decrease in overall spin (Volonteri et al. 2005; Fanidakis et al. 2011).

Our model comparison isolates how uncertain off-centre physics modulates this baseline evolution. The chaotic prescription, which suppresses hydrodynamic spin-up while black holes wander, yields the lowest high-redshift spins and the most rapid convergence to low-spin states, effectively bracketing a lower limit on coherent growth. Conversely, the assumed central model provides an upper limit by allowing uninterrupted alignment and spin-up even during displacement. The gas adhesion model sits between these extremes physically–preserving a "memory" disc enables continued spin-up off-centre, but the frozen orientation prevents perfect tracking of evolving central gas–but produces similar results to the assumed central presecription. This ordering mirrors the analytic dichotomy between fully coherent disc accretion and "chaotic accretion" scenarios composed of randomly oriented, short episodes (King & Pringle 2006; Sesana et al. 2014), placing quantitative bounds on how much black hole wandering can dilute spin growth in a cosmological context.

The seed mass experiment further clarifies that early high spins in the 10^5&Adhesion model owe to elevated effective Eddington ratios produced by lowering the initial mass while holding $\dot{M}$ fixed. Smaller seeds traverse a phase of persistently high f_{Edd}, spending a larger fraction of their early lifetimes in the thin-disk regime and reaching high spins before self-regulation sets in. This behaviour parallels semi-analytic spin-growth models in which the spin-up timescale scales with $\frac{M}{\dot{M}}$ and can be "short-circuited" for low-mass seeds embedded in gas-rich hosts (e.g. Dubois et al. 2014; Ricarte et al. 2023). Once masses exceed $\sim 10^{7-8} M_\odot$, the memory of the initial seed is largely erased, and different prescriptions converge-demonstrating that present-day spin distributions are more sensitive to the integrated coherence of early accretion than to the precise seeding choice.

Mergers play a comparatively secondary role in setting the bulk spin distribution: remnant spins cluster near moderate values expected from aligned or mildly misaligned mass ratios, and the strong pre-merger spin and orbital alignment we measure limits the production of low-spin remnants. This agrees with fitting-formula predictions (e.g. Rezzolla et al. 2008) applied to semi-analytic alignment models (Bogdanovic et al. 2007), and implies that, within the explored parameter space, stochastic accretion mode transitions–not merger-driven spin resets–govern the long-term decline in ensemble spin. Overall, our results integrate naturally into existing analytic and semi-analytic frameworks while providing simulation-anchored constraints on the magnitude of off-centre suppression and the persistence of early coherent spin-up.

We can gauge the relative importance of mergers and gas accretion by comparing the remnant spin distribution from individual coalescences (Fig. 4.5) with the full black hole population spin distributions at fixed redshift (Fig. 4.10). The remnant PDF is narrowly unimodal, peaking at $a \sim 0.7$ with relatively little support at very low or very high spins. By contrast, the population PDFs at the same epochs are much broader, with a substantial low-spin tail and a clear redshift evolution toward lower spins at late times. This mismatch in shape persists at all redshifts considered.

If mergers dominated the setting of spins, repeated coalescences would drive the population toward the same form as the remnant distribution (with some modest broadening from accretion driven spin-up/spin-down), and the two PDFs would look similar. The fact that the population distributions are both broader and strongly time-dependent indicates that inter-merger accretion must substantially move spins between merger events–spinning black holes up during coherent, thin-disk phases and spinning them down during hot/MAD phases via jet torques–so that the time-averaged spin demographics are governed primarily by the accretion duty cycle rather than by the instantaneous remnant outcomes. While this comparison does not yield a precise fractional attribution, it implies that accretion is the dominant channel shaping the ensemble spin distribution.

Finally, a central and novel contribution of this thesis is the explicit and detailed treatment of angular momentum accretion by wandering (off-centre) black holes. In many cosmological suites, black holes are artificially recentered each timestep to the local gravitational potential minimum, keeping them near galaxy centres for an unrealistically large fraction of their lifetimes; by contrast, Romulus replaces recentering with a physically motivated subgrid dynamical-friction prescription that permits long-lived off-centre orbits and realistic wandering (Tremmel et al. 2015). Here we directly address this limitation by systematically exploring three physically distinct off-centre accretion scenarios–assumed central, chaotic, and gas adhesion–each capturing a different extreme of unresolved accretion geometry and coherence. Our findings reveal that off-centre accretion is not merely a minor correction but a critical driver of the global spin distribution, significantly modulating the black hole spin demographics across cosmic time. Thus, we demonstrate that explicitly modelling off-centre growth phases is essential for producing realistic and predictive spin evolution models, underscoring a previously overlooked aspect of SMBH-galaxy co-evolution.

5.2 Relation to Previous Spin-Evolution Studies

Our results broadly support the canonical picture emerging from prior post-processed and live spin studies: prolonged, coherent thin-disk accretion at early times drives rapid spin-up, whereas the growing prevalence of low-f_{Edd}, hot/chaotic phases at late times reduces spins toward $z = 0$ (e.g. Dubois et al. 2014). In the luminosity-selected subset, we likewise recover a high-spin tail comparable to that inferred for bright AGN samples, consistent with the view that observed spins are drawn from a biased, high-accretion fraction of the population. Where we differ from several earlier works is in the global spin distribution of the full population: our ensemble median spin at low redshift is lower, and the high-mass, high-spin incidence is more modest.

Two modelling choices plausibly explain these differences. First, our post-processing adopts a relatively long averaging window for accretion rates, whereas prior post-processed studies typically use shorter cadences or instantaneous rates (e.g. Dubois et al. 2014; Bustamante & Springel 2019), which will be adopted in further work on this model for black hole spin. Because spin responds to the instantaneous angular momentum flux, temporal smoothing damps short, coherent episodes that can efficiently spin up seeds and intermediate-mass black holes, biasing the population toward lower spin magnitudes. Second, Romulus seeds black holes at $10^6 M_\odot$, skipping the possible earlier growth where $\Delta M/M_{BH}$ is largest and spin-up is most rapid; while our $10^5 M_\odot$ seed mass experiment boosts relative accretion rates to account for this, it does not account for $10^6 M_\odot$ seeds forming later than $10^5 M_\odot$ seeds, so some early accretion in gas-rich environments is still missed. Several comparison works initialise at lower masses or track early growth more explicitly, naturally producing faster rise to high spins (e.g. Sala et al. 2024; Beckmann et al. 2025). These two effects together likely account for much of the discrepancy at the low-mass and early-time end.

Compared to live spin implementations, our results show similar trends but fail to produce very high magnitude spins, particularly in high mass, luminous AGN. Live models that couple spin to radiative/jet efficiencies and update orientation during the run typically report faster early spin-up and more extended high-spin tails, especially at $z \gtrsim 2$ (e.g. Sala et al. 2024; Beckmann et al. 2025). This is expected: on-the-fly coupling creates feedback loops (e.g. spin-dependent jet power) that can accelerate the transition between accretion states and preserve high spins in gas-rich phases, whereas our post-processing cannot alter the underlying gas supply in real time. Once luminosity selection is imposed, our results resemble those reported by live-spin studies for observable AGN in that the observable population is biased towards higher spins (Beckmann et al. 2025), suggesting that selection effects erase some of the methodological differences in the bright subset.

On spin-galaxy alignment, our finding of strong pre-merger alignment with the host's central angular momentum is consistent with both post-processed and live-spin results that tie the disc axis to kpc-scale gas in settled systems (e.g. Beckmann et al. 2024; Peirani et al. 2024). Where our work extends the literature is in explicitly testing off-centre phases: by bracketing accretion geometry for wandering black holes (assumed-central, chaotic, adhesion), we show that displaced growth can materially shift spin demographics even if alignment remains strong. In particular, the chaotic prescription lowers early spins and accelerates convergence to low a, while the adhesion model preserves some coherent memory off-centre. This emphasis on wandering complements prior studies that either reposition black holes to galaxy centres or lack explicit off-nuclear treatments.

Overall, several previous analyses reproduce the observed prevalence of high spin AGN more readily than our predictions. Given differences in time averaging, seed mass, and whether spin-feedback coupling is live or post-processed, this residual discrepancy is unsurprising and points to clear tests: shorter variability timescales in the spin pipeline, finer treatment of early ($< 10^6 M_\odot$) growth, and explicit modelling of off-centre reservoirs. In this sense, our results are complementary to Dubois et al. (2014) and Bustamante & Springel (2019) (post-processed), and to Sala et al. (2024); Beckmann et al. (2025) (live), while our off-centre analysis speaks directly to alignment-focused comparisons in Beckmann et al. (2024); Peirani et al. (2024). Together, these studies outline a convergent physical picture while clarifying which assumptions most strongly control the quantitative spin demographics.

5.3 Comparison with Observations

Observed SMBH spins inferred from X-ray reflection spectroscopy are heavily concentrated in luminous, nearby Seyfert nuclei and a handful of bright quasars, typically with masses $M_{BH} \sim 10^{6-8} M_{\odot}$ and moderate-to-high Eddington ratios (e.g. Reynolds 2014). These samples show an apparent prevalence of high spins ($a \gtrsim 0.7$), which has sometimes been interpreted as evidence that most active SMBHs retain long-lived coherent accretion. In our full simulated population, by contrast, the median spin at $z = 0$ is low because the majority of black holes spend only a small fraction of their lifetimes in the thin-disk regime at late times. However, when we impose a bolometric luminosity cut, the luminosity-selected subsample recovers systematically higher spins and earlier assembly histories: precisely the bias seen in current spin catalogues. Thus the high observed spins are more consistent with our model once selection effects are applied; they do not require the entire underlying population to be rapidly spinning.

The predicted spin-mass relation likewise exhibits an observational bias. In the intrinsic distribution, higher-mass ($> 10^8 M_{\odot}$) black holes have already transitioned to low Eddington ratios and experienced substantial spin-down. Yet these objects often fall below the luminosity threshold for reliable reflection spectroscopy and are therefore under-represented observationally. The observed mass range is dominated by systems in which thin-disk duty cycles remain non-negligible, naturally skewing the measured spins upward relative to the global mean. A similar effect arises in the spin-Eddington plane: our simulation shows a positive correlation between instantaneous f_{Edd} and spin at $z \sim 1$, but by $z = 0$ many low-f_{Edd} objects have low spins and low luminosities, and hence are effectively invisible to current techniques. The surviving observational sample preferentially inhabits the high-f_{Edd}, high-spin quadrant, reinforcing the apparent correlation.

Residual tensions remain at the extreme high-spin end ($a \gtrsim 0.9$), where a subset of observed AGN show spins close to the Thorne limit. Our model under-produces such objects, even after luminosity selection. This shortfall may reflect overly efficient jet-driven spin-down, or the absence of super-Eddington episodes that could prolong coherent growth. Given current measurement uncertainties and small sample sizes, in both observations and the simulation, these discrepancies are not yet decisive, but they highlight model ingredients–self-consistent spin coupling to accretion state, magnetic flux evolution–that warrant refinement.

5.4 Implications for Gravitational Wave Probes

Our merger catalogue shows that progenitor spins are typically well aligned with each other and with the binary orbital angular momentum at coalescence, producing remnant spins clustered at moderate values. For LISA, this implies a source population dominated by low-precession waveforms (small effective precession parameter) and positive effective spin χ_{eff} at high redshift–where coherent accretion has already spun up the components–transitioning toward lower χ_{eff} at later times as hot accretion and jet torques reduce spin magnitudes. Such a redshift evolution is directly testable through LISA's measurement of inspiral phasing and ringdown frequencies: aligned, moderate-spin remnants yield relatively narrow distributions of quasi-normal mode frequencies, simplifying stacking analyses but limiting sensitivity to extreme-spin tests of strong gravity. Strong alignment also suppresses gravitational-wave recoil kicks, favouring remnant retention in galactic nuclei and enabling subsequent hierarchical mergers; this enhances the high-z event rate relative to scenarios with widespread spin misalignment, potentially boosting the number of merger systems accessible to LISA.

Our results show that the remnant spins do not randomise completely: spin retains a measurable imprint of how long a black hole spent in the thin-disk (high-f_{Edd}) regime versus hot/chaotic accretion. In practice, it correlates with both the super critical mass and lifetime fractions at higher redshifts, indicating that coherent growth leaves a lasting "memory" in the spin even after subsequent mergers. This has direct implications for LISA: joint constraints on χ_{eff} during the inspiral and remnant spin from the ringdown, together with accurately measured masses and redshifts, can be inverted to infer the time-averaged accretion state of the progenitors. In other words, high-z spin measurements provide a new lever arm to reconstruct MBH growth histories–distinguishing systems built primarily through coherent, radiatively efficient accretion from those dominated by chaotic, low-f_{Edd} episodes.

5.5 Model Limitations

Several layers of idealisation and numerical constraint qualify the interpretation of our spin-evolution results. First, the subgrid accretion prescriptions treat radiatively efficient (thin-disk) and radiatively inefficient (hot/MAD) flows with simplified, parametrised spin-up/down laws. The hydrodynamic spin-up and electromagnetic (jet) spin-down terms–while spin dependent in our implementation–are still calibrated to a limited set of GRMHD simulations and ignore potential additional dependencies on magnetic flux accumulation history, disc thickness transitions, or episodic super-Eddington phases. If real systems sustain higher magnetic fluxes or experience prolonged slim-disc episodes, our model may under-produce very high spins; conversely, if jet torques are stronger than assumed, the low-spin tail would be further enhanced. The gas adhesion prescription also assumes a lossless, indefinitely coherent reservoir during off-centre phases; in practice, stripping, turbulence, or continued star formation could deplete or reorient this reservoir, shifting the adhesion model closer to the chaotic limit and lowering early spins.

Second, although radiative and jet efficiencies vary with spin, the model neglects other possible feedback couplings–for example, spin-modulated clearing of the central gas supply, that could accelerate late-time spin decline. Additionally, black hole recoil following mergers is not applied in the post-processing; while our strong spin alignment suppresses large kicks, rare misaligned events could displace remnants and transiently alter their accretion states.

The analysis also inherits numerical uncertainties from the underlying simulation. The central gas angular momentum vector is measured within a fixed 1.4 kpc aperture; unresolved sub-kpc structure or resolution-dependent turbulence may introduce orientation errors that propagate into the alignment torque integration, but we do not directly quantify this dispersion. The 0.1 Gyr output cadence is mitigated by RK4 sub-stepping, yet any rapid accretion rate transitions occurring on shorter timescales (e.g. brief chaotic episodes) are temporally smeared. Our merger identification criterion (coalescence at the softening scale) collapses the unmodelled hardening and gravitational-wave inspiral phases; if realistic delays are long, some high-redshift mergers would instead occur later, modestly shifting the redshift distribution of remnant spins.

Finally, the seed-mass experiment is not self-consistent with the hydrodynamic run: we rescale the initial M_{BH} but leave the simulated accretion rates (and hence feedback histories) unchanged. Prior work on Romulus shows that accretion is largely gas-supply limited and correlates more strongly with the host's star formation than with M_{BH} itself (e.g. Ricarte et al. 2019), so imposing a lower seed mass effectively acts like an early "boost factor" (akin to the $\sim 100\times$ boosts used in other simulations; e.g. Springel et al. 2005) rather than fundamentally changing the fuelling physics. Because subgrid accretion is unresolved, this is not inherently unrealistic, but it does mean the associated feedback modulation is missing: a smaller black hole accreting the same $\dot{M}$ would couple less energy/momentum to the interstellar medium than in the original run. This caveat emphasises that the exercise clarifies causal pathways–how sustained high f_{Edd} phases imprint on spin–rather than providing a fully self-consistent alternative history. Crucially, these limitations affect the quantitative extremes (very high spins, intrinsic scatter) more than the qualitative trend of early coherent spin-up followed by late-time decline.

Our spin evolution model adopts a 0.1 Gyr averaging window for black hole accretion rates, considerably longer than the typical short-term variability seen in simulated AGN, which often occurs on Myr timescales or shorter (e.g. Novak et al. 2011). Because black hole spin is sensitive to the instantaneous direction and magnitude of angular momentum supplied by the accretion flow, smoothing over periods longer than relevant disc timescales may reduce the frequency and coherence of spin-up episodes, biasing the model toward lower spin magnitudes. Indeed, previous spin evolution models that used shorter averaging windows (or instantaneous rates) typically found more rapid spin-up, particularly for seed and intermediate-mass black holes (e.g. Dubois et al. 2014; Bustamante & Springel 2019; Beckmann et al. 2025).

Physically, the most relevant timescale is the viscous timescale, which determines how rapidly gas flows inward and therefore how quickly an accretion disc's angular momentum profile and orientation can change. For typical AGN thin-discs, the viscous timescale at the ISCO is on the order of thousands of years or shorter, increasing outward to millions of years at parsec scales (e.g. King 2008). Thus, averaging accretion rates on 0.1 Gyr intervals significantly exceeds these characteristic disc timescales, potentially washing out intermittent high-$\dot{M}$, high-spin-up phases.

On the other hand, Romulus employs a comparatively large seed mass of $10^6\,M_\odot$ resulting in black holes forming later, compared to lower mass seeds, since higher density gas is required to collapse into a seed. This means the earliest phase of rapid spin-up in gas-rich environments is effectively unresolved. Thus, even if a shorter averaging window were employed, the earliest and potentially most rapid spin-up episodes might already be missed in this particular simulation, limiting achievable maximum spins regardless of accretion averaging. Nevertheless, a more rigorous future treatment employing finer time resolution–ideally comparable to viscous timescales–would help quantify how significantly averaging affects predicted spin distributions, particularly for high-mass black holes observed to have high spins (e.g. Reynolds 2014).

In summary, the dominant uncertainties in our analysis are systematic rather than statistical, and where they bias the results, we can estimate the qualitative effects. The three off-centre prescriptions act as a bracketing experiment, with the spread between chaotic and assumed-central providing an envelope. The 0.1 Gyr averaging of $\dot{M}$–much longer than typical AGN variability (Myr–sub-Myr; Novak et al. 2011)–together with Romulus's $10^6 M_\odot$ seeds likely makes our high-spin incidence a lower bound. Adopting the 1.4 kpc nuclear aperture provides a stable reference unlikely to erase the strong pre-merger alignment. The lower-seed experiment is not fully self-consistent but effectively mimics an early 'boost' to accretion (Springel et al. 2005); because Romulus accretion is largely supply-limited (Ricarte et al. 2019), our causal conclusions remain intact. Consequently, while the extremes (very high a) remain uncertain, the qualitative conclusions–higher typical spins at early times, decline toward $z = 0$, strong alignment, and accretion dominating over mergers–are robust and consistent with prior studies (e.g. Dubois et al. 2014; Bustamante & Springel 2019; Beckmann et al. 2025; Sala et al. 2024), justifying the approach for the purposes of this work.

5.6 Future Directions

Several extensions would strengthen and refine the conclusions presented here. A natural next step is to perform fully self-consistent spin-evolution simulations, embedding the accretion-mode classification, spin magnitude/orientation updates, and merger remnant prescriptions directly into the hydrodynamic time integration. Coupling spin to accretion and feedback on the fly would allow black hole spin to influence gas supply in real time, producing feedback loops (e.g. spin-dependent jet clearing) that the current post-processing cannot capture. This implementation would also enable dynamic adjustment of alignment timescales based on instantaneous disc properties rather than fixed analytic forms.

Improving the treatment of off-centre accretion is another priority. The adhesion model's assumption of an indefinitely coherent reservoir can be relaxed by tracking a finite mass and angular momentum budget that depletes through accretion, mixes via turbulence, and is replenished only upon re-centring. Alternatively, local gas properties around wandering black holes could be sampled directly at higher resolution (e.g. via zoom-in re-simulations) to empirically calibrate stochastic angular momentum delivery and the frequency of thin-disk episodes during displacement. Incorporating a physically motivated diffusion term for spin orientation while off-centre would further bracket realistic coherence loss.

Finally, an extended parameter survey exploring seed mass distributions, alignment timescale normalisations, jet torque strengths, and alternative merger-delay prescriptions would quantify degeneracies and identify combinations that remain consistent with existing spin and AGN demographics. Coupling these runs with synthetic observational pipelines (producing mock X-ray reflection spectra and LISA gravitational-wave catalogues) would allow direct modelling of selection functions, tightening constraints on the underlying population. Together, these developments would move the field toward predictive models of MBH spin evolution across cosmic time.

Chapter 6

Conclusion

This thesis has developed and applied a post-processing framework to evolve the spins of MBHs within the Romulus25 cosmological simulation. Building on two complementary formalisms–accretion-driven spin evolution (Ricarte et al. 2023) and merger spin prescriptions with orientation tracking (Dubois et al. 2014)–the model integrates spin magnitude and direction through thin-disk and hot (MAD) regimes, incorporates spin-dependent radiative and jet efficiencies, and treats mergers using numerical-relativity-based fits. A central innovation is the exploration of three alternative off-centre accretion prescriptions (assumed central, chaotic, and gas adhesion) combined with two seed masses, allowing us to bracket uncertainties associated with black hole wandering and early growth.

The resulting spin histories reveal a coherent narrative: early, gas-rich epochs support prolonged thin-disk accretion that rapidly spins up low-mass seeds; declining gas supply and increasing hot-mode duty cycles thereafter drive a population-wide spin decrease toward the present day. Off-centre accretion phases also materially shape the spin distribution. Our three prescriptions show that how a displaced black hole accretes- -retaining a coherent reservoir, accreting chaotically, or being treated as central- -shifts the population toward or away from high spins. Modelling this off-nuclear growth is therefore essential, as it imprints directly on spin demographics and the electromagnetic and gravitational wave signals we aim to interpret. Seed mass experiments show that elevated early Eddington ratios amplify initial spin-up, with memory largely lost once black holes grow beyond

$\sim 10^{7-8}M_{\odot}$. Mergers, occurring predominantly between aligned progenitors, generate moderate remnant spins and play a secondary role in shaping the ensemble distribution.

When observational selection is imposed (luminosity cuts approximating current X-ray reflection measured spin samples), the model produces a bias toward high spins at high and moderate redshifts. These selection effects reconcile apparent tensions between observed high spins and the low median spins of the full simulated population, and they generate testable predictions for the spin-mass-Eddington relations accessible to future facilities. Strong spin alignment further implies low gravitational-wave recoil velocities, favouring remnant retention and enhancing the hierarchical merger channel relevant for LISA.

The analysis remains subject to limitations: idealised off-centre gas retention, simplified magnetic-flux evolution, instantaneous merger coalescence, and the artificial seed-mass rescaling. Nonetheless, these approximations primarily affect the extreme high-spin tail and do not undermine the central conclusion that early coherent accretion imprints high spins which are subsequently eroded by hot-mode,jet-driven spin-down. Future work incorporating self-consistent on-the-fly spin evolution, finite adhesion reservoirs, and improved magnetohydrodynamic jet torque prescriptions will refine these results and strengthen links to gravitational-wave and electromagnetic observables.

In summary, this thesis provides a physically grounded, simulation-anchored view of MBH spin evolution across cosmic time, quantifying the roles of coherent accretion, wandering, and mergers, and establishing a framework for interpreting current and forthcoming spin measurements in the broader context of black hole-galaxy co-evolution.

Bibliography

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, Physical Review Letters, 116, 061102

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, The Astrophysical Journal Letters, 882, L24

Anglés-Alcázar, D., Davé, R., Faucher-Giguère, C.-A., Özel, F., & Hopkins, P. F. 2017, Monthly Notices of the Royal Astronomical Society, 464, 2840

Bambi, C., Brenneman, L. W., Dauser, T., et al. 2021, Space Science Reviews, 217

Barausse, E., Dvorkin, I., Tremmel, M., Volonteri, M., & Bonetti, M. 2020, The Astrophysical Journal, 904, 16, publisher: IOP ADS Bibcode: 2020ApJ...904...16B

Barausse, E. & Rezzolla, L. 2009, The Astrophysical Journal, 704, L40

Bardeen, J. M. 1970, The Astrophysical Journal, 162, 71, publisher: IOP ADS Bibcode: 1970ApJ...162...71B

Beckmann, R. S., Dubois, Y., Volonteri, M., et al. 2025, Monthly Notices of the Royal Astronomical Society, 536, 1838

Beckmann, R. S., Smethurst, R. J., Simmons, B. D., et al. 2024, Monthly Notices of the Royal Astronomical Society, 527, 10867

Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, Nature, 287, 307, ADS Bibcode: 1980Natur.287..307B

Berti, E. & Volonteri, M. 2008, The Astrophysical Journal, 684, 822, publisher: IOP Publishing

Blandford, R. D. & Znajek, R. L. 1977, Monthly Notices of the Royal Astronomical Society, 179, 433, publisher: OUP ADS Bibcode: 1977MNRAS.179..433B

Blecha, L., Cox, T. J., Loeb, A., & Hernquist, L. 2011, Monthly Notices of the Royal Astronomical Society, 412, 2154, publisher: OUP ADS Bibcode: 2011MNRAS.412.2154B

Bogdanovic, T., Reynolds, C. S., & Miller, M. C. 2007, The Astrophysical Journal, 661, L147

Booth, C. M. & Schaye, J. 2009, Monthly Notices of the Royal Astronomical Society, 398, 53

Bromm, V. & Loeb, A. 2003, The Astrophysical Journal, 596, 34, publisher: IOP Publishing

Bryan, G. L. & Norman, M. L. 1998, The Astrophysical Journal, 495, 80

Buonanno, A., Kidder, L. E., & Lehner, L. 2008, Physical Review D, 77

Bustamante, S. & Springel, V. 2019, Monthly Notices of the Royal Astronomical Society, 490, 4133

Campanelli, M., Lousto, C., Zlochower, Y., & Merritt, D. 2007, The Astrophysical Journal, 659, L5–L8

Chael, A., Lupsasca, A., Wong, G. N., & Quataert, E. 2023, The Astrophysical Journal, 958, 65

Chapon, D., Mayer, L., & Teyssier, R. 2013, Monthly Notices of the Royal Astronomical Society, 429, 3114, publisher: OUP ADS Bibcode: 2013MNRAS.429.3114C

Chen, N., Ni, Y., Tremmel, M., et al. 2022, Monthly Notices of the Royal Astronomical Society, 510, 531

Colpi, M., Holley-Bockelmann, K., Bogdanovic, T., et al. 2019, Astro2020 science white paper: The gravitational wave view of massive black holes

Costa, T., Sijacki, D., Trenti, M., & Haehnelt, M. G. 2014, Monthly Notices of the Royal Astronomical Society, 439, 2146, publisher: OUP ADS Bibcode: 2014MNRAS.439.2146C

Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, Monthly Notices of the Royal Astronomical Society, 486, 2827–2849

Devecchi, B. & Volonteri, M. 2009, The Astrophysical Journal, 694, 302, publisher: IOP ADS Bibcode: 2009ApJ...694..302D

Dotti, M., Colpi, M., Pallini, S., Perego, A., & Volonteri, M. 2012, The Astrophysical Journal, 762, 68, publisher: The American Astronomical Society

Dubois, Y., Beckmann, R., Bournaud, F., et al. 2021, Astronomy & Astrophysics, 651, A109, publisher: EDP Sciences

Dubois, Y., Devriendt, J., Slyz, A., & Teyssier, R. 2012, Monthly Notices of the Royal Astronomical Society, 420, 2662

Dubois, Y., Volonteri, M., & Silk, J. 2014, Monthly Notices of the Royal Astronomical Society, 440, 1590, publisher: OUP ADS Bibcode: 2014MNRAS.440.1590D

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, The Astrophysical Journal, 875, L1, publisher: IOP ADS Bibcode: 2019ApJ...875L...1E

Fabian, A. C. 2012, Annual Review of Astronomy and Astrophysics, 50, 455, ADS Bibcode: 2012ARA&A..50..455F

Fanidakis, N., Baugh, C. M., Benson, A. J., et al. 2011, Monthly Notices of the Royal Astronomical Society, 419, 2797–2820

Ferrarese, L. & Merritt, D. 2000, The Astrophysical Journal, 539, L9

Feruglio, C., Maiolino, R., Piconcelli, E., et al. 2010, Astronomy and Astrophysics, 518, L155, ADS Bibcode: 2010A&A...518L.155F

Fiore, F., Feruglio, C., Shankar, F., et al. 2017, Astronomy and Astrophysics, 601, A143, ADS Bibcode: 2017A&A...601A.143F

Gaspari, M., Tombesi, F., & Cappi, M. 2020, Nature Astronomy, 4, 10

Gebhardt, K., Bender, R., Bower, G., et al. 2000, The Astrophysical Journal, 539, L13, publisher: IOP ADS Bibcode: 2000ApJ...539L..13G

Georgakakis, A., Aird, J., Schulze, A., et al. 2017, Monthly Notices of the Royal Astronomical Society, 471, 1976–2001

Gerosa, D. & Berti, E. 2017, Physical Review D, 95

Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, The Astrophysical Journal, 689, 1044, publisher: IOP ADS Bibcode: 2008ApJ...689.1044G

Gillessen, S., Eisenhauer, F., Trippe, S., et al. 2009, The Astrophysical Journal, 692, 1075, publisher: IOP ADS Bibcode: 2009ApJ...692.1075G

Greene, J. E., Strader, J., & Ho, L. C. 2020, Annual Review of Astronomy and Astrophysics, 58, 257

Group, L. C. W. W., Afshordi, N., Akçay, S., et al. 2023, Waveform Modelling for the Laser Interferometer Space Antenna

Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, The Astrophysical Journal, 742, 76

Ho, L. C. 2008, Annual Review of Astronomy and Astrophysics, 46, 475

Ho, L. C. 2009, The Astrophysical Journal, 699, 626, publisher: The American Astronomical Society

Hopkins, P. F., Hernquist, L., Cox, T. J., et al. 2006, The Astrophysical Journal Supplement Series, 163, 1, publisher: IOP ADS Bibcode: 2006ApJS..163....1H

Inayoshi, K., Visbal, E., & Haiman, Z. 2020, Annual Review of Astronomy and Astrophysics, 58, 27

Katz, M. L., Kelley, L. Z., Dosopoulou, F., et al. 2020, Monthly Notices of the Royal Astronomical Society, 491, 2301, publisher: OUP ADS Bibcode: 2020MNRAS.491.2301K

Kelley, L. Z., Blecha, L., & Hernquist, L. 2017, Monthly Notices of the Royal Astronomical Society, 464, 3131

King, A. 2008, New Astronomy Reviews, 52, 253

King, A. R., Lubow, S. H., Ogilvie, G. I., & Pringle, J. E. 2005, Monthly Notices of the Royal Astronomical Society, 363, 49

King, A. R. & Pringle, J. E. 2006, Monthly Notices of the Royal Astronomical Society, 373, L90

King, A. R., Pringle, J. E., & Hofmann, J. A. 2008, Monthly Notices of the Royal Astronomical Society, 385, 1621

King, A. R., Pringle, J. E., & Livio, M. 2007, Monthly Notices of the Royal Astronomical Society, 376, 1740–1746

Klein, A., Barausse, E., Sesana, A., et al. 2016, Physical Review D, 93, 024003, publisher: American Physical Society

Knollmann, S. R. & Knebe, A. 2009, The Astrophysical Journal Supplement Series, 182, 608–624

Komissarov, S. 2001, Monthly Notices of the Royal Astronomical Society, 326, L41

Kormendy, J. & Ho, L. C. 2013, Annual Review of Astronomy and Astrophysics, 51, 511, ADS Bibcode: 2013ARA&A..51..511K

Kormendy, J. & Richstone, D. 1995, Annual Review of Astronomy and Astrophysics, 33, 581, ADS Bibcode: 1995ARA&A..33..581K

Kroupa, P. 2001, Monthly Notices of the Royal Astronomical Society, 322, 231

Liska, M., Tchekhovskoy, A., Ingram, A., & Klis, M. v. d. 2019, Monthly Notices of the Royal Astronomical Society, 487, 550

Loeb, A. & Rasio, F. A. 1994, The Astrophysical Journal, 432, 52, publisher: IOP ADS Bibcode: 1994ApJ...432...52L

Lousto, C. O. & Zlochower, Y. 2011, Physical Review Letters, 106

Lynden-Bell, D. 1969, Nature, 223, 690, ADS Bibcode: 1969Natur.223..690L

Madau, P. & Rees, M. J. 2001, The Astrophysical Journal, 551, L27

Magorrian, J., Tremaine, S., Richstone, D., et al. 1998, The Astronomical Journal, 115, 2285, publisher: IOP ADS Bibcode: 1998AJ....115.2285M

Mallick, L., Fabian, A. C., García, J. A., et al. 2022, Monthly Notices of the Royal Astronomical Society, 513, 4361–4379

Marconi, A., Risaliti, G., Gilli, R., et al. 2004, Monthly Notices of the Royal Astronomical Society, 351, 169

Mashchenko, S., Couchman, H. M. P., & Wadsley, J. 2006, Nature, 442, 539–542

McConnell, N. J. & Ma, C.-P. 2013, The Astrophysical Journal, 764, 184

McKinney, J. C., Tchekhovskoy, A., & Blandford, R. D. 2012, Monthly Notices of the Royal Astronomical Society, 423, 3083, publisher: OUP ADS Bibcode: 2012MNRAS.423.3083M

McNamara, B. R. & Nulsen, P. E. J. 2007, Annual Review of Astronomy and Astrophysics, 45, 117, ADS Bibcode: 2007ARA&A..45..117M

Menon, H., Wesolowski, L., Zheng, G., et al. 2015, Computational Astrophysics and Cosmology, 2, 1

Miller, J. 2007, Annual Review of Astronomy and Astrophysics, 45, 441–479

Narayan, R., Chael, A., Chatterjee, K., Ricarte, A., & Curd, B. 2022, Monthly Notices of the Royal Astronomical Society, 511, 3795–3813

Narayan, R., Igumenshchev, I. V., & Abramowicz, M. A. 2003, Publications of the Astronomical Society of Japan, 55, L69, publisher: OUP ADS Bibcode: 2003PASJ...55L..69N

Narayan, R. & Yi, I. 1994, The Astrophysical Journal, 428, L13, publisher: IOP ADS Bibcode: 1994ApJ...428L..13N

Novak, G. S., Ostriker, J. P., & Ciotti, L. 2011, The Astrophysical Journal, 737, 26

Oppenheimer, J. R. & Snyder, H. 1939, Physical Review, 56, 455, publisher: American Physical Society

Peirani, S., Suto, Y., Beckmann, R. S., et al. 2024, Astronomy amp; Astrophysics, 686, A233

Penna, R. F., Narayan, R., & Sądowski, A. 2013, Monthly Notices of the Royal Astronomical Society, 436, 3741

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, Astronomy & Astrophysics, 571, A20

Rees, M. J. 1984, Annual Review of Astronomy and Astrophysics, 22, 471, ADS Bibcode: 1984ARA&A..22..471R

Reynolds, C. S. 2014, Space Science Reviews, 183, 277

Reynolds, C. S. 2019, Nature Astronomy, 3, 41

Reynolds, C. S. & Nowak, M. A. 2003, Physics Reports, 377, 389

Rezzolla, L., Barausse, E., Dorband, E. N., et al. 2008, Physical Review D, 78, 044002

Ricarte, A., Narayan, R., & Curd, B. 2023, The Astrophysical Journal, 954, L22, publisher: IOP ADS Bibcode: 2023ApJ...954L..22R

Ricarte, A. & Natarajan, P. 2018, Monthly Notices of the Royal Astronomical Society, 481, 3278

Ricarte, A., Natarajan, P., Narayan, R., & Palumbo, D. C. M. 2025, The Astrophysical Journal, 980, 136

Ricarte, A., Tremmel, M., Natarajan, P., & Quinn, T. 2019, Monthly Notices of the Royal Astronomical Society, 489, 802–819

Rosas-Guevara, Y. M., Bower, R. G., Schaye, J., et al. 2015, Monthly Notices of the Royal Astronomical Society, 454, 1038

Sala, L., Valentini, M., Biffi, V., & Dolag, K. 2024, Astronomy Astrophysics, 685

Salpeter, E. E. 1964, The Astrophysical Journal, 140, 796, publisher: IOP ADS Bibcode: 1964ApJ...140..796S

Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, Monthly Notices of the Royal Astronomical Society, 446, 521

Schnittman, J. D. 2011, Classical and Quantum Gravity, 28, NTRS Author Affiliations: NASA Goddard Space Flight Center NTRS Report/Patent Number: GSFC.JA.6417.2012 NTRS Document ID: 20120013628 NTRS Research Center: Goddard Space Flight Center (GSFC)

Sesana, A., Barausse, E., Dotti, M., & Rossi, E. M. 2014, The Astrophysical Journal, 794, 104

Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2005, The Astrophysical Journal, 623, 23, publisher: IOP Publishing

Shakura, N. I. & Sunyaev, R. A. 1973, Astronomy and Astrophysics, 24, 337, ADS Bibcode: 1973A&A....24..337S

Sijacki, D., Springel, V., Matteo, T. D., & Hernquist, L. 2007, Monthly Notices of the Royal Astronomical Society, 380, 877

Silk, J. & Rees, M. J. 1998, Quasars and galaxy formation, ISSN: 0004-6361 Volume: 331 ADS Bibcode: 1998A&A...331L...1S

Soares, G. & Nemmen, R. 2020, Monthly Notices of the Royal Astronomical Society, 495, 981

Springel, V., Di Matteo, T., & Hernquist, L. 2005, Monthly Notices of the Royal Astronomical Society, 361, 776, publisher: OUP ADS Bibcode: 2005MNRAS.361..776S

Tchekhovskoy, A., Narayan, R., & McKinney, J. C. 2011, Monthly Notices of the Royal Astronomical Society: Letters, 418, L79

Thorne, K. S. 1974, The Astrophysical Journal, 191, 507

Tombesi, F., Meléndez, M., Veilleux, S., et al. 2015, Nature, 519, 436, ADS Bibcode: 2015Natur.519..436T

Tremmel, M., Governato, F., Volonteri, M., & Quinn, T. R. 2015, Monthly Notices of the Royal Astronomical Society, 451, 1868–1874

Tremmel, M., Karcher, M., Governato, F., et al. 2017, Monthly Notices of the Royal Astronomical Society, 470, 1121

Volonteri, M. 2010, Astronomy and Astrophysics Review, 18, 279, publisher: Springer ADS Bibcode: 2010A&ARv..18..279V

Volonteri, M., Haardt, F., & Madau, P. 2003, The Astrophysical Journal, 582, 559

Volonteri, M., Madau, P., Quataert, E., & Rees, M. J. 2005, The Astrophysical Journal, 620, 69, publisher: IOP Publishing

Volonteri, M., Sikora, M., Lasota, J.-P., & Merloni, A. 2013, The Astrophysical Journal, 775, 94

Weinberger, R., Springel, V., Hernquist, L., et al. 2017, Monthly Notices of the Royal Astronomical Society, 465, 3291, publisher: OUP ADS Bibcode: 2017MNRAS.465.3291W

Yuan, F., Gan, Z., Narayan, R., et al. 2015, The Astrophysical Journal, 804, 101, publisher: The American Astronomical Society

Yuan, F. & Narayan, R. 2014, Annual Review of Astronomy and Astrophysics, 52, 529, ADS Bibcode: 2014ARA&A..52..529Y