

---

THE IMPACT OF AGN OUTFLOWS  
ON MULTIPHASE GAS: LINKING  
OBSERVATIONS AND SIMULATIONS

---



Samuel Ruthven Ward  
Munich, 2024



---

**The Impact of AGN Outflows on Multiphase Gas:  
Linking Observations and Simulations**

**Samuel Ruthven Ward**

---

**Dissertation**

at the

**Faculty of Physics**

of the

**Ludwig Maximilians University  
Munich**

submitted by

**Samuel Ruthven Ward**

born in Sidcup, United Kingdom

Munich, 24<sup>th</sup> July 2024



---

**Die Auswirkungen von AGN Ausflüssen auf  
Mehrphasengas: Verknüpfung von Beobachtungen  
und Simulationen**

**Samuel Ruthven Ward**

---

**Dissertation**

an der

**Fakultät für Physik**

der

**Ludwig-Maximilians-Universität  
München**

vorgelegt von

**Samuel Ruthven Ward**

geboren in Sidcup, Vereinigtes Königreich

München, den 24. Juli 2024



Erstgutachter: Prof. Dr. Volker Springel  
Zweitgutachter: Prof. Dr. Eric Emsellem  
Tag der mündlichen Prüfung: 01. Oktober 2024





# Declaration

The work described in this thesis was undertaken by the author while enrolled as a research student at Ludwig-Maximilians-Universität München, under the supervision of **Prof. Dr. Volker Springel**.

The research presented here was carried out between September 2020 and July 2024 at the European Southern Observatory, Garching bei München, Germany and Newcastle University, Newcastle upon Tyne, United Kingdom, under the joint supervision of **Dr. Chris Harrison**, **Dr. Tiago Costa** and **Dr. Vincenzo Mainieri**.

The author acknowledges funding from the Excellence Cluster ORIGINS. The ORIGINS cluster is funded by the Deutsche Forschungsgemeinschaft (DFG; German Research Foundation) under Germany's Excellence Strategy: EXC-2094-390783311.



*And all the science I don't understand,  
it's just my job five days a week*

*–Rocket Man, Elton John*



# Contents

<b>List of Figures</b>	<b>xiii</b>
<b>List of Tables</b>	<b>xiv</b>
<b>List of Acronyms</b>	<b>xvii</b>
<b>Zusammenfassung</b>	<b>xix</b>
<b>Abstract</b>	<b>xxi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Galaxy Evolution . . . . .	3
1.1.1 Structure Formation in $\Lambda$ CDM . . . . .	3
1.1.2 Galaxy Classification . . . . .	5
1.1.3 Quenching . . . . .	6
1.1.4 The Need for Baryonic Feedback . . . . .	6
1.2 Active Galactic Nuclei . . . . .	8
1.2.1 Supermassive Black Holes . . . . .	8
1.2.2 Accretion . . . . .	8
1.2.3 Energy Coupling . . . . .	9
1.2.4 Accretion Discs . . . . .	10
1.2.5 Launching Mechanisms . . . . .	12
1.2.6 Observational Signatures . . . . .	13
1.2.7 Theorised Feedback Impact . . . . .	16
1.3 Thesis Overview . . . . .	17
<b>2 Computational Methodologies</b>	<b>19</b>
2.1 A Brief History of Numerical Simulations . . . . .	21
2.2 Hydrodynamics . . . . .	22
2.2.1 Equations of Fluid Dynamics . . . . .	22
2.2.2 Hydrodynamic Codes . . . . .	23
2.3 Cosmological Simulations . . . . .	25
2.3.1 Subgrid Models . . . . .	26
2.3.2 Modelling Black Holes . . . . .	27
2.3.3 Interpretation Challenges . . . . .	29
2.4 A Resolved Model for AGN Winds . . . . .	31
2.4.1 Analytic Background . . . . .	31
2.4.2 Implementation in AREPO . . . . .	33

<b>3</b>	<b>AGN Feedback in Cosmological Simulations</b>	<b>35</b>
3.1	Introduction . . . . .	37
3.2	Methods . . . . .	40
3.2.1	Simulations . . . . .	40
3.2.2	Observational Samples . . . . .	45
3.2.3	Key Quantities . . . . .	46
3.2.4	AGN Selection Criteria . . . . .	47
3.2.5	Quenching & Gas-Depletion Definitions . . . . .	48
3.2.6	Correlation Coefficients . . . . .	49
3.3	Results . . . . .	50
3.3.1	Correlations with $L_{\text{bol}}$ . . . . .	51
3.3.2	AGN and non-AGN in the $f_{H_2}-M_*$ plane . . . . .	53
3.3.3	$f_{H_2}$ Distributions in AGN . . . . .	55
3.3.4	Gas-Depleted Fractions of AGN . . . . .	56
3.4	Discussion . . . . .	58
3.4.1	Cosmological models predict AGN preferentially live in gas rich, star-forming galaxies . . . . .	58
3.4.2	Different feedback modes lead to different simulation predictions at $z = 0$ . . . . .	60
3.4.3	Limitations and Outlook . . . . .	62
3.5	Conclusions . . . . .	64
3.6	Chapter Appendix: sSFR Plots . . . . .	65
<b>4</b>	<b>Multiphase Outflows from AGN Winds</b>	<b>69</b>
4.1	Introduction . . . . .	71
4.2	Numerical Simulations . . . . .	73
4.2.1	Hydrodynamic Code . . . . .	73
4.2.2	Galaxy Setup . . . . .	74
4.2.3	Setting the ISM Structure . . . . .	74
4.2.4	AGN Wind Model . . . . .	76
4.2.5	Simulation Suite . . . . .	76
4.2.6	Calculating Outflow Properties . . . . .	78
4.3	Results I: Multiphase Outflows from Small-Scale Winds . . . . .	79
4.3.1	Multiscale Outflow Structure . . . . .	79
4.3.2	Multiphase Gas Energetics . . . . .	82
4.3.3	Parameter Variation . . . . .	86
4.4	Results II: Implications for Observed Outflow Properties . . . . .	89
4.4.1	Challenges for Measuring Outflow Properties . . . . .	89
4.4.2	Challenges for Interpreting Momentum Fluxes and Kinetic Luminosities . . . . .	93
4.4.3	Outflow Scaling Relations with AGN Luminosity . . . . .	97
4.5	Conclusions & Outlook . . . . .	101
4.6	Chapter Appendix: Numerical Convergence . . . . .	103
<b>5</b>	<b>X-ray Emission from Wind-ISM Interactions</b>	<b>107</b>
5.1	Introduction . . . . .	109
5.2	Methods . . . . .	110
5.2.1	The ACDC Simulations . . . . .	110
5.2.2	Bremsstrahlung Emission . . . . .	111
5.2.3	Synthetic Observations . . . . .	112

5.3	Results . . . . .	112
5.3.1	Wind-Cloud Interactions Produce X-ray Emission . . . . .	112
5.3.2	Wind Tracer Phases . . . . .	114
5.3.3	Total X-ray Contribution by Wind Phase . . . . .	115
5.3.4	Radial Evolution . . . . .	117
5.3.5	Time Evolution . . . . .	118
5.3.6	Scaling Relations between AGN and X-ray Luminosities . . . . .	119
5.4	Discussion: Observations with Chandra . . . . .	121
5.5	Conclusions & Outlook . . . . .	123
<b>6</b>	<b>Summary and Outlook</b>	<b>125</b>
6.1	Conclusions . . . . .	127
6.1.1	AGN preferentially reside in gas-rich galaxies . . . . .	127
6.1.2	ISM structure is key to understanding multiphase outflows . . . . .	128
6.1.3	Wind-cloud mixing produces luminous X-ray emission . . . . .	129
6.2	Outlook . . . . .	129
6.2.1	Linking Simulations & Observations to Constrain Feedback . . . . .	129
6.2.2	Understanding Cold Gas in Outflows . . . . .	130
6.2.3	X-ray and Radio Emission from Wind-ISM Interactions . . . . .	133
6.2.4	The Next Decade in Galaxy Simulations . . . . .	135
6.3	Final Remarks . . . . .	137
	<b>Bibliography</b>	<b>139</b>
	<b>Publication List</b>	<b>155</b>
	<b>Acknowledgements</b>	<b>157</b>





# List of Figures

1.1	Stephan’s Quintet, as seen by <i>JWST</i> . . . . .	3
1.2	The main sequence of star formation. . . . .	5
1.3	Efficiency of star formation as a function of halo mass, demonstrating the need for baryonic feedback. . . . .	7
1.4	A schematic overview of the different processes involved in AGN feedback . . . . .	11
1.5	Scaling relations between AGN luminosity and outflow properties from <a href="#">Fiore et al. (2017)</a> . . . . .	14
2.1	Holmberg’s early simulation of a galaxy merger, using lamps. . . . .	21
2.2	The increase in the number of particles used in cosmological simulations in the last 60 years . . . . .	22
2.3	A schematic showing various numerical methods to calculate fluid flows. . . . .	23
2.4	The EAGLE cosmological simulation. . . . .	25
2.5	The effect of subgrid models for SMBH feedback in cosmological simulations. . . . .	29
2.6	The structure of an AGN wind-driven outflow. . . . .	32
2.7	The implementation of a resolved AGN wind in the AREPO code. . . . .	33
3.1	The gas fraction main sequence, demonstrating that galaxies hosting AGN are preferentially gas-rich and star-forming. . . . .	39
3.2	AGN luminosity as a function of black hole mass in the three simulations. . . . .	43
3.3	Calculation of the correlation coefficients. . . . .	49
3.4	Correlation coefficients between $L_{\text{bol}}$ and host galaxy properties. . . . .	51
3.5	The gas fraction main sequence for the three simulations. . . . .	53
3.6	The distribution of gas fractions at $z = 0$ and $z = 2$ . . . . .	56
3.7	The gas-depleted fraction in the simulations for different AGN selection criteria. . . . .	57
3.8	The gas fraction as a function of $L_{\text{bol}}$ for the simulations and observational samples. . . . .	60
3.10	The distribution of sSFR at $z = 0$ and $z = 2$ . . . . .	67
3.11	The quenched fraction in the simulations for different AGN selection criteria. . . . .	67
3.12	Specific star formation rate as a function of $L_{\text{bol}}$ for the simulations and observational samples. . . . .	68
4.1	The initial conditions for the ACDC simulations, showing the range of clumping factors. . . . .	74
4.2	A summary figure, showing the results of our fiducial simulation at a range of scales. . . . .	80

4.3	An AGN wind interacting with a smooth disc. . . . .	81
4.4	Phase diagrams showing the gas temperature against the radial velocity. . . . .	83
4.5	The cumulative momentum rate and kinetic energy coupling efficiency as a function of temperature. . . . .	84
4.6	The distribution of radial velocities for simulations of different clump sizes. . . . .	86
4.7	The variation in mass outflow rate due to different setup parameters. . . . .	88
4.8	The radial evolution of the outflow. . . . .	92
4.9	The time evolution of the outflow. . . . .	95
4.10	Momentum flux against velocity as a diagnostic for energy- or momentum-driving. . . . .	96
4.11	Scaling relations between the AGN luminosity, mass outflow rate and kinetic energy coupling efficiency, compared to observational samples. . . . .	98
4.12	The distribution of cell diameters as a function of mass resolution. . . . .	104
4.13	The numerical convergence of global outflow properties. . . . .	105
4.14	The results from a high-resolution simulation, showing how the size of the resulting cloudlets depends on numerical resolution. . . . .	106
5.1	An overview of the location of X-ray emission in the ACDC simulations. . . . .	113
5.2	The distribution of X-ray luminosity as a function of wind tracer density. . . . .	114
5.3	The cumulative X-ray luminosity as a function of outflow velocity, showing the relative contribution from different wind phases. . . . .	116
5.4	The radial evolution of the X-ray producing outflow. . . . .	117
5.5	The time evolution of the global X-ray luminosity. . . . .	118
5.6	Scaling relation between AGN luminosity and Bremsstrahlung emission, compared to nuclear emission from the AGN and contamination from star formation. . . . .	119
5.7	Synthetic observations using the <i>Chandra</i> X-ray observatory. . . . .	121
6.1	Spatially resolved kinematics and stellar populations from recent IFU observations. . . . .	131
6.2	A novel targeted refinement method for cold, outflowing gas. . . . .	132
6.3	EHT images of M87 and SgrA* . . . . .	137

# List of Tables

- 3.1 A summary of the subgrid implementations for AGN feedback in IllustrisTNG, EAGLE and SIMBA. . . . . 42
- 3.2 Sample sizes for our AGN selection at  $z = 0$  and  $z = 2$  . . . . . 46
- 4.1 The suite of simulations explored in this chapter. . . . . 77



# List of Acronyms

ACDC – AGN in Clumpy Discs  
ACIS – Advanced CCD Imaging Spectrometer  
ADAF – Advection-dominated accretion flow  
AGB – Asymptotic giant branch  
AGN – Active galactic nucleus / nuclei  
ALMA – Atacama Large Millimeter Array  
APEX – Atacama Pathfinder Experiment  
AMR – Adaptive mesh refinement  
BCG – Brightest cluster galaxy  
BH – Black hole  
BOLA - Boundary Layer for AGN  
CCD – Charge-coupled device  
CDM – Cold dark matter  
CGM – Circumgalactic medium  
CIE – Collisional excitation equilibrium  
CMB – Cosmic microwave background  
CO – Carbon monoxide  
CR – Cosmic ray  
CSA – Canadian Space Agency  
DESI – Dark Energy Spectroscopic Survey  
eEOS – Effective equation of state  
EHT – Event Horizon Telescope  
ELT – Extremely Large Telescope  
ESA – European Space Agency  
ESO – European Southern Observatory  
FR-I / II – Fanaroff-Riley Class I / II  
FWHM – Full width at half maximum  
GRMHD – General-relativistic-magneto-hydrodynamics  
ICM – Intracluster medium

IFU – Integral field unit  
IGM – Intergalactic medium  
IGrM – Intragroup medium  
IR – Infrared  
ISM – Interstellar medium  
JWST – James Webb Space Telescope  
LERG / HERG – Low / high excitation radio galaxy  
LLAGN – Low luminosity AGN  
LMU – Ludwig-Maximilians-Universität München  
MFM – Meshless finite mass  
MHD – Magneto-hydrodynamics  
MNRAS – Monthly Notices of the Royal Astronomical Society  
MS – Main sequence of star-forming galaxies  
NASA – National Aeronautics and Space Administration  
NIR / FIR – Near-infrared / far-infrared  
NOEMA – Northern Extended Millimetre Array  
PAH – Polyaromatic hydrocarbon  
PSB – Post-starburst galaxy  
PSF – Point-spread function  
QFeedS – Quasar Feedback Survey  
RL / RQ – Radio loud / quiet  
SAM – Semi-analytic model  
SDSS – Sloan Digital Sky Survey  
Sgr – Sagittarius  
SKA – Square Kilometre Array  
SMBH – Supermassive black hole  
SN / SNe – Supernova / supernovae  
SPH – Smoothed particle hydrodynamics  
SSD – Shakura-Sunyaev disc  
SFR – Star formation rate  
sSFR – Specific star formation rate  
STScI – Space Telescope Science Institute  
UFO – Ultra-fast outflow  
UV – Ultraviolet  
VLT – Very Large Telescope  
XRB – X-ray binary

# Zusammenfassung

Es wird vermutet, dass Rückkopplung von aktiven galaktischen Kernen (AGN vom englischen active galactic nucleus) eine zentrale Rolle in der Entwicklung von Galaxien spielt, da theoretische Modelle ohne sie außerstande sind entscheidende Beobachtungen wie die Häufigkeit massiver, erloschener Galaxien zu reproduzieren. Starke Ausflüsse, die von Jets, Winden in der Akkretionsscheibe und/oder Strahlungsdruck angetrieben werden, sind eine Art, wie AGN ihre Muttergalaxien beeinflussen können. Es wird angenommen, dass diese Ausflüsse die Sternentstehung regulieren oder unterdrücken – ein Prozess, der auch 'negative Rückkopplung' genannt wird. Allerdings mangelt es an direkten Beobachtungsnachweisen, die zeigen, dass AGN die Sternentstehung oder den Anteil an molekularem Gas in Galaxien verringern. Dies liegt daran, dass AGN vorzugsweise in gasreichen Galaxien, in denen Sterne entstehen, gefunden werden, was in scheinbarem Widerspruch zur Theorie der negative Rückkopplung steht. Darüber hinaus haben Beobachtungsstudien Schwierigkeiten schlüssige Messungen darüber zu liefern, wie effizient AGN Ausflüsse mit dem mehrphasigen Gas koppeln. Sie verlassen sich oft auf vereinfachte analytische Modelle, um wichtige Eigenschaften der Ausflüsse herzuleiten. In dieser Arbeit nutze ich sowohl kosmologische als auch idealisierte Simulationen, um die Auswirkung der Rückkopplung von AGN auf ihre Muttergalaxien zu untersuchen. Dann repliziere ich häufig verwendete Beobachtungstests, um deren Wirksamkeit in der Beurteilung der Auswirkungen von AGN und der Eigenschaften mehrphasiger Ausflüsse zu untersuchen.

Zunächst untersuche ich drei hochmoderne kosmologische Simulationen (IllustrisTNG, EAGLE und SIMBA), welche sich alle auf AGN Rückkopplungsmodellen stützen, um realistische Galaxienpopulationen zu erzeugen. Ich stelle fest, dass AGN sich üblicherweise in Galaxien befinden, die den gleichen oder einen höheren Anteil an Gas aufweisen als solche, die kein AGN beherbergen. Diese Ergebnisse stehen im Einklang mit Beobachtungen in der Fachliteratur bei  $z=0$  und  $z=2$ . Des Weiteren finde ich messbare Unterschiede zwischen Simulationen, die auf verschiedenen Ansätzen zur Modellierung der Rückkopplung basieren. Es ist jedoch nicht immer klar, zu welchem Anteil das Verhalten dieser Simulationen physikalisch gerechtfertigt oder durch die numerische Aufmachung der Modelle bedingt ist. Zudem bedeutet die relativ grobe Auflösung der kosmologischen Simulationen, dass das Zusammenspiel zwischen dem interstellaren Medium (ISM) der Galaxien und dem AGN Ausfluss nicht aufgelöst werden kann.

Um unser Verständnis darüber wie AGN Ausflüsse mit dem ISM koppeln zu verbessern, habe ich eine Reihe idealisierter Galaxiesimulationen entwickelt. Diese beinhalten ein physikalisch-motiviertes AGN Wind Modell und eine manuell einstellbare, mehrphasige ISM Struktur, die auf Sub-Parsec Skalen aufgelöst ist. Das Zusammenspiel zwischen AGN Winden und einem klumpenhaften ISM startet einen mehrphasigen Ausfluss mit kleinen, kalten Gaswolken, die von schnelleren heißen Winden mitgerissen werden. Ich stelle fest, dass die kalte Phase die Masse des

Ausflusses dominiert, während die heiße Phase die Energie dominiert. Ein lange bestehendes Problem in diesem Forschungsgebiet ist das Überleben kalter Wolken wenn diese im Ausfluss beschleunigt werden. Ich stelle jedoch fest, dass ein effizientes Durchmischen und Abkühlen am Übergang zwischen den Phasen den Wolken ein Überleben von mehr als 5 Myr ermöglicht. Ich vergleiche eine klumpenhafte Anordnung mit einem in einer homogenen Scheibe gestarteten Wind und zeige, dass sich die Energetik und Morphologie des Ausflusses signifikant unterscheiden. Die gleichmäßige Scheibe führt zu einer Schale abkühlenden Gases, in der sich beide Phasen gemeinsam und mit derselben Geschwindigkeit bewegen. Ist das ISM klumpenhafte, so resultiert die Fähigkeit des heißen Windes den Niederdruck-Regionen zu entweichen in einer signifikant geringeren Impulsverstärkung und Effizienz der Kopplung von kinetischer Energie als im Falle der gleichmäßigen Scheibe. Unter Verwendung gängiger Beobachtungsansätze zeige ich, wie dies zu einer Fehlklassifizierung von Ausflüssen als impulsgetrieben anstatt energiegetrieben führen kann. Ich prognostiziere zudem Skalierungsbeziehungen zwischen Ausfluss Eigenschaften und der AGN Leuchtkraft und vergleiche diese mit Beobachtungsdaten.

Diese Simulationen zeigen außerdem, dass das Mischen kalter Wolken mit dem AGN Wind das relative Abkühlen innerhalb des gemischten Materials verstärken kann, was wiederum zu heller, den Wolken in Schweifen hinterherhinker Bremsstrahlung führen kann. Das gemischte Material dominiert die erwartete Röntgenstrahlung, was im Gegensatz zu früheren Arbeiten steht, welche ein homogenes ISM annahmen und feststellten, dass das geschockte ISM die meiste Röntgenstrahlung erzeugte. Ich finde eine starke Skalierungsbeziehung zwischen der Röntgenstrahlung und der AGN Leuchtkraft und erörtere ob dieser vorhergesagter Emissionsmechanismus durch Beobachtungen von der durch Sternentstehung und von der AGN Akkretionsscheibe erzeugter Röntgenstrahlung getrennt werden kann. Zum Schluss zeige ich eine simulierte Beobachtung dieser Strahlung mit dem *Chandra* Weltraumteleskop.

**Insgesamt zeigt diese Arbeit, dass Beobachtungen weiterhin mit dem weiten Bild der negativen Rückkopplung übereinstimmen. Dennoch hebt sie auch hervor, dass viele der zeitgenössischen Beobachtungsdiagnostiken nur beschränkt Einblicke in die Effekte der AGN Rückkopplung bieten. Daher erfordert die Verbesserung unserer Beobachtungsarten sowohl ein fundierteres theoretisches Verständnis des Zusammenspiels zwischen AGN Ausflüssen und der umgebenden Galaxie als auch einen verstärkten Dialog zwischen den Beobachtungs- und den Modellierungsgemeinschaften.**



# Abstract

Feedback from active galactic nuclei (AGN) is believed to play a critical role in the evolution of galaxies, as without it, theoretical models are unable to reproduce key observations such as the prevalence of massive quenched galaxies. One way in which AGN can impact their host galaxies is through powerful outflows, driven by jets, accretion disc winds and/or radiation pressure. These are expected to regulate or suppress star formation – a process called ‘negative feedback’. However, direct observational evidence that AGN reduce the star formation or molecular gas content of galaxies is lacking, as, on a population level, AGN are found to preferentially reside in gas-rich, star forming galaxies, seemingly in contradiction to this theory. Furthermore, observational studies struggle to provide conclusive measurements on how efficiently AGN outflows couple to the multiphase gas, and often rely on simplified analytical models to derive important outflow properties. In this thesis, I use both cosmological and idealised simulations to study the effect of AGN feedback on galaxies, and then replicate commonly-used observational tests to investigate their efficacy in evaluating the impact of AGN and the properties of multiphase outflows.

First, I analyse three state-of-the-art cosmological simulations (IllustrisTNG, EAGLE and SIMBA) which all rely on AGN feedback models to create realistic galaxy populations. I find that AGN are typically located in galaxies with equal or higher molecular gas fractions to non-AGN galaxies. These results agree with literature observations at  $z \approx 0$  and  $z \approx 2$ . I also find quantifiable differences between the simulations, driven by their varied approach to modelling feedback. However, it is not always clear to what extent the behaviour of these simulations is physically justified, or driven by the numerical choices of the models. Additionally, the relatively coarse resolution of cosmological simulations means that the interaction between the interstellar medium (ISM) of galaxies and the AGN outflow cannot be resolved.

To improve our understanding of how AGN outflows couple to the ISM, I developed a suite of idealised galaxy simulations. These feature a physically-motivated AGN wind model, and a manually-set, multiphase ISM structure, resolved on sub-parsec scales. The interaction between the AGN wind and clumpy ISM launches a multiphase outflow with small, cold gas clouds entrained in a faster-moving hot wind. I find that the cold phase dominates the mass of the outflow, but the hot phase dominates the energy. A longstanding problem in this field is the survival of cold clouds when they are accelerated in an outflow. However, I find that the efficient mixing and cooling at the interface between the phases allows the clouds to survive on  $> 5$  Myr timescales. I compare the clumpy setup with a wind launched in a homogeneous disc and show that the energetics and morphology of the outflow differ significantly, with the smooth disc resulting in a shell of cooling gas with both phases moving almost co-spatially at the same velocity. When the ISM is clumpy, the ability of the hot wind to ‘vent’ out of low density regions results in a significantly lower

momentum boost and kinetic energy coupling efficiency than the homogeneous case. Using standard approaches from observational work, I show how this could lead to the misclassification of the outflow as momentum-driven rather than energy-driven. I also predict scaling relations between the outflow properties and AGN luminosity and compare these to observational samples.

Additionally, these simulations show that mixing between the cold clouds and the AGN wind can enhance radiative cooling in this mixed material, which I show leads to luminous Bremsstrahlung emission extending in tails behind the clouds. This mixed material dominates the expected X-ray emission, in contrast to previous analytic work, which assumed a homogeneous ISM, that found that the shocked ISM produced the most X-ray emission. I find a strong scaling between the X-ray emission and the AGN luminosity and discuss whether this predicted emission mechanism could be observationally separated from X-rays due to star formation and the AGN accretion disc itself. Finally, I show a mock observation of this emission using the *Chandra* space telescope.

**Overall, this thesis shows that observations remain consistent with the broad picture of negative feedback from AGN. However, it also highlights that many contemporary observational diagnostics offer only limited insight into the effect of AGN feedback. Therefore, improving our observational probes requires both a more fundamental theoretical understanding of the interaction between AGN outflows and the surrounding galaxy, and increased dialogue between observational and computational communities.**

# 1

## Introduction

*With golden string  
our universe was clothed with light.  
Pulling at the seams,  
our once barren world now brims with life.  
That we may fall in love  
every time we open up our eyes.  
I guess space, and time,  
takes violent things, angry things  
and makes them kind.*

*–Sun, Sleeping at Last*



**Figure 1.1:** Stephan's Quintet, a stunning example of ongoing interactions between galaxies. This was one of the first images released after launch of the *James Webb Space Telescope (JWST)* on Christmas Day 2021. The top galaxy (NGC 7319) contains an active galactic nucleus (AGN) at its centre (Image credit: NASA, ESA, CSA & STScI).

## 1.1 Galaxy Evolution

Christmas Day 2021 saw the launch of the *James Webb Space Telescope (JWST)*, which, much as *Hubble* did three decades before, revealed afresh the beauty of the Universe to a new generation. Figure 1.1 presents one of the first images publicly released by *JWST*, showing five galaxies engaged in a complex cosmic dance. Intriguingly, the galaxy at the top of the image contains a mysterious object at its centre; an accreting supermassive black hole, consuming gas from its surroundings and lighting up to become an active galactic nucleus (AGN). The immense power released from this process can be launched back into its surroundings, giving these objects the ability to profoundly shape the evolution of their host galaxies.

This thesis is the story of the life – and death – of galaxies such as these, and the elusive yet powerful role played by the supermassive black holes lurking at their hearts.

### 1.1.1 Structure Formation in $\Lambda$ CDM

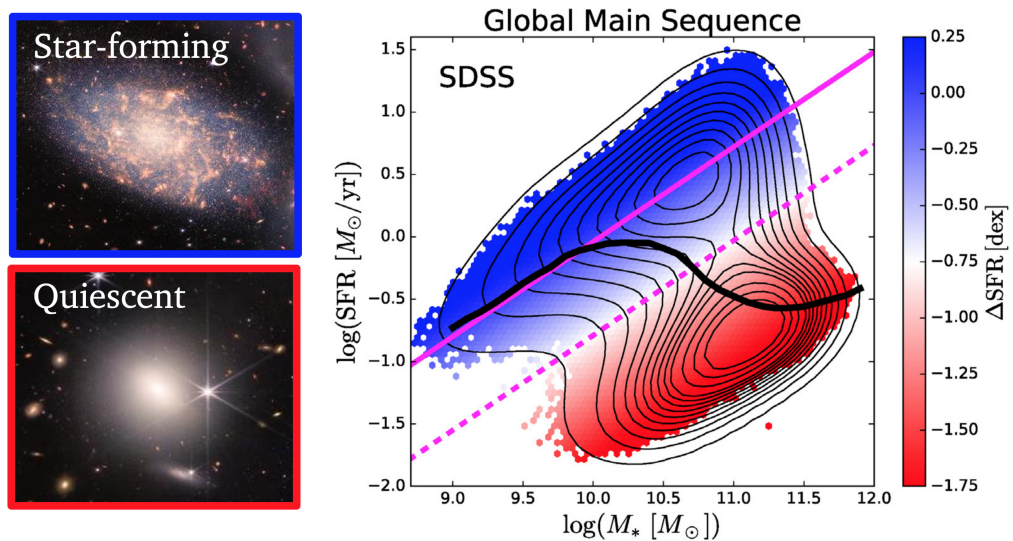
Our story begins barely a fraction of a second after the birth of the Universe, as fluctuations in the hot plasma planted the seeds for the formation of galaxies.

The current leading cosmological theory describing this structure growth is  $\Lambda$ CDM. In this model, the Universe is composed of three main components: dark energy ( $\Lambda$ ), which drives the late-time expansion of the Universe and has a mass density of  $\Omega_\Lambda \approx 0.69$ ; ‘cold’ dark matter (CDM), responsible for gravitational collapse that leads to structure formation ( $\Omega_{\text{DM}} \approx 0.22$ ); and baryonic matter which makes up all the stars, galaxies and the gas between them that can be observed in the Universe ( $\Omega_b \approx 0.05$ ). These parameters can be accurately measured from observations of the Cosmic Microwave Background (CMB; [Planck Collaboration 2020](#)).

In this paradigm, the Universe began 13.8 billion years ago in a hot Big Bang. This was followed by a period of rapid inflation, which both ensured the Universe is isotropic and homogeneous on large scales (the ‘Cosmological Principle’), but also amplified the initial quantum fluctuations in the hot plasma, creating local density perturbations. At a redshift of around  $z \approx 3600$ , the Universe became matter-dominated, allowing the dark matter to start clustering in these initial perturbations. Once the density contrast of these regions reached  $\delta\rho/\bar{\rho} \gtrsim 1$ , the growth continued non-linearly, decoupling from the cosmic expansion, and forming virialised dark matter halos. These halos merged with each other, forming larger and larger structures.

Initially, the baryonic component was coupled to a radiation field which prevented it from collapsing into these dark matter halos. However, at a redshift of  $z \approx 1100$ , the Universe became cool enough for hydrogen atoms to combine with electrons to form neutral atoms. This ‘recombination’ created the cosmic microwave background – the first light we can observe in the Universe. It also decoupled them from the radiation component, allowing them to gravitationally interact with the dark matter overdensities. Unlike dark matter, gas can cool to dissipate its gravitational energy, permitting it to collapse quickly to the centres of these dark matter halos, creating the conditions needed to form proto-galaxies. It is these three components – initially small dark matter perturbations, halo merging and gas accretion – which are the foundation of the ‘bottom-up’ or ‘hierarchical’ model of galaxy formation that is the dominant theory today ([White & Rees, 1978](#); [Blumenthal et al., 1984](#); [White & Frenk, 1991](#); [Cole et al., 2000](#)).

By  $z \approx 20$ , the conditions in these proto-galaxies were cool and dense enough to allow the formation of the first stars (‘Population III’; [Kashlinsky & Rees 1983](#)) and galaxies (*JWST* detections at  $z > 10$ : [Arrabal Haro et al. 2023](#); [Curtis-Lake et al. 2023b](#); [Robertson et al. 2024](#)), bringing the first light into the Universe since recombination. By  $z \approx 6$  the radiation from stars and accreting black holes had excited the neutral gas around them, causing the Universe to become ‘reionised’. This was followed by a steep rise in star formation, reaching a peak at  $z \approx 2$  (‘Cosmic Noon’; [Madau & Dickinson 2014](#)), resulting in around half of all stellar mass in galaxies being formed by  $z \approx 1$  ([Bundy et al., 2005](#)). However, in the 10 billion years since Cosmic Noon, the star formation rate density in the Universe has been steadily decreasing, making our current cosmic era an ‘epoch of galaxy quenching’ ([Curtis-Lake et al., 2023a](#)). This creates a problem, as models based on standard hierarchical structure formation overpredict the the number density of high-mass galaxies compared to what is observed (e.g., [Read & Trentham, 2005](#)). This suggests that there must be some additional processes regulating star formation in galaxies. In Section 1.1.4, we will argue that this role is filled by baryonic feedback process, but first, we need to understand the structures within which stars are created – galaxies.



**Figure 1.2:** The main sequence of star formation. On the left, we show examples of a star-forming and quiescent galaxy, taken from Stephan’s Quintet (Figure 1.1). On the right, we show a plot from [Bluck et al. \(2020\)](#) to show the bimodality in star formation rates seen in SDSS data. The magenta line is a fit to the main sequence fit from [Renzini & Peng \(2015\)](#) and the dashed magenta shows the minima of this bimodality. The sold black shows the median position of galaxies as a function of stellar mass.

### 1.1.2 Galaxy Classification

As Figure 1.1 shows, there is a stunning diversity of galaxies in the night’s sky, with mergers, gas accretion and star formation all etching out a unique history for each galaxy. However, on a population level, galaxies can be classified into two main groups, based on star formation rates (e.g., [Brinchmann et al., 2004](#); [Noeske et al., 2007](#)), colour (e.g., [Baldry et al., 2004](#)), and morphology (e.g., [Schawinski et al., 2014](#)).

**Star-forming** galaxies have high levels on ongoing star formation ( $\text{SFR} \gtrsim 1 M_{\odot} \text{ yr}^{-1}$ ) and are rich in atomic (HI) and molecular gas ( $\text{H}_2$ ). They typically contain a rotationally-supported disc of gas and stars, which may feature prominent spiral arms and/or a bar (‘spiral’ or ‘late-type’ galaxies). Their ongoing star formation results in their light being dominated by young, massive stars, giving them a blue colour. **Quiescent** galaxies are typically elliptical in morphology (‘early-type’) and have low gas fractions. They have low levels of star formation, which can due to a lack of fuel and/or inefficient conversion of gas into stars. With few young, massive stars being created, their light is dominated by older stellar populations, making them redder in colour.

Figure 1.2 shows a plot of stellar mass against star formation rate for galaxies in the Sloan Digital Sky Survey (SDSS). We can see a clear bimodal separation in star formation rate between the two classes of galaxies. Galaxies in the ‘blue cloud’ of this plot (with high SFRs) are referred to as being on the ‘main sequence’ of star formation – a fit to this relation this is shown by the magenta line ([Renzini & Peng 2015](#), see also [Speagle et al. 2014](#); [Whitaker et al. 2014](#)). Galaxies lying significantly below this line are then defined as quiescent and have low ongoing levels of star formation – they have been ‘quenched’. The black line shows the median position of a galaxy in this plane and shows that higher stellar mass galaxies are more likely to be quenched. As a galaxy grows its stellar mass, it becomes more likely to move

from the main sequence to the quiescent region. This means that, through some process, the gas fraction and star formation rate of the galaxy drops. This process is called galaxy quenching.

### 1.1.3 Quenching

Galaxies are complex systems and are constantly cycling gas to and from their surrounding environment (the ‘baryon cycle’). Stars are formed within the interstellar medium (ISM) of the galaxy. This process required cold clouds of molecular hydrogen ( $\text{H}_2$ ), which are then able to collapse to form new stars. To maintain a sufficient cold gas reservoir, the ISM needs to be replenished by a steady supply of inflowing gas from the circumgalactic medium (CGM) which surrounds the galaxy for 10 – 100 kpc (Tumlinson et al., 2017; Langan et al., 2023). Beyond this, at megaparsec scales, lies the gas situated between groups and clusters of galaxies, called the intragroup (IGrM) or intracluster (ICM) medium (e.g., Fabian, 1994). Star formation is therefore both dependent on the continual flow of gas between these vast spatial scales and the sub-parsec interactions and cooling of molecular gas clouds. For a galaxy to quench, one or more of these processes must be disrupted somehow (Man & Belli, 2018).


On the CGM scale, gas inflows can be disrupted, ‘starving’ the galaxy from new fuel (e.g., Feldmann & Mayer, 2015). The halo can be heated, for example, by virial shocks (Rees & Ostriker, 1977), which prevents the accretion of gas onto the galaxy. In the ISM, the star formation efficiency of the cold gas could be suppressed; for example, a large stellar bulge (Martig et al., 2009) or a bar (Khoperskov et al., 2018) could increase the turbulence of the gas, stabilising the disc from fragmentation (‘morphological’ quenching). A rapid burst of star formation, perhaps triggered by a major merger (Mihos & Hernquist, 1996; Langan et al., 2024) or disc instability (Zolotov et al., 2015), could consume the available fuel, leading to the suppression of future star formation. Additionally, the environment a galaxy lives in plays an important role. Ram-pressure stripping can quench satellite galaxies (Martín-Navarro et al., 2021) and gravitationally heat the halo of central galaxies as a satellite infalls (Khochfar & Ostriker, 2008). The timescale of quenching can also differ significantly between galaxies, possibly reflecting the vastly different spatial scales these processes operate on. The timescales can range from  $\approx 100$  Myr for rapid quenching pathways, to  $> 1$  Gyr for slow quenching (e.g., Wu et al., 2018; Belli et al., 2019; Akins et al., 2022), although this is hard to constrain.

However, for many of these mechanisms, even if they could reduce a galaxy’s star formation temporarily, a fresh inflow of gas could rejuvenate star formation. Therefore, to fully quench the galaxy the CGM gas supply and/or ISM conditions must remain suppressed. This leads us to the question: are these processes sufficient to quench galaxies to the level seen observationally?

### 1.1.4 The Need for Baryonic Feedback

In Figure 1.3 we show the mean baryon to star conversion efficiency ( $\epsilon_\star = M_\star/M_h$ ) as a function of halo mass (Harrison 2017, see also Henriques et al. 2019). The grey line shows the trend from a hierarchical model of galaxy formation without feedback included (Somerville et al., 2008) showing an efficient conversion of baryons to stars. This process peaks at  $\epsilon_\star \approx 60\%$ , and has a slight drop at the low-mass end (due to environmental effects on satellite galaxies) and a larger suppression at the high mass end (due to the long cooling times of massive halos). However,





**Figure 1.3:** A plot of how efficiently galaxies convert their baryons to stars (Harrison, 2017). The green line shows a semi-empirical fit to observational data from Moster et al. (2013) and the other lines show the results from a semi-analytic galaxy formation model (Somerville et al., 2008). At the high mass end, AGN feedback is an essential component of our theoretical models to sufficiently suppress star formation enough to match the observed Universe.

when compared to a semi-empirical fit based on halo abundance matching Moster et al. (2013), we can see that the galaxy formation model over-predicts the star formation efficiency at all halo masses. Galaxy quenching based on environmental or secular processes alone (grey line) is insufficient to explain the observed population of massive quenched galaxies (green line). There must therefore be additional processes operating within galaxies, altering the evolution of the baryonic component in such a way to reduce the efficiency of star formation. These processes are referred to as baryonic feedback, specifically negative feedback, as they are reducing the star formation rate of galaxies. At the low mass end, this feedback is dominated by star formation processes; namely, energy injection by supernova explosions and massive stellar winds (e.g., Dekel & Silk, 1986; Heckman et al., 1990; Somerville & Primack, 1999; Hayward & Hopkins, 2017; Girichidis et al., 2020). The purple line in Figure 1.3 shows the result when star formation feedback is included, demonstrating good agreement between the model and observed data. However, it fails to reproduce the suppression at the high-mass end. Massive galaxies ( $M_{\star} \gtrsim 10^{10} M_{\star}$ ), therefore, require a more energetic feedback mechanism. The leading explanation is feedback from accreting supermassive black holes, which are believed to lie at the centres of all massive galaxies. These Active Galactic Nuclei (AGN) can launch huge jets and winds, injecting significant energy into their host galaxies, and, as shown by the black line in Figure 1.3, are theoretically able to quench massive galaxies (e.g., Bower et al., 2006; Croton et al., 2006; Fabian, 2012).

The rest of this thesis is devoted to exploring the mechanisms, successes, challenges and future of the theory of AGN feedback.

## 1.2 Active Galactic Nuclei

The processes associated with AGN feedback span a vast range of spatial scales. The Schwarzschild radius of a  $10^9 M_\odot$  black hole is  $5 \times 10^{-6}$  pc (approximately the size of the orbit of Uranus) and yet it can launch jets that span hundreds of kiloparsecs. This represents a dynamic scale of over ten orders of magnitude and encompasses a broad range of physical processes. In this Section, we will present an overview of these processes. Starting at the black hole itself, we will introduce how black holes accrete via an accretion disc; how this disc can release energy back into the galaxy; the observable signatures of AGN activity; and finally the theorised feedback impact on the host galaxy.

### 1.2.1 Supermassive Black Holes

Black holes are enigmatic objects; singularities with a gravitational field so strong even light cannot escape. Most black holes are formed from the remnants of massive stars ( $M_{\text{BH}} \gtrsim 8 M_\odot$ ; ‘stellar-mass’ black holes) but other are much larger, with masses  $M_{\text{BH}} \approx 10^{6-10} M_\odot$ . These ‘supermassive’ black holes (SMBHs)<sup>1</sup> are believed to lie at the centre of every massive galaxy and their masses correlate strongly with the stellar velocity dispersion of their host galaxies (Kormendy & Ho, 2013). Even our own Milky Way hosts one, as evidenced by the precession of stellar orbits close to the centre of our galaxy (Ghez et al., 2008; Genzel et al., 2010), and direct imaging with the Event Horizon Telescope (EHT; Event Horizon Telescope Collaboration et al., 2022). The origin of these objects is still shrouded in mystery, with recent *JWST* observations finding evidence for their existence within a few hundred millions years of the Big Bang (e.g., Larson et al., 2023), posing new challenges for our models of early black hole seeding and growth (Inayoshi et al., 2020). However, despite their large masses, the gravitational influence of a central SMBH is negligible compared to the host galaxy (Peebles, 1972). Using values for the mass of the SMBH ( $M_{\text{BH}}$ ) and stellar velocity dispersion ( $\sigma_\star$ ) in M87 (Gültekin et al., 2009), we find a radius of gravitational influence of:

$$r_g = \frac{GM_{\text{BH}}}{\sigma_\star^2} \approx \frac{G \cdot 7 \times 10^9 M_\odot}{(375 \text{ km s}^{-1})^2} \approx 110 \text{ pc} \quad (1.1)$$

which is less than 0.3% of the stellar radius of M87 ( $R \approx 40$  kpc) – and M87 has one of the most massive black holes we know! Thus, if SMBHs are to have a large impact on the evolution of their host galaxies, it must be through another channel than purely gravitational influence.

### 1.2.2 Accretion

If the supermassive black hole has ready access to a gas supply, it can accrete this material, growing its mass. Reducing the angular momentum of infalling gas sufficiently for it to be accreted can be challenging, and is likely to be triggered both by galaxy mergers (e.g., Treister et al., 2010), which are expected to dominate at high redshift, and secular processes such as bar instabilities (e.g., Schawinski et al., 2011; Verwilghen et al., 2024), which dominate at low redshift. During periods of

<sup>1</sup>From now on, we will not consider stellar mass black holes, so the terms black hole, supermassive black hole and SMBH will be used interchangeably.

mass accretion, the SMBH releases energy (either radiatively or mechanically, see Section 1.2.5) and becomes known as an **Active Galactic Nucleus** (AGN). The power released at a given accretion rate ( $\dot{M}_{\text{accr}}$ ) is given by:

$$L_{\text{AGN}} = \epsilon_r \dot{M}_{\text{accr}} c^2 \quad (1.2)$$

leading to an increase in black hole mass of:

$$\dot{M}_{\text{BH}} = (1 - \epsilon_r) \dot{M}_{\text{accr}} \quad (1.3)$$

where  $\epsilon_r$  is the radiative efficiency of the black hole. This can vary between  $\epsilon_r \approx 5 - 40\%$  depending on the spin of the black hole (Thorne, 1974), but is typically assumed to be around  $\epsilon_r \approx 10\%$  (Soltan, 1982; Yu & Tremaine, 2002). This represents an incredibly efficient process: for comparison, in stellar nuclear fusion, only around  $\epsilon_r \approx 0.7\%$  of the rest mass is converted to energy.

The **Eddington luminosity** is the accretion rate for a spherical system at which the resulting radiation pressure from the SMBH accretion balances the infalling gravitational force:

$$L_{\text{Edd}} = \frac{4\pi G M_{\text{BH}} m_p c}{\sigma_T} \approx 1.26 \times 10^{46} \left( \frac{M_{\text{BH}}}{10^8 M_{\odot}} \right) \text{ erg s}^{-1} \quad (1.4)$$

where  $\sigma_T$  is the Thompson cross-section and  $m_p$  is the proton mass. The luminosity of the AGN is often quoted as a fraction of this Eddington luminosity:

$$\lambda_{\text{Edd}} = \frac{L_{\text{AGN}}}{L_{\text{Edd}}} \quad (1.5)$$

which is known as the **Eddington ratio**. A ratio of  $\lambda_{\text{Edd}}=1$ , therefore, represents an idealised maximum limit on the accretion rate in the case of spherical symmetry – the **Eddington limit**.

### 1.2.3 Energy Coupling

Over their lifetimes, AGN can release a huge amount of energy. Assuming their mass growth is dominated by accretion, we can estimate the total energy released by  $E_{\text{AGN}} \approx \epsilon_r M_{\text{BH}} c^2$ . We can compare this to the binding energy of the galaxy (approximated as  $E_{\text{bind}} \approx M_{\star} \sigma_{\star}^2$ , where  $\sigma_{\star}$  is the velocity dispersion of stars in the bulge). Using the observationally-derived scaling between black hole and stellar mass from Kormendy & Ho (2013) and reasonable estimates for the stellar dispersion, this yields an energy ratio of:

$$\frac{E_{\text{AGN}}}{E_{\text{bind}}} \approx 500 \left( \frac{\epsilon_r}{0.1} \right) \left( \frac{M_{\text{BH}}/M_{\star}}{5 \times 10^{-3}} \right) \left( \frac{\sigma_{\star}}{300 \text{ km s}^{-1}} \right)^{-2} \quad (1.6)$$

This demonstrates that the energy released from black hole mass accretion can exceed the binding energy of the host galaxy by orders of magnitude, suggesting that AGN could have a large impact on the evolution of their host galaxies. However, this impact will be strongly dependent on how efficiently this released energy couples to the gas in the galaxy. We can define a feedback efficiency ( $\epsilon_f$ ) to quantify the fraction of the AGN power that couples to the galaxy:

$$\dot{E}_{\text{feed}} = \epsilon_f L_{\text{AGN}} = \epsilon_f \epsilon_r \dot{M}_{\text{accr}} c^2 \quad (1.7)$$

The feedback efficiency is a free parameter and there is no clear consensus on its value (see discussion in [Harrison et al., 2018](#)). In one of the first simulations of AGN feedback, [Di Matteo et al. \(2005\)](#) calibrated the parameter to match the normalisation of the  $M_{\text{BH}} - \sigma$  relation, obtaining  $\epsilon_f \approx 5\%$ . Cosmological simulations often use a value of  $\epsilon_f \approx 10\%$ , however, this is chosen for numerical reasons due to the unresolved AGN feedback models used in these simulations (e.g., [Booth & Schaye, 2009](#)), so caution must be applied when comparing to observations (see discussion in Chapter 2.3.3).

Analytic work assuming a fast nuclear wind as a feedback mechanism (based on [King, 2003](#)) found that the feedback efficiency scaled strongly with the wind speed, from  $\epsilon_f \approx 0.0005 - 5\%$  for velocities of  $v_w \approx 1000 - 30000 \text{ km s}^{-1}$ . Work done against the gravitational potential and ambient pressure can also reduce the energy available to be transferred to the gas ([Richings & Faucher-Giguère, 2018a](#)). However, these theoretical studies were performed in spherically symmetric setups within a homogeneous medium. It is not clear how the coupling would change if the gas were distributed inhomogeneously and the wind were able to escape along low-density channels.

Observational works have attempted to measure the kinetic coupling efficiency of AGN outflows, finding a large range of values from  $\dot{E}_k/L_{\text{AGN}} = 0.001 - 10\%$  (e.g., [Fiore et al., 2017](#)). However, the wide range of assumptions, techniques and outflow tracers used render any comparison highly challenging ([Veilleux et al., 2020](#); [Harrison & Ramos Almeida, 2024](#)).

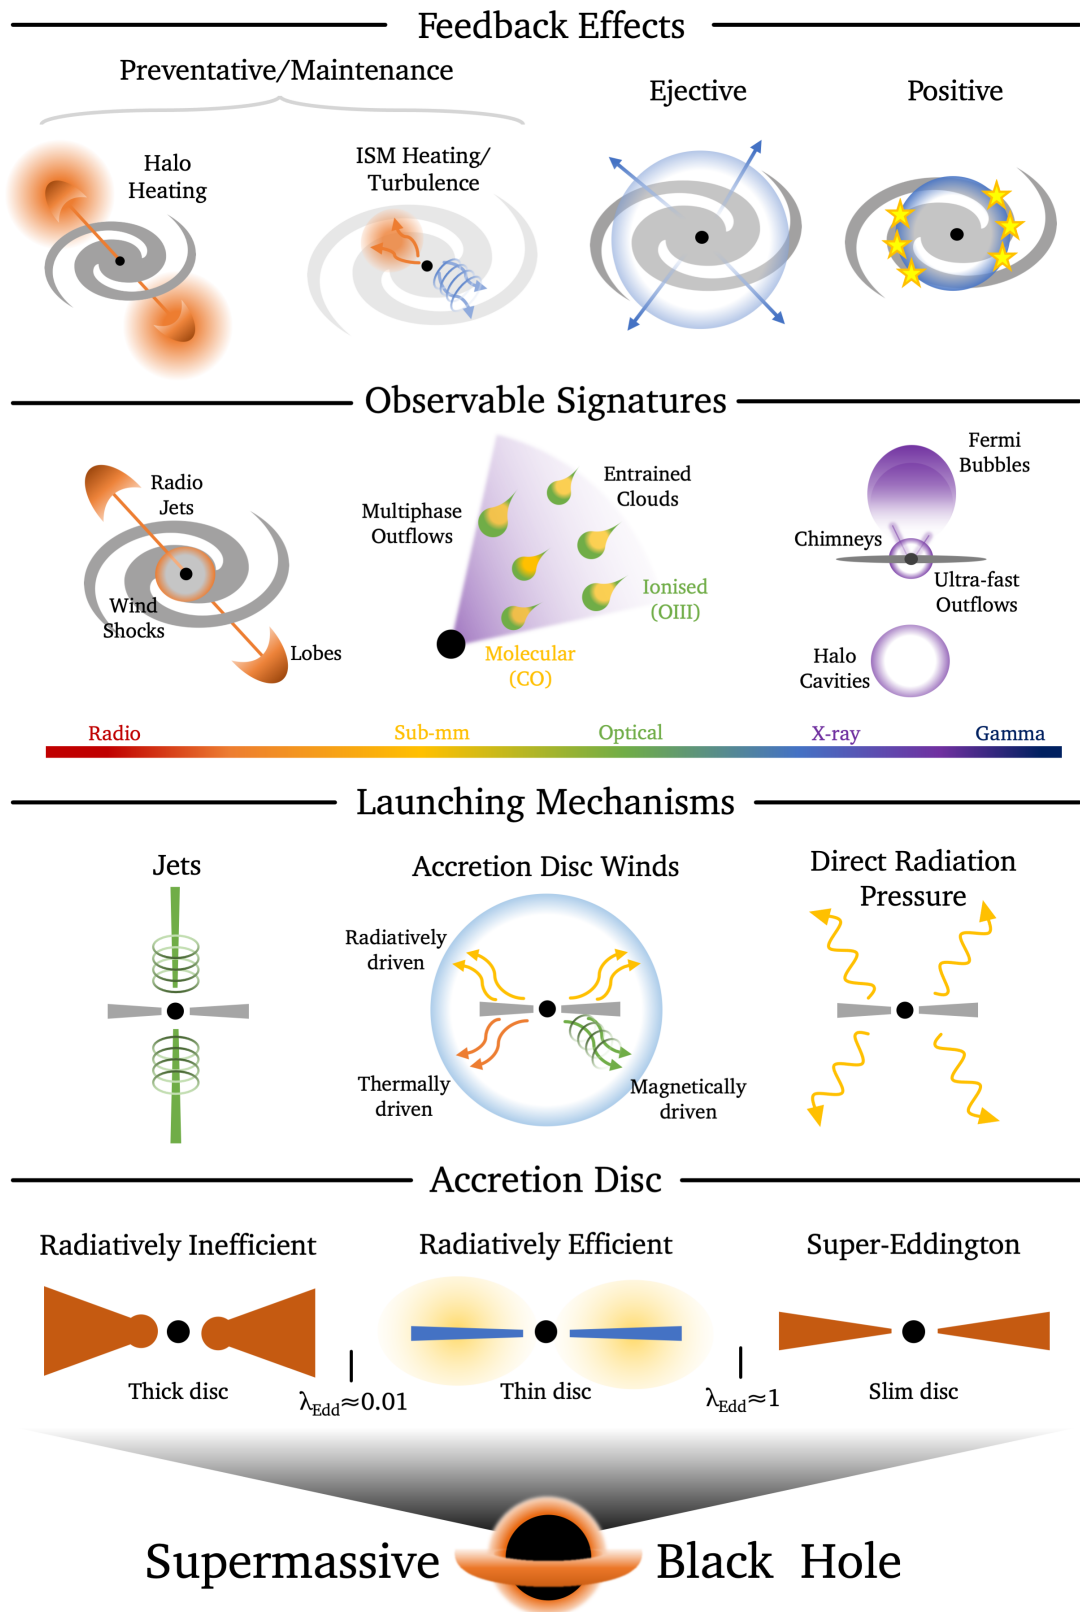
In Chapters 2 & 3, we describe in more detail the feedback prescriptions in cosmological simulations and discuss some of the challenges when comparing to observations. In Chapter 4, we address the issue of AGN feedback in an inhomogeneous environment and discuss the implications this has for measuring observed outflow properties. Now we have motivated the potential for AGN feedback to affect its host galaxy, we describe the various physical processes responsible for gas accretion onto the SMBH and energy launching from it.

#### 1.2.4 Accretion Discs

Gas accretion onto SMBHs is mediated through an **accretion disc** which orbits the black hole. Recent studies have suggested there may be three distinct accretion disc regimes, depending on the accretion rate of the black hole. We show a schematic overview of these regimes in Figure 1.4.

The classic example of an accretion disc is the Shakura-Sunyaev disc (SSD; [Shakura & Sunyaev 1973](#)) which have Eddington ratios of  $\lambda_{\text{Edd}} \approx 0.01 - 1$ . The disc is geometrically thin, optically thick, and accretes relatively cold gas. The disc is radiatively efficient ( $\eta_r \approx 10\%$ ), as all the heat generated by viscosity within the disc is immediately radiated away. The energy output is dominated by radiation, but there may also be non-negligible mechanical components. Radiatively efficient discs can be found in quasars, Seyfert galaxies and high-excitation radio galaxies (HERGs; [Best & Heckman 2012](#)).

If the accretion rate drops to  $\lambda_{\text{Edd}} \lesssim 0.01$ , the disc can be described as an advection-dominated accretion flow (ADAF; [Narayan & Yi 1994](#)). These discs are hot (close to the virial temperature), optically thin and have a quasi-spherical (or ‘thick’) geometry. They are radiatively inefficient ( $\epsilon_r \propto \dot{M}_{\text{acc}}/\dot{M}_{\text{Edd}}$ ), as the dissipated energy is advected into the SMBH faster than it can be radiated away ([Abramowicz &](#)



**Figure 1.4:** A schematic overview of the different processes involved in AGN feedback, showing the broad range of scales and mechanisms that can affect the evolution of galaxies.

[Fragile, 2013](#)). These discs can be found in low-excitation radio galaxies (LERGs) and low-luminosity AGN (LLAGN), such as the SMBH at the centre of the Milky Way (Sgr A\*).

The third accretion regime is when the accretion rate approaches or exceeds the Eddington limit  $\lambda_{\text{Edd}} \gtrsim 1$ , resulting in ‘Super-Eddington’ accretion. As the accretion rate increases, so too does the optical thickness of the gas, reducing the ability of the disc to radiate its dissipated energy ([Yuan & Narayan, 2014](#)). This results in a ‘super-critical’ advection-dominated flow ([Madau et al., 2014](#)), with a radiative efficiency that decreases with increasing accretion rate ( $\epsilon_r \propto (\dot{M}_{\text{acc}}/\dot{M}_{\text{Edd}})^{-1}$ ). These discs are slightly geometrically thicker than thin SSDs, resulting in the name ‘slim discs’ ([Sądowski, 2009](#)). It is difficult to find examples of Super-Eddington accretion onto black holes in the local Universe, but it may be essential in explaining the rapid growth of SMBHs at high redshift (e.g., [Madau et al., 2014](#); [Larson et al., 2023](#); [Lupi et al., 2024](#)).

### 1.2.5 Launching Mechanisms

The huge amount of gravitational, radiative and magnetic energy contained in accretion discs can result in hugely energetic outbursts of material into the wider galaxy. There are broadly three main mechanisms mediating this release of energy, as sketched in Figure 1.4.

**Jets** are narrow beams of relativistic particles ([Blandford et al., 2019](#)). Due to the rapid rotation of the black hole, frame-dragging twists the magnetic field of the highly-magnetised accretion disc into a poloidal geometry, aligned with the angular momentum vector of the SMBH. This accelerates relativistic plasma in collimated beams along this axis which are known as Blandford-Znajek jets ([Blandford & Znajek, 1977](#)). This effect has been successfully replicated in accretion-disc-scale general relativistic magnetohydrodynamic (GRMHD) simulations (e.g., [McKinney, 2006](#); [Sądowski et al., 2015](#); [Chatterjee et al., 2019](#)). Jets are commonly considered to originate from radiatively inefficient, advection-dominated accretion flows, however, recent studies have found that thin discs can also launch jets ([Liska et al., 2019, 2020](#)), possibly helping to explain the bright radio emission seen in some quasars. It is worth noting that the alignment of the SMBH spin (and thus the jet direction) may not match the angular momentum vector of the galaxy as a whole. Furthermore, torques due to accretion and SMBH mergers can change the spin alignment, causing the jet to precess (e.g., [Talbot et al., 2024](#)). This can cause the jet to become inclined into the galactic plane, which has important implications for the effect it has on the host galaxy (e.g., [Mukherjee et al., 2018](#); [Tanner & Weaver, 2022](#)), which we will address later.

**Accretion disc winds** are fast ( $v_w/c \approx 0.1 - 0.3$ ), wide-angle winds which may be radiatively-, magnetically-, or thermally-driven from the SMBH accretion disc or dusty torus. Radiatively efficient thin discs release a large amount of energy in the form of radiation. This can couple to nearby gas through electron scattering or, more efficiently, ultraviolet (UV) line driving, propelling a wide-angle wind ([Stevens & Kallman, 1990](#); [Murray et al., 1995](#); [Proga et al., 2000](#); [Mizumoto et al., 2021](#)). In radiatively inefficient discs, winds can be launched thermally ([Almeida & Nemmen, 2020](#); [Almeida et al., 2023](#)) and magnetically ([Yuan et al., 2015](#)), often alongside jets ([Yang et al., 2021](#)). These small-scale driving forces sweep up surrounding ambient gas, shocking it to high temperatures and creating a wide-angle or spherical outflowing shock front which then propagates into the host galaxy ([King, 2003, 2005](#);

Zubovas & King, 2012; Faucher-Giguère & Quataert, 2012). The resulting outflow can either be momentum-driven, in the case of efficient cooling, or energy-driven, where the outflow expands adiabatically (Costa et al., 2014b). We discuss these two cases in more detail in Chapter 2, where we also describe a numerical implementation of this model (Costa et al., 2020) which will be used in Chapters 4 & 5.

**Radiation pressure** acts directly on gas and dust at large (100–1000 pc) scales rather than being mediated through a wind as in the previous case (e.g., Fabian, 1999; Thompson et al., 2015; Bieri et al., 2017; Ishibashi et al., 2018). The optical and UV radiation is absorbed and re-emitted in the infrared (IR). If the surrounding gas has low density, the IR photons escape, and the AGN struggles to produce strong ( $\dot{p} > L_{\text{AGN}}/c$ ) outflows (e.g., Ciotti & Ostriker, 2007, 2012; Novak et al., 2012). However, if the column density of the gas is high ( $N \gtrsim 10^{23-24} \text{ cm}^{-2}$ ), the gas becomes optically thick to IR radiation and the photons can scatter multiple times, leading to much more powerful outflows (e.g., Costa et al., 2018a). However, Costa et al. (2018b) suggest that only the brightest  $L_{\text{AGN}} \gtrsim 10^{47} \text{ erg s}^{-1}$  quasars are able to launch fast outflows via radiation pressure alone, perhaps limiting this feedback channel to only the more extreme systems.

## 1.2.6 Observational Signatures


The energy emitted by AGN in the form of jets, winds and radiation can have observable effects on the host galaxy on kiloparsec scales. We briefly summarise some of multiwavelength signatures observed in active galaxies.

### Radio Structures

Some of the most striking evidence for AGN activity is the detection of large radio jets, spanning kiloparsec, or even megaparsec scales (Willis et al., 1974). These large-scale jets are launched from a central AGN, remain collimated as they escape the galaxy, and terminate in extended lobes and hotspots in the CGM or ICM (Begelman et al., 1984). At radio frequencies below 10 GHz, the flux is dominated by non-thermal emission from relativistic electrons (Condon & Yin, 1990), allowing radio observations to trace some of the most energetic processes occurring in galaxies. Fanaroff & Riley (1974) introduced a morphological classification of radio jets into ‘limb-brightened’ (FR-I) if the emission is brightest along the inner jet, or ‘edge-brightened’ (FR-II) if the jet has produced bright hotspots at the edge of the radio lobes.

Around 10% of optically-selected AGN are classified as ‘radio-loud’ (RL), meaning they are dominated by bright radio emission associated with powerful jets.<sup>2</sup> The remaining  $\sim 90\%$  are termed ‘radio-quiet’ (RQ). The origin of radio emission in RQ AGN is still debated, but may be produced by star formation (e.g., Condon & Yin, 1990; Padovani et al., 2015), low-powered jets, winds, and/or accretion disc coronae (Panessa et al., 2019). For example, Jarvis et al. (2019); Jarvis et al. (2021) showed that even radio-quiet quasars can host extended jets. Some numerical works have shown that even low-powered jets can have a significant effect on their galaxy’s ISM if they become inclined into the galaxy disc, driving multiphase outflows (Mukherjee et al., 2018; Tanner & Weaver, 2022). This effect has been seen observationally, with both Venturi et al. (2021); Girdhar et al. (2022) finding inclined jets that are causing enhanced velocity dispersion perpendicular to the jet direction.

<sup>2</sup>There are a variety of definitions for radio loudness (e.g., Kellermann et al., 1989; Padovani et al., 2017) discussion of which is beyond the scope of this thesis (see also Harrison & Ramos Almeida 2024).



**Figure 1.5:** An observational sample from [Fiore et al. \(2017\)](#) showing scaling relations between the mass outflow rate (*left*) and kinetic energy rate (*right*) with AGN luminosity, for a range of outflow gas phases.

A potentially important, but underappreciated, source of radio emission can be produced by wind shocks in the ISM. Shocks can accelerate electrons to relativistic velocities, resulting in synchrotron emission ([Nims et al., 2015](#)). Wind shocks have been invoked to explain the observed correlation between ionised outflows and radio emission in quasars (e.g., [Zakamska & Greene, 2014](#); [Jarvis et al., 2021](#); [Petley et al., 2022](#); [Fawcett et al., 2023](#)) and may be caused by accretion disc winds, or low-powered jets driving outflows into the ISM.

### Multiphase Outflows

One of the key observational indicators of AGN activity is the presence of multiphase outflows. Once the central AGN has released its energy through jets, accretion disc winds and/or direct radiation pressure on dust, these driving mechanisms move beyond the nucleus and interact with the surrounding medium, sweeping up gas to form a large-scale outflow.<sup>3</sup> These outflows not only span a wide range of spatial scales, but also contain gas in different phases, resulting in a range of observational tracers that can be used to study them ([Veilleux et al., 2020](#); [Harrison & Ramos Almeida, 2024](#)).

The coldest gas is in the **molecular** phase ( $\text{H}_2$ ;  $T \approx 10^{1-3}$  K). Molecular hydrogen can be probed directly through roto-vibrational  $\text{H}_2$  lines in the infrared (warm molecular; e.g., [Dasyra & Combes 2011](#); [Rupke & Veilleux 2013](#)) or indirectly using tracers such as hydroxyl (OH) in the far-infrared (e.g., [Sturm et al., 2011](#); [González-Alfonso et al., 2017](#)) or carbon monoxide (CO) in the sub-millimetre (e.g., [Cicone et al., 2014](#); [Molyneux et al., 2019](#); [Lamperti et al., 2022](#)). Sub-mm interferometers such as ALMA and NOEMA have allowed the molecular phase to be mapped at high resolution, revealing sub-kpc structure (e.g., [Girdhar et al., 2024](#)).

At slightly warmer temperatures ( $T \approx 10^{2-3}$  K), we can probe the **neutral** phase of the gas (H I), either directly, through the 21 cm line (e.g., [Morganti et al., 2005](#)) or indirectly via sodium doublet absorption (Na I D; [Rupke et al., 2005](#); [Concas](#)

<sup>3</sup>We note that the terminology in the literature can be unclear, with ‘wind’ and ‘outflow’ sometimes being used interchangeably. We follow [Harrison & Ramos Almeida \(2024\)](#) and define a ‘wind’ strictly as the accretion disc launching mechanism, and an ‘outflow’ as the material swept up by a wind, jet and/or radiation pressure on galaxy scales.



et al., 2019). Furthermore [C II] fine-structure emission can be used to probe cold gas, although it is debated whether this traces the neutral or molecular phase (see discussion in Zanella et al. 2018).

**Warm ionised** gas ( $T \approx 10^{3-5}$  K) is traced through emission lines in the rest-frame UV, optical (commonly [O III]) and NIR (e.g., Murray et al., 1995; Harrison et al., 2014; Rose et al., 2018; Riffel et al., 2023). The advent of integral field units on large optical telescopes such as the VLT have allowed detailed mapping of optical-NIR emission line distributions on sub-kpc scales (e.g., Girdhar et al., 2022).

**Hot ionised** gas ( $T \approx 10^{5-8}$  K) is hard to study due to its low density, but has been observed using X-ray absorption lines (e.g. Longinotti et al., 2013; Tombesi et al., 2013; Chartas et al., 2021). We will briefly touch on other uses of X-rays to study AGN in the next section.

Studies generally conclude that molecular outflows are slower ( $v \lesssim 500$  km s<sup>-1</sup>) and more compact ( $r \lesssim 1 - 2$  kpc) whereas warm ionised outflows are faster ( $v \approx 500 - 1000$  km s<sup>-1</sup>) and extend to greater radii ( $r \approx 1 - 10$  kpc). There have been a range of studies looking at how the properties of outflows, such as the mass outflow rate and kinetic energy coupling, scale with AGN luminosity in different gas phases (an example from Fiore et al. 2017 is shown in Figure 1.5; see also Leung et al. 2019; Bischetti et al. 2019; Ramos Almeida et al. 2022; Lamperti et al. 2022; Musiimenta et al. 2023). However, there are many challenges when making these measurements such as de-blending the host galaxy from the outflow contribution and constraining conversion factors between line fluxes and total gas masses (Harrison et al., 2018; Veilleux et al., 2020; Davies et al., 2020; Holden et al., 2023; Harrison & Ramos Almeida, 2024). We discuss these challenges further in Chapter 4.

## X-ray Emission

X-ray observations have detected small-scale, highly ionised winds in the centres of some AGN-hosting galaxies. These are known as ultra-fast outflows (UFOs) due to their high velocities ( $v_w/c \approx 0.1 - 0.3$ ) and have been detected in  $\approx 40\%$  of AGN (Pounds et al., 2003a,b; Tombesi et al., 2011, 2012). These UFOs are likely directly tracing accretion disc winds from the AGN, and the simultaneous detections of UFOs and galaxy-scale outflows lends weight to the argument for these small-scale winds as powerful driving mechanisms (Tombesi et al., 2015; Veilleux et al., 2017; Sirressi et al., 2019). In Chapter 2 we describe a numerical implementation of a UFO-like wind (Costa et al., 2020) and in Chapters 4 & 5 we present the results of such a wind interacting with a clumpy ISM.

As we mentioned in the previous section, X-ray-producing gas on larger scales is hard to study due to its low density. However, there have been several successes in detecting kiloparsec-scale X-ray bubbles in the CGM of nearby galaxies (e.g., Croston et al., 2008; Greene et al., 2014; Lansbury et al., 2018). Excitingly, our own Milky Way shows similar signatures of AGN activity. In 2010, the *Fermi* gamma ray satellite detected two bubbles extending 25 kpc above and below the galactic plane (Su et al., 2010) which were later also observed with eROSITA in the X-ray (Predehl et al., 2020), providing direct evidence for previous outbursts from Sgr A\* (Zubovas et al. 2011; Zhang & Guo 2020; Yang et al. 2021, although others have suggested a starburst origin, see for example Lacki 2014). X-ray observations have also revealed narrow ‘chimneys’, around 160 pc in length, linking the galaxy centre to the *Fermi*/eROSITA bubbles (Ponti et al., 2019, 2021). In Chapter 5, we explore some of these X-ray signatures of wind-driven outflows.

Moving to larger scales, the intracluster medium (ICM) of the most massive dark matter halos are permeated by hot X-ray emitting gas. Observations of this gas – first with *ROSAT*, and later with *Chandra* and *XMM-Newton* – revealed cavities in the X-ray emission from the ICM (Boehringer et al., 1993) with around 95% of clusters showing these holes (Dunn & Fabian, 2006; Fabian, 2012). These cavities were found to be spatially coincident with radio lobes from powerful AGN, demonstrating that these jets were transferring the AGN power out to these large scales. The possible feedback affect this has on massive galaxies is described in the next section.

### 1.2.7 Theorised Feedback Impact

We have seen that galaxies contain SMBHs; these SMBHs can become AGN when they accrete; AGN can launch winds and jets; and that these winds and jets can drive large-scale outflows into the ISM and CGM of the galaxy. But what effect does this have on the evolution of the host galaxy? As shown in Figure 1.3, according to the current paradigm of galaxy evolution, the population-level impact of AGN feedback on star formation must overall be negative. Figure 1.4 shows the main mechanisms that have been suggested to explain how AGN can negatively impact their host’s star formation:

At the largest scales, AGN feedback can heat the gas around galaxies (at CGM, IGM and ICM scales), preventing accretion of gas onto the galaxy and thus starving it of future fuel for star formation. This type of AGN feedback has been called the **preventative** or **maintenance mode**. This effect is clearest around massive galaxies residing in the centres of clusters and groups. The hot gas in these massive halos should have relatively short cooling times ( $\lesssim 1$  Gyr) resulting in cooling flows of up to  $\approx 1000 M_{\odot} \text{ yr}^{-1}$  (Fabian, 1994). However, such cold gas flows has not been observed at the levels expected (e.g., O’Dea et al., 1998; David et al., 2001; Peterson et al., 2001, 2003) and the star formation of brightest cluster galaxies (BCGs) remains low. This ‘cooling-flow problem’ can be solved by invoking AGN feedback – powerful radio jets from the central AGN can pierce through their host galaxy and dissipate their kinetic energy in the halo. This heating effect of the AGN appears able to balance and offset the cooling of the halo (Fabian, 2012). Preventative feedback can also occur within the ISM of the galaxy itself, by heating or creating turbulence in the gas that prevents it from forming stars (e.g., Almeida & Nemmen, 2020; Mercedes-Feliz et al., 2023). Ionising radiation from the AGN can also directly heat and photo-ionise gas up to CGM scales (e.g., Arrigoni Battaia et al., 2019; Costa et al., 2022).

Another way the AGN can affect the host galaxy is through the **ejective mode** of feedback. This removes gas from the galaxy from the centre outward. As mentioned, there have been a range of observational studies showing outflows in AGN-hosted galaxies, with outflow rates that can exceed the star formation rate, suggesting that they are powerful enough to affect the long-term future of star formation within galaxies (e.g., Fiore et al., 2017; Bischetti et al., 2019). Other studies have also observed central gas depletion in AGN-hosting galaxies (e.g., Ellison et al., 2021) in support of this idea. This mode of feedback has also been invoked to explain the observed correlation between black hole mass and stellar velocity dispersion, as ejective wind models can produce relations with scalings in the range  $M_{\text{BH}} \propto \sigma_{\star}^{4-5}$  (Silk & Rees, 1998; King, 2003, 2005), which is remarkably close to the observed  $M_{\text{BH}} \propto \sigma_{\star}^{4.3}$  (Kormendy & Ho, 2013). Jet-launched outflows have also been shown to produce high mass outflow rates (e.g., Talbot et al., 2022). However, population-level

studies have not found evidence of depleted molecular gas or reduced star formation in AGN-hosted galaxies, as might be expected from an ejective mode of feedback (e.g., Stanley et al., 2015; Rosario et al., 2018; Shangguan et al., 2018; Koss et al., 2021; Zhuang et al., 2021). Additionally, how efficiently the outflow couples to the ISM to drive large outflows is unconstrained, with a wide range of values having been observed (e.g., Fiore et al., 2017; Bischetti et al., 2019; Lamperti et al., 2022) and theoretical expectations being limited by an unresolved ISM structure. These uncertainties will be the focus of Chapters 3 & 4 of this thesis.

A third form of feedback is **positive** feedback where the AGN activity increases the galaxy’s star formation rate. An AGN outflow can compress the ISM as it propagates, triggering star formation (e.g., Cresci et al. 2015a,b; Bessiere & Ramos Almeida 2022 although see Scholtz et al. 2020 for a counter-argument). There have even been studies showing star formation occurring within an outflow (Maiolino et al., 2017). However, as we have motivated in this chapter, the overall effect of AGN feedback is expected to be negative, so positive feedback is probably either rare, or limited to local effects. For example, Mercedes-Feliz et al. (2023) performed zoom-in simulations of a galaxy with and without an AGN wind and, although they found some local enhancement in star formation due to the outflow, the overall effect of the wind was negative due to gas ejection and preventative heating.

We note that some studies have termed the preventative mode the ‘jet/radio’ mode, and the ejective mode the ‘quasar/wind’ mode. However, this clear dichotomy has been challenged by recent studies finding that both jets and winds can launch outflows (Girdhar et al., 2022; Talbot et al., 2022); quasars can also host jets (Jarvis et al., 2019); winds can produce radio emission (Nims et al., 2015); and that winds can act preventatively (Almeida & Nemmen, 2020; Mercedes-Feliz et al., 2023). Therefore, when discussing AGN feedback, we avoid the terms ‘jet’ or ‘quasar’ mode, and focus our nomenclature on the *effect* the feedback is having on the galaxy.

### 1.3 Thesis Overview

AGN feedback has thus been established as an essential component of modern galaxy formation theory; both in semi-analytic models (e.g., Kauffmann & Haehnelt, 2000; Bower et al., 2006; Croton et al., 2006; Somerville et al., 2008) and cosmological hydrodynamic simulations (Springel et al., 2005a; Vogelsberger et al., 2014a; Hirschmann et al., 2014; Khandai et al., 2015; Schaye et al., 2015; Weinberger et al., 2018; Davé et al., 2019; Dubois et al., 2021; Wellons et al., 2023). It has been successfully invoked to produce realistic populations of quenched galaxies (Cattaneo et al., 2009; Sijacki et al., 2015; Terrazas et al., 2017; Akins et al., 2022), explain the self-regulation of black hole growth (Di Matteo et al., 2005; Sijacki et al., 2007; Habouzit et al., 2021), alter the structure and dynamics of galaxies (Dubois et al., 2013; Choi et al., 2018; van der Vlugt & Costa, 2019; Irodoutou et al., 2022), and solve the ‘cooling problem’ in galaxy clusters (Cattaneo et al., 2009; McCarthy et al., 2011).

However, as we have discussed throughout this chapter, there are still many uncertainties and open questions about these processes. In particular:

- **Why don’t we see a ‘smoking gun’ for AGN-driven quenching on a population level?** Observational studies show that AGN are preferentially located in gas-rich, star-forming galaxies (Mainieri et al., 2011; Rosario et al., 2018; Bischetti et al., 2019; Valentino et al., 2021; Ramos Almeida et al., 2022),

which some have suggested is in tension with the theory of negative AGN feedback (Trump et al., 2015; Shangguan et al., 2018; Schulze et al., 2019; Koss et al., 2021; Ji et al., 2022). In Chapter 3 we investigate this claim by applying similar tests to those performed by observers to three cosmological simulations, all of which utilise effective feedback models.

- **How do AGN winds couple to the ISM?** There are a plethora of observational studies that investigate how multiphase outflows correlate with AGN properties (e.g., Fiore et al. 2017; Bischetti et al. 2019; Lamperti et al. 2022; Ramos Almeida et al. 2022; Musiimenta et al. 2023, see Veilleux et al. 2020; Harrison & Ramos Almeida 2024 for reviews). However, poorly-constrained observational limitations and a lack of theoretical expectations make it challenging to interpret the results obtained. To help overcome this, in Chapter 4, we present a suite of idealised simulations of an AGN wind model embedded in a spatially-resolved, clumpy ISM to investigate how efficiently the wind couples to the gas in different ISM conditions and how this affects its ability to launch multi-phase outflows.
- **What are the observational signatures of wind-cloud interactions?** One of the hardest outflow phases to study is the hot phase, due to its low density. However, it is critical to understand this phase as it is expected to dominate the energy budget of the outflow (Ward et al., 2024). In Chapter 5, we use our clumpy disc simulations to predict the Bremsstrahlung emission expected from the wind-ISM interactions. We compare our results to analytic expectations in a homogeneous medium (Nims et al., 2015) and determine if this emission would be observable with the *Chandra* X-ray observatory.


These questions will be the focus of the rest of this thesis.

# 2

## Computational Methodologies

*I wrote it down in the winter of 1610.  
Just a secret under lock and key until then.  
While collecting the stars, I connected the dots.  
I don't know who I am, but now I know who I'm not.  
I'm just a curious speck that got caught up in orbit.  
Like a magnet it beckoned my metals toward it.*

*–Jupiter, Sleeping at Last*



**Figure 2.1:** Holmberg’s ‘simulation’ of a galaxy merger, performed using an arrangement of lamps, each with its own light detector (Holmberg, 1941). Despite the primitive setup, Holmberg was able to explain the formation of tidal arms which had thus far been an enigma.


## 2.1 A Brief History of Numerical Simulations

The first numerical ‘simulation’ in astrophysics wasn’t performed on a computer, but instead utilised lamps, photocells and manual calculation (Holmberg, 1941). Realising that both gravity and light intensity follow an inverse-square law, Holmberg’s innovation was to fit each lamp with a photocell to record the flux from every other lamp, allowing the acceleration felt by each particle to be calculated, and the lamp to be moved accordingly. This painstaking process had to be repeated for each of the 76 lamps for every time interval. Holmberg used this method to simulate the merger of two spiral galaxies and, despite the primitive nature of the setup, succeeded in studying effects such as tidal arms that arise from these interactions (Figure 2.1).

In the 1960s, the increased power and availability of computational facilities birthed the field of computational astrophysics. Early examples included less than 100 particles, but laid the groundwork for future studies (e.g., von Hoerner, 1960; Aarseth, 1963, 1966). In 1972, Toomre & Toomre (1972) revisited the question of tidal interactions, two decades after Holmberg’s pioneering work shuffling lamps around on the floor of his laboratory. The meteoric rise of computing power in the late 20<sup>th</sup> and early 21<sup>st</sup> centuries allowed rapidly increasing particle numbers to be used in these ‘N-body’ simulations (Figure 2.2). As computing power developed, so too did more efficient algorithms, such as mesh and tree methods, that improved the particle count further, culminating in the Millennium simulation, containing 10 billion particles (Springel et al., 2005b).

However, these N-body methods only include interactions due to gravity. This is a fair approximation for dark matter, which acts as a collisionless fluid due to its weak self-interaction, and dictates the formation of structure on cosmological scales (as introduced in Chapter 1). However, the formation and evolution of galaxies also depends heavily on the baryonic component. Baryons have a small mean free path compared to the length scale over which their thermodynamic properties vary, meaning they act as a collisional fluid and must be treated hydrodynamically, not just gravitationally.

There are two approaches used to investigate the effect of gas physics on galaxy evolution. The first is to use semi-analytic models to relate the properties of galaxies to the underlying dark matter structure; ‘painting-on’ the baryonic physics to the



**Figure 2.2:** The number of particles used by a selection of large-scale simulations in the last fifty years, showing the rapid increase in computational power and algorithm efficiency. Figure credit to Dr Florent Leclercq. Full references and reproducible code can be accessed via this [Github repository](#) or [blog post](#).

N-body simulation (e.g. [Somerville et al., 2008](#); [Henriques et al., 2015](#)) These use approximate prescriptions of mechanisms relevant to galaxy formation, such as gas accretion into halos, gas cooling and collapse to form stars, and feedback from star formation and black holes. The advantage of semi-analytic models is that they are relatively quick to run, allowing a large sample size of galaxies to be generated and a wide range of physical models tested (see [Benson 2010](#) for a review).

The second approach is to explicitly calculate the hydrodynamical evolution of the baryons. This is much more computationally expensive, but allows a self-consistent simulation of the formation and evolution of galaxies. This thesis will therefore focus on the full hydrodynamic method.

## 2.2 Hydrodynamics

### 2.2.1 Equations of Fluid Dynamics

By treating baryons as a collisional fluid, we can take the continuum assumption and compute their macroscopic evolution by using the fluid dynamic equations. We treat the baryons as an ideal fluid, and track their velocity ( $\vec{v}$ ), density ( $\rho$ ), pressure ( $P$ ) and internal energy ( $U$ ).

Firstly, the requirement for mass conservation gives us the continuity equation<sup>1</sup>:

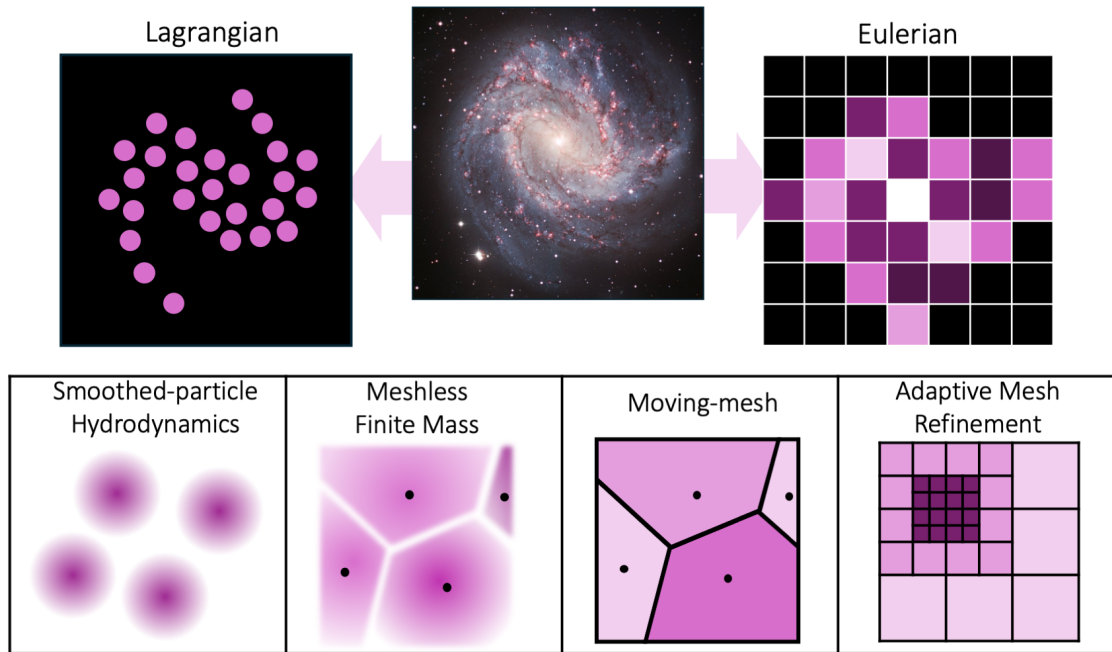
$$\frac{d\rho}{dt} + \rho \nabla \cdot \vec{v} = 0 \quad (2.1)$$

momentum conservation yields:

$$\frac{d\vec{v}}{dt} = -\frac{\nabla P}{\rho} - \nabla \Phi \quad (2.2)$$

<sup>1</sup>We have presented these equations in their Lagrangian form, using the material/substantial derivative:  $\frac{dQ}{dt} = \frac{\partial Q}{\partial t} + \vec{v} \cdot \nabla Q$ , where  $Q = Q(t, x, y, z)$  is some hydrodynamic quantity.





**Figure 2.3:** A schematic showing the two broad approaches to discretising astrophysical fluids: by splitting into mass elements (*Lagrangian*) or volume elements (*Eulerian*). In the bottom row, we show four example of approaches taken by modern hydrodynamic codes. The galaxy shown is M83, image credit: ESO.

and energy conservation requires:

$$\frac{dU}{dt} = -\frac{P}{\rho} \nabla \cdot \vec{v} - \frac{\Lambda(U, \rho)}{\rho} \quad (2.3)$$

Also included are source terms for the gravitational potential ( $\Phi$ ) and energy losses due to radiative cooling ( $\Lambda$ ). Finally, this set of equations can be closed by assuming an equation of state for an ideal, monoatomic ( $\gamma = 5/3$ ) gas:

$$P = (\gamma - 1) \rho U \quad (2.4)$$

Additionally, some studies extend these hydrodynamic equations with terms for magnetism (MHD; Pakmor & Springel 2013); radiation transport (RHD; Rosdahl et al. 2013); cosmic rays (CRs; Pfrommer et al. 2017); and, in environments with extreme gravity such as AGN accretion discs, general relativity (GRMHD; Liska et al. 2022). Full discussion of these is beyond the scope of this thesis.

### 2.2.2 Hydrodynamic Codes

To solve these equations numerically, the continuous fluid needs to be discretised and either set to move with the flow, or fixed in place. There are two main approaches for this which are shown in Figure 2.3: **Lagrangian** and **Eulerian**. In modern astrophysics there is a wide array of hydrodynamic codes available that take one of these two approaches, or attempt to combine their advantages. We will now briefly describe a few of the most popular approaches in numerical studies of galaxy formation.

In the **Lagrangian** formulation, the fluid is split into separate mass elements which are free to move with the fluid as it evolves. The elements act as interpolation points for the fluid calculations, with quantities at each point generally averaged over some number of neighbouring particles. The most popular family of Lagrangian


codes is **smoothed-particle hydrodynamics (SPH)**, where each mass element is smoothed over a smoothing kernel. Properties are obtained via a kernel-weighted sum over neighbouring particles within a smoothing length. Advantages of SPH are that it is fully Galilean invariant and exactly conserves mass, energy, entropy, and linear and angular momentum. Additionally, the spatial resolution is naturally adaptive, with more particles in denser regions and fewer in low density regions, which makes it particularly suitable for galaxy formation simulations where there is a large dynamic range of densities (although poor resolution in low-density regions can be a problem). Lagrangian codes struggle when dealing with shock fronts where entropy is not be conserved. To overcome this, artificial viscosity generally has to be added to achieve the necessary dissipation and to prevent particle interpenetration. However, this can suppress fluid instabilities and result in poor shock capturing, although there have been recent attempts to improve these problems, such as density-independent reformulations, or time-dependent artificial viscosity (for reviews of these methods see [Springel, 2010a](#); [Somerville & Davé, 2015](#); [Hopkins, 2015](#)). Some example SPH codes are GADGET and its descendants ([Springel, 2005](#)), which is used in the EAGLE cosmological simulations ([Crain et al., 2015](#); [Schaye et al., 2015](#)); and SWIFT ([Schaller et al., 2024](#)), which utilises the SPHENIX density-energy formulation ([Borrow et al., 2022](#)).

In an **Eulerian** formulation, the simulation domain is split into volume elements which remain stationary while the fluid flows through them. A [Godunov \(1959\)](#) scheme<sup>2</sup> is used to reconstruct the hydrodynamic properties at each cell boundary and then solve the exact or approximate 1D Riemann problem to give the flux across the interface. The main advantage of this scheme is better capturing of shocks and fluid instabilities than Lagrangian codes, without the need for artificial viscosity. However, this system is not Galilean invariant and does not explicitly conserve angular momentum.<sup>3</sup> The geometry of the mesh may imprint grid artefacts onto the result and quantities can become overmixed, suffering from numerical diffusion. Furthermore, on a uniform grid the spatial resolution is constant across the domain. For problems with a large dynamic range (e.g., galaxy evolution simulations), this results in poorer resolution in regions of high density compared to Lagrangian codes. One method to overcome this problem is **adaptive mesh refinement (AMR)** which superimposes a hierarchy of increasingly fine grids in regions of interest (usually in high-density regions). An example AMR code is RAMSES ([Teyssier, 2002](#)), used in the Horizon-AGN ([Dubois et al., 2014](#)) and NewHorizon ([Dubois et al., 2021](#)) simulations.

**Moving-mesh** codes attempt to combine some of the advantages of both Lagrangian and Eulerian formulations. They feature an unstructured mesh which moves with the fluid ('quasi-Lagrangian'). The mesh can refine or de-refine adaptively, usually with the aim of keeping each cell at a constant mass, improving the spatial resolution in high density regions. Quantities are advected across cell faces as in an Eulerian formulation, resulting in good handling of shocks and contact discontinuities. An example is the AREPO code ([Springel, 2010b](#); [Pakmor et al., 2016](#)) which is used in the Illustris ([Vogelsberger et al., 2014a](#)) and IllustrisTNG ([Springel et al., 2018](#)) cosmological simulations.

<sup>2</sup>Godunov's original scheme is only first-order accurate; however, there are many methods to extend it to higher-order accuracies (e.g., [van Leer, 1979](#)).

<sup>3</sup>The geometry and drift velocity of an Eulerian grid can be specified in such a way to conserve angular momentum (e.g., a radial grid), but this is only used in systems with a high degree of azimuthal symmetry. Most cosmological or galaxy evolution simulations use a Cartesian grid which does not guarantee angular momentum conservation.



**Figure 2.4:** The EAGLE cosmological simulation, demonstrating the vast dynamic range that can be achieved with these techniques. Slices are shown at decreasing spatial scales of 100 cMpc, 10 cMpc, and 60 ckpc. Figure taken from [Schaye et al. \(2015\)](#).

Finally, **meshless finite mass (MFM)** codes use a similar approach to moving-mesh by attempting to combine the two formulations. In MFM, the volume is partitioned with a continuous kernel which is integrated over when the Riemann problem is solved at the ‘effective’ faces between cells. An example is the GIZMO code ([Hopkins, 2015](#)) which is used in its MFM mode by the SIMBA cosmological simulations ([Davé et al., 2019](#)), and by the FIRE suite of zoom-in simulations ([Hopkins et al., 2014](#)).

In Chapter 3, we present an analysis of a sample of cosmological simulations which use a range of hydrodynamic solvers, as mentioned above. In Chapters 4 & 5, we present an idealised galaxy simulation which was run using the AREPO moving-mesh code.

## 2.3 Cosmological Simulations

Cosmological simulations<sup>4</sup> combine large box sizes – typically around  $(100 \text{ cMpc})^3$  – with a hydrodynamic treatment of baryonic physics, allowing them to trace the evolution of  $\sim 10^4$  galaxies across cosmic time. Figure 2.4 shows the results from one such simulation (EAGLE; [Schaye et al. 2015](#); [Crain et al. 2015](#)). We can see that, on large scales, the dark matter clusters together to form a cosmic ‘web’. In the densest nodes of this web, gas has cooled and collapsed to form galaxies, which can then host star formation (top-right inset of Figure 2.4). Over the last decade, there have been a range of these simulations, which require large, international teams

<sup>4</sup>Strictly speaking, the term ‘cosmological simulation’ also includes N-body simulations of large-scale structure (as discussed in Section 2.1). However, in this thesis, we generally use this term to refer to simulations that also include hydrodynamic processes.

to develop and analyse due to their complexity and computational intensity. Some examples include OWLS (Schaye et al., 2010); Illustris (Vogelsberger et al., 2014b,a); Horizon-AGN (Dubois et al., 2014); Magneticum (Hirschmann et al., 2014); EAGLE (Schaye et al., 2015; Crain et al., 2015); IllustrisTNG (Pillepich et al., 2017; Springel et al., 2018; Naiman et al., 2018; Nelson et al., 2018; Marinacci et al., 2018); SIMBA (Davé et al., 2019); and FLAMINGO (Schaye et al., 2023).

### 2.3.1 Subgrid Models

The addition of hydrodynamic calculations to cosmological-scale simulations adds significant computational cost. As Figure 2.2 shows, simulations featuring hydrodynamics (black squares) have 1 – 2 orders of magnitude fewer particles than modern N-body simulations. Additionally, the mass resolutions of these simulations is usually of the order  $10^{6-7} M_{\odot}$ , corresponding to spatial resolutions of around a kiloparsec within galaxies. This is significantly larger than the scale of many physical processes, such as star formation or SMBH accretion and feedback, that make important contributions to the evolution of the baryonic component in galaxies. Such mechanisms are therefore included in these simulations in the form of **subgrid models**, representing effects happening at scales below the resolution limit. These models use numerical prescriptions to mimic the emergent effect of various processes on scales that can be resolved by the simulations.

#### The Interstellar Medium and Star Formation

The interstellar medium (ISM) in galaxies is maintained by a host of processes such as (supersonic) turbulence, radiation pressure, thermal instabilities, and cosmic ray feedback that are impossible to replicate at the resolution of most galaxy formation simulations. Without this support, the ISM overcools and collapses.

To overcome this problem, simulations often use an ‘effective equation of state’ (eEOS) formalism (e.g., Springel & Hernquist, 2003) to account for these unresolved processes. In this model, the cold, dense gas is converted into stars which lowers the density of the ambient gas, reducing radiative losses and allowing its temperature to be increased by supernovae. This heating provides additional pressure support for the ISM, which acts against further star formation. This self-regulating cycle allows the temperature of the ISM to be approximated as a function of the density only, and parameterised by a single variable which is the efficiency of the conversion of mass into stellar particles. This can then be calibrated on observational constraints, such as the Kennicutt-Schmidt law ( $\Sigma_{\text{SFR}} \propto \Sigma_{\text{gas}}^{1.4}$ ; Kennicutt 1998). This approach has been used successfully in cosmological simulations for more than a decade (e.g., Vogelsberger et al., 2013).

In addition to an eEOS, simulations also include prescriptions for feedback from supernovae (SNe) or AGB star explosions (e.g., Rosdahl et al., 2017). This normally takes the form of energy injection, either directly at the site of star formation, or as hydrodynamically de-coupled particles that deposit their energy elsewhere (e.g., Springel & Hernquist, 2003; Smith et al., 2024). SNe events also pollute the surrounding gas with metals, to model the effect of stellar nucleosynthesis. The implementation of SF feedback has some similarities with the modelling of AGN which we now discuss.

### 2.3.2 Modelling Black Holes

A key aspect of subgrid modelling in cosmological simulations is their treatment of SMBHs. As discussed in Chapter 1.1.4, these simulations are unable to reproduce key observations without the implementation of feedback from SMBHs. There are three main aspects of modelling SMBHs in cosmological simulations: ‘seeding’ new black holes, accretion of gas onto the black hole, and feedback from the SMBH back into the surrounding galaxy.<sup>5</sup>

#### Seeding

To create a new SMBH, a sink particle of a set seed mass is spawned in the centre of a galaxy, usually once the halo has reached a certain mass. A higher choice for the minimum halo mass results in black hole being seeded later in cosmic time. Additionally, the mass of the black hole seed can be chosen to represent a stellar mass black hole (from a Population III star;  $M_{\text{BH}} \approx 10^{2-3}$ ) or a direct collapse scenario ( $M_{\text{BH}} \approx 10^5$ ). In the case that the black hole mass is below the mass resolution of the simulation, some models track the black hole mass independently from the dynamical mass of the sink particle. [Booth & Schaye \(2009\)](#) found that differing choices for minimum halo and seed masses had a much smaller contribution to the effect of AGN feedback on their host galaxies than other subgrid parameters such as accretion models. Additionally, the choice of these parameters mostly affects low-mass galaxies (see e.g., [Habouzit et al., 2021](#)) where AGN feedback is believed to be less effective<sup>6</sup>. Therefore, the discussion in this thesis will mostly neglect variations in seeding models.

#### Gas Accretion

Black holes grow their mass by accreting nearby gas. This accretion also powers the feedback effects such as winds and jets that are launched back into their host galaxy. However, the size of SMBH accretion discs ( $\approx 10^{-4}$  pc) is far below the spatial resolution of galaxy simulations, so this must also be modelled in a subgrid fashion.

Accretion models involve converting nearby gas mass into black hole mass. Some number of nearest-neighbour cells are selected to find the average density of gas near the black hole. Then, an accretion formula is used to calculate the amount of mass accreted onto the black hole, which is removed from the gas and added to the black hole particle mass. Most simulations use the Bondi-Hoyle-Lyttleton formula ([Bondi, 1952](#)):

$$\dot{M}_{\text{BH}} = \frac{4\pi\rho G^2 M_{\text{BH}}^2}{c_s^3} \quad (2.5)$$

where  $\rho$  and  $c_s$  are the density and sound speed of the surrounding gas respectively. However, this equation assumes spherical accretion and neglects angular momentum, self-gravity and radiative losses. Some simulations include additional equations to model different ‘modes’ of accretion (as discussed in Chapter 1.2.4). For example,

<sup>5</sup>Simulations also need to account for SMBH mergers and dynamics, which are challenging to model due to the small timesteps required to resolve these processes (for a contemporary approach, see [Mannerkoski et al. 2023](#)), but a full discussion is beyond the scope of this thesis.

<sup>6</sup>However, some recent simulations have suggested AGN feedback in dwarf galaxies may be important than previously thought ([Koudmani et al., 2021, 2022](#); [Sharma et al., 2023](#))

SIMBA uses the Bondi-Hoyle-Lyttleton equation for accreting hot gas and a torque-limited accretion model (Anglés-Alcázar et al., 2015, 2017) for the cold gas. Most cosmological simulations cap the accretion either at, or just above, the Eddington limit (e.g., Davé et al., 2019). Recent models have also attempted to add terms for super-Eddington accretion (e.g., Rennehan et al., 2024; Lupi et al., 2024).

However, a longstanding problem with accretion models in cosmological-scale simulations is that, due to poor spatial resolution, the nearest-neighbouring cells can be far away from the black hole. This has the effect of reducing the mean density of the nearby gas, resulting in a lowered accretion rate. To overcome this, some simulations apply a ‘boost factor’ to increase the accretion rate (e.g., Vogelsberger et al., 2013). Additionally, simulations are unable to resolve dynamical effects and structures near the black hole such as the fuelling of the central gas reservoir or small-scale shears due to interaction with a bar (e.g., Emsellem et al., 2015; Verwilghen et al., 2024). To resolve these issues, hyper-refinement schemes have been developed to improve the spatial resolution in regions near the black hole to sub-pc scales (Curtis & Sijacki, 2015; Anglés-Alcázar et al., 2021).

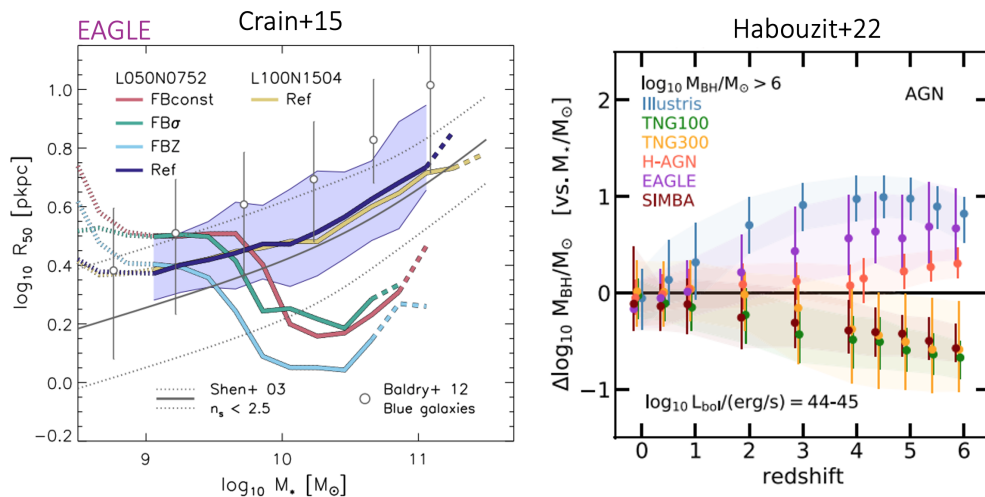
### AGN Feedback

Finally, cosmological simulations incorporate prescriptions for AGN feedback. There are a variety of approaches and, arguably, this is where simulations differ most in their subgrid modelling.

As the SMBH accretes mass from its surroundings, some of the liberated energy will be available for feedback processes. The amount of energy injected back into the galaxy is expressed as the feedback efficiency ( $\epsilon_f$ ; see Equation 1.7), which, in cosmological simulations, is a free parameter which can be tuned to adjust the strength of the feedback. This released energy can then take the form of thermal or kinetic energy injection. The energy can either be directly injected into the material surrounding the black hole, or decoupled and allowed to propagate into the galaxy before depositing its energy.

In **thermal** injection, the temperature of the neighbouring cells is increased by  $\Delta T$ . This loosely models radiative heating from the AGN, shock-heating from an AGN wind, or heating from radio jets; for example, in the Illustris simulation, thermal energy is injected directly into the galaxy halo to emulate observed radio bubbles Sijacki et al. (2007). A problem with thermal feedback is that, if the cooling time of the gas the energy is injected into is shorter than its sound-crossing time the energy will be quickly radiated away (see discussion in e.g., Dalla Vecchia & Schaye, 2008; Creasey et al., 2011). This is known as numerical overcooling. Some simulations, such as EAGLE compensate for this effect by ‘pulsing’ the feedback; storing the emitted energy until  $\Delta T$  is large enough to raise the temperature of the gas to the level where it won’t overcool.

The second method of energy injection is **kinetic**. In this system, neighbouring particles are given a kinetic ‘kick’ of momentum in a certain direction. This is used as an analogy to a kinetic wind or jet ejecting gas from the galaxy. In some simulations (e.g., TNG), this directional kick is in a random direction, leading to an isotropically-averaged outflow. In others, (e.g., SIMBA) the angular momentum of the black hole is tracked and the kinetic kicks are performed along the polar directions of the spin. As the inputted kinetic energy cannot be radiated away, unlike in the thermal case, this form of feedback can be more efficient.



**Figure 2.5:** The effect of subgrid models for SMBH feedback in cosmological simulations. *Left: an example of subgrid model calibration;* the half-light radius of disc galaxies in the EAGLE simulation is shown for different variations of the subgrid stellar feedback model. Despite all four models being calibrated to match the galaxy stellar mass function, only the reference model is able to reproduce galaxy sizes (Crain et al., 2015). *Right: observable differences between models;* black hole mass offset compared to the mean  $M_{\text{BH}} - M_{\star}$  relation at  $z = 0$ . The choice of subgrid models in the different simulations results in the black holes being over/under-massive at high redshifts. These differences could potentially be observable with JWST, providing a pathway to differentiate between the models used (Habouzit et al., 2022b).

In practice, cosmological simulations may use a combination of these methods and there is no clear consensus on the best way to model AGN feedback. Due to the coarse spatial resolution and the need for computational efficiency, these subgrid models are by necessity phenomenological in nature, rather than seeking to model the effect of AGN from first-principles. Some simulations seek to emulate the theorised split between ‘jet/maintenance’ and ‘quasar/wind/ejective’ modes of AGN feedback (see Chapter 1.2.7), whereas others, such as EAGLE, use just a single mode of feedback, reducing the number of free parameters. Despite these different philosophies, all these simulations are able to reproduce realistic populations of galaxies, demonstrating significant degeneracy between the models. In Chapter 3, we describe in more detail the subgrid models used for AGN feedback in three contemporary cosmological simulations: IllustrisTNG, EAGLE and SIMBA.

### 2.3.3 Interpretation Challenges

#### Calibration of subgrid models

The subgrid models for star formation and black hole feedback in cosmological simulations are phenomenological in nature and represent simplifications of the complex physical processes below the resolution limit. They are designed to have ideally only a few free parameters, which can then be calibrated on certain observables in order to reproduce realistic galaxy populations (see discussion in Crain & van de Voort, 2023). Most commonly, the galaxy stellar mass function or stellar-to-halo mass ratio is used as a calibrator for both SMBH and star formation feedback (e.g., Schaye et al., 2015). Additionally, IllustrisTNG adjusts the feedback model to better match the cosmic star formation rate density (Weinberger et al., 2017), EAGLE considers the sizes of produced disc galaxies (Crain et al., 2015), and SIMBA tunes the SMBH accretion efficiency to match the  $M_{\text{BH}} - M_{\star}$  relation (Davé et al., 2019).

Additionally, resolution plays a role in calibrating these models. As shown by [Bourne et al. \(2015\)](#), low resolution simulations result in a more efficient ejection of material as high density gas clumps, which are more resistant to feedback, are washed out in lower resolution. Therefore, the subgrid models are often recalibrated for each resolution explored.

Care must be taken, therefore, when analysing cosmological simulations to ensure that calibration diagnostics are not mistaken for observational predictions (see discussion in [Crain & van de Voort, 2023](#)). For example, the AGN feedback efficiency (see Section 1.2.3) in cosmological simulations is often around  $\epsilon_f \approx 10\%$ , which is set for calibration reasons outlined above. However, as discussed in [Harrison et al. \(2018\)](#), this has been misinterpreted as an observational prediction for the kinetic coupling efficiency of outflows, with observational studies then evaluating the impact of observed outflows based on whether they match this value. In Chapter 4, we discuss the kinetic coupling predicted from a physically motivated AGN model operating in an inhomogeneous ISM and compare it to observational samples.

### Comparing simulations

Although all cosmological simulations can reproduce generally realistic galaxy properties, studies have found differences between them. Finding variations between cosmological simulations, identifying the aspects of the subgrid modelling that lead to these differences, and linking these to testable, observational predictions is critical in improving our models for AGN feedback.

In a series of works, [Habouzit et al. \(2021, 2022a,b\)](#) compare a range of cosmological simulations, focusing on the  $M_{\text{BH}} - M_*$  relation. They find that differences at the low- $M_*$  end are mostly driven by variations in the SF feedback, SMBH seeding models, and numerical resolution, and that the high- $M_*$  end and the normalisation are most affected by the AGN feedback models (see also [Booth & Schaye 2010](#)). The right panel of Figure 2.5 shows whether the simulations predict over- or under-massive black hole at  $z \gtrsim 5$  which has important implications for observations using *JWST* ([Habouzit et al., 2022b](#)). They also present predictions for the quasar luminosity function and the constraints this places on forthcoming X-ray missions ([Habouzit et al., 2022a](#)). [Voit et al. \(2024\)](#) also investigated the  $M_{\text{BH}} - M_*$  in IllustrisTNG and EAGLE and found that the flatter relation at the high- $M_*$  end was caused by the switch to a kinetic mode of feedback in IllustrisTNG. Another regime where the simulations differ significantly is in the circumgalactic medium (CGM). [Davies et al. \(2020\)](#) compared the gas content in the CGM around galaxies in IllustrisTNG and EAGLE and found that low-mass halos were more gas-rich in IllustrisTNG than in EAGLE. This is due to the high-accretion rate feedback mode in IllustrisTNG being less efficient at expelling gas than EAGLE’s ‘burstier’ mode.

However, it is not always obvious to what extent the behaviour of these models is physically justified, or instead driven by numerical choices. This can make it confusing for observers seeking to test these models, as it is not clear whether a certain output is a genuine prediction of the model, or has been set for numerical reasons, such as calibration. Increasing the dialogue between observational and computational communities is therefore crucial in progressing our understanding.

In this thesis, we present a comparative study between cosmological simulations, as well as an idealised simulation with a physically motivated wind model. We also compare these to observational samples to evaluate common observational probes of AGN feedback. The challenges discussed in this section lead to two key questions:



- **Can we break the model degeneracy with observable predictions?** In Chapter 3, we outline the feedback models in TNG, EAGLE and SIMBA in more detail and study the effects they have on the star formation rate and molecular gas content of galaxy populations.
- **How can we improve our understanding of AGN feedback to more closely tie subgrid models to physical models?** In Chapters 4 & 5 we present idealised simulations & observable predictions for a physically-motivated wind model.

## 2.4 A Resolved Model for AGN Winds

As we have discussed, one of the problems with AGN feedback in cosmological simulations is the poor spatial resolution, leading to the impact of AGN being modelled phenomenologically, rather than from first-principles. However, there have been recent improvements in developing numerical implementations of more physical feedback models. In this section, we will briefly outline one such model for AGN winds and explain its implementation in the AREPO code.


### 2.4.1 Analytic Background

King (2003, 2005) introduced a model describing the interaction of sub-relativistic accretion disc wind with the ambient medium surrounding the black hole. This was motivated by X-ray observations of ultra-fast outflows (UFOs) in the nuclei of quasars (e.g., Pounds et al. 2003a,b; Tombesi et al. 2011; see Chapter 1.2.6 for more details). The theory behind this model was further developed by Zubovas & King (2012); Faucher-Giguère & Quataert (2012); Costa et al. (2014b); Costa et al. (2020), including the predicted thermal and non-thermal emission by Nims et al. (2015).

In this model, the accretion disc and initial launching mechanism are unresolved, but assumed to launch a fast, spherical outflow, moving at speeds  $v \approx (0.1 - 0.3)c$  into the surrounding medium, matching observations of UFOs. This is called the ‘free-expansion’ phase, or the ‘inner wind’. This supersonic inner wind drives a forward shock into the ambient ISM, sweeping up an outflowing shell. If this is decelerated, it drives a strong reverse shock back towards the black hole which thermalises a large fraction of the wind’s kinetic energy. Figure 2.6 shows the resulting structure in increasing distance from the black hole: the freely-expanding, unshocked, supersonic wind; the shocked wind; the shocked ambient medium being pushed in front of the outflow; and finally the undisturbed ambient medium beyond the outflow.

The evolution of the outflow is affected by how efficiently the shocked wind cools radiatively. This results in two limits, shown in Figure 2.6. In the case where the wind shock cools rapidly, we recover a **momentum-driven** outflow. The wind shock region loses energy due to radiative cooling and collapses into a thin shell. The freely expanding wind can then be thought to collide directly with the ISM, directly transfer its momentum to the outflow. However, if the wind shock is unable to radiate its energy away effectively, then an **energy-driven** outflow is produced. In this case, the outflow is driven by the adiabatic expansion of the hot wind, which may conserve some or all of the energy of the inner wind.

Despite having only relatively simple assumptions, this UFO-driven wind model is able to replicate the observed  $M_{\text{BH}}-\sigma_{\text{b}}$  relation, which is often taken as evidence for a co-evolving relationship between SMBHs and their host galaxies (Kormendy & Ho,



**Figure 2.6:** The structure of an AGN-wind driven outflow in a homogeneous medium. The freely-expanding wind sweeps up an outflowing shell of the shocked ambient medium. This drives a reverse shock back into the wind. If the shocked wind can cool efficiently, a momentum-driven solution is obtained (*left panel*), otherwise the outflow remains energy-driven (*right panel*). Figure modified from [Costa et al. \(2014b\)](#).

2013). In the momentum-driven limit, a relation of  $M_{\text{BH}} \propto \sigma_{\text{b}}^4$  is recovered, and in the energy-driven limit,  $M_{\text{BH}} \propto \sigma_{\text{b}}^5$  ([Costa et al., 2014b](#)). In reality, outflows are likely to be somewhere between these two limits, making these predictions remarkably close to the observed relation of  $M_{\text{BH}} \propto \sigma_{\text{b}}^{4.4}$ .

We can quantify which regime we are in by calculating the momentum ‘boost’ above the initial wind input. If the small-scale momentum of the wind is given by

$$\dot{p}_{\text{w}} = \frac{L_{\text{AGN}}}{c} \quad (2.6)$$


then in the case of energy-conservation ( $\dot{M}_{\text{OF}} v_{\text{OF}}^2 = \dot{M}_{\text{w}} v_{\text{w}}^2$ ), we would expect a momentum boost of

$$\frac{\dot{p}_{\text{OF}}}{L_{\text{AGN}}/c} = \frac{v_{\text{w}}}{v_{\text{OF}}} \quad (2.7)$$

whereas in the momentum-conserving case ( $\dot{M}_{\text{OF}} v_{\text{OF}} = \dot{M}_{\text{w}} v_{\text{w}}$ ), we would only expect:

$$\frac{\dot{p}_{\text{OF}}}{L_{\text{AGN}}/c} = 1 \quad (2.8)$$

In practice, the maximally energy-conserving case won’t be achieved due to the outflow having to do work against gravity and the ambient gas pressure. Analytic models have recovered momentum boosts as large as  $\dot{p}_{\text{OF}} \approx 20(L_{\text{AGN}}/c)$  ([King, 2003](#); [Zubovas & King, 2012](#); [Faucher-Giguère & Quataert, 2012](#)) for winds expanding in an isotropic medium. Inspired by these works, several observational studies have attempted to quantify the momentum boost of observed outflows as a diagnostic of whether they are energy- or momentum-driven (e.g., [Cicone et al., 2014](#); [Carniani et al., 2015](#); [Fiore et al., 2017](#); [Musumanta et al., 2023](#)). However, there are several observational challenges when calculating this parameter, such as poorly constrained inclination angles, conversion factors, and only being able to capture one phase of the gas, depending on which observational tracer is used (see discussion in [Harrison & Ramos Almeida 2024](#)).



**Figure 2.7:** The implementation of this analytic wind model in the AREPO code (Costa et al., 2020). *Left panel:* the resolved wind structure, showing the freely-expanding wind (1); shocked wind (2); shocked ambient medium (3); and the undisturbed ambient medium. The lower-right subpanel shows the injected wind tracer fluid, revealing a clear contact discontinuity between the shocked wind and the shocked ambient medium. *Middle panel:* cell structure in AREPO. Two concentric spheres of cells are fixed in place and the wind properties are set at the interface between them, allowing the wind to be injected across the cell boundaries (*right panel*).

Additionally, a major uncertainty is what effect a realistic, inhomogeneous ISM structure would have on the energetics of the outflow. Analytic work assumes a smooth ambient medium, and simulations generally do not have the spatial resolution to capture the substructure of the ISM on parsec scales. This could cause the evolution of the outflow to diverge significantly from the shell-like morphology assumed in most observational works. To investigate this, in Chapter 4 we calculate the momentum boost for an AGN wind embedded in a clumpy disc and discuss the observational implications.

### 2.4.2 Implementation in AREPO

Encouraged by the successes of this analytic wind model, Costa et al. (2020) devised a numerical scheme to implement it into the AREPO code – the BOLA model (BOundary Layer for AGN; Figure 2.7). Two concentric spheres of cells are centred on the position of the black hole to create a boundary layer. These cells are fixed in place and not allowed to refine or derefine. The mass, momentum and energy flux of the wind are set at the interface between these two layers and the wind geometry can be carefully set by letting the fluxes vary with solid angle. The fluxes are set at the boundary as predicted by analytic models or UFO observations. This makes it possible to stitch together the resolved large-scale outflow with physics operating below the resolution scale. The properties of the wind can be parameterised by just three variables: the luminosity of the AGN ( $L_{\text{AGN}}$ ), the wind speed ( $\beta = v_w/c$ ) and the momentum flux (usually taken to be  $\tau = 1$ ), giving equations for the mass outflow rate:

$$\dot{M}_w = \frac{\tau L_{\text{AGN}}}{\beta c^2} \quad (2.9)$$

the momentum rate:

$$\dot{p}_w = \tau \frac{L_{\text{AGN}}}{c} \quad (2.10)$$

and the energy rate:

$$\dot{E}_w = \frac{\tau\beta}{2} L_{\text{AGN}} \quad (2.11)$$

The initial temperature of the wind can also be set, but, as the wind is quickly shocked-heated to high temperatures and the flow is highly supersonic, this choice makes a negligible difference to the outflow. Alongside the wind parameters, a passive, conserved scalar is also injected to track the advected mass of the wind fluid. This ‘wind tracer’ can be seen in the left-hand panel of Figure 2.7, clearly highlighting the contact discontinuity between the shocked wind and the shocked ambient medium.

Using this model, the thermalisation of the inner wind is explicitly resolved, rather than just being emulated by a thermal energy dump, as in the phenomenological models used in cosmological simulation. This therefore represents a significant increase in realism when modelling the effect of AGN wind feedback. In Chapter 4 we investigate how an outflow launched from this wind model is affected by an inhomogeneous, clumpy ambient medium, and in Chapter 5 we explore the expected X-ray emission from these wind-cloud interactions.

3

# AGN Feedback in Cosmological Simulations

Cosmological simulations predict that AGN preferentially live in gas-rich, star-forming galaxies despite effective feedback

Samuel Ruthven Ward; C.M.Harrison; T.Costa; V.Mainieri  
MNRAS, 514 (2022)

---

Negative feedback from active galactic nuclei (AGN) is the leading mechanism for the quenching of massive galaxies in the vast majority of modern galaxy evolution models. However, direct observational evidence that AGN feedback causes quenching on a population scale is lacking. Studies have shown that luminous AGN are preferentially located in gas-rich and star-forming galaxies, an observation that has sometimes been suggested to be in tension with a negative AGN feedback picture. We investigate three of the current cosmological simulations (ILLUSTRISTNG, EAGLE and SIMBA) along with post-processed models for molecular hydrogen gas masses and perform similar tests to those used by observers. We find that the simulations predict: (i) no strong negative trends between  $L_{\text{bol}}$  and  $f_{H_2}$  or sSFR; (ii) both high-luminosity ( $L_{\text{bol}} \geq 10^{44}$  erg s $^{-1}$ ) and high-Eddington ratio ( $\lambda_{\text{Edd}} \geq 1\%$ ) AGN are preferentially located in galaxies with high molecular gas fractions and sSFR; and (iii) that the gas-depleted and quenched fractions of AGN host galaxies are lower than a control sample of non-active galaxies. These three findings are in qualitative agreement with observational samples at  $z = 0$  and  $z = 2$  and show that such results are not in tension with the presence of strong AGN feedback, which all simulations we employ require to produce realistic massive galaxies. However, we also find quantifiable differences between predictions from the simulations, which could allow us to observationally test the different subgrid feedback models.

---

### 3.1 Introduction

The population of galaxies in our universe exhibits a bimodal distribution, split into star-forming (or ‘blue cloud’) galaxies, which have high levels of star formation and blue colours due to young massive stars; and quiescent (or ‘red sequence’) galaxies with low star formation, and red colours from their older stellar population (e.g. [Baldry et al., 2004](#); [Schawinski et al., 2014](#); [Bluck et al., 2020](#)). To create the population of quiescent galaxies, some process is required to ‘quench’ star formation. For star formation to occur, gas needs to cool, forming molecular hydrogen ( $H_2$ ) clouds, and collapse gravitationally. Thus, any mechanism that acts to quench a galaxy must remove the gas, prevent it from cooling, stabilise it against collapse, or destroy the densest gas phase. A strong candidate for this mechanism is the energy released by the supermassive black holes (SMBHs) that are located at the centres of all massive galaxies ([Kormendy & Ho, 2013](#)). If these SMBHs have a readily available fuel supply of cold gas, they can grow rapidly through gas accretion and light up to become ‘Active Galactic Nuclei’ (AGN). Such AGN are able to release vast amounts of energy ( $L_{\text{bol}} \simeq 10^{42-48} \text{ erg s}^{-1}$ ) and even if only a small fraction of this power can couple to the gas in a galaxy, it has the potential to unbind gas, prevent cooling and have a significant influence on the evolution of their host galaxies through a process known as ‘AGN feedback’ ([Fabian, 2012](#)).

Indeed, AGN feedback is essential to our current theoretical understanding of galaxy evolution and is deeply embedded in cosmological models and simulations (e.g. [Somerville et al., 2008](#); [Schaye et al., 2015](#); [Khandai et al., 2015](#); [Dubois et al., 2016](#); [Weinberger et al., 2018](#); [Davé et al., 2019](#)). This process is necessary in the models to reproduce the observed galaxy bimodality; without it, purely environmental effects, or self-regulation from star formation, are not enough to sufficiently suppress star formation in the most massive galaxies ([Bower et al., 2006](#); [McCarthy et al., 2011](#); [Beckmann et al., 2017](#)). Furthermore, AGN feedback is required in cosmological models and simulations to explain a variety of other observational results, such as a solution to the ‘cooling problem’ in galaxy clusters, reproducing observed galaxy sizes, and it could be critical in determining galaxy structures and dynamics ([Sijacki et al., 2007](#); [Cattaneo et al., 2009](#); [Choi et al., 2018](#); [van der Vlugt & Costa, 2019](#); [Irodou et al., 2022](#)).

Despite the success of AGN feedback, its numerical implementation varies considerably across cosmological simulations and there is no clear consensus as to the most realistic model ([Crain et al., 2015](#); [Weinberger et al., 2018](#); [Davé et al., 2019](#); [Costa et al., 2020](#), for more details see Section 3.2.1). Observational tests are crucial to refine or rule out different models, as well as establish whether AGN feedback is effective at quenching galaxies, as theoretical models have long suggested ([Silk & Rees, 1998](#); [King, 2003](#)).

From the observational perspective, many studies have shown that luminous AGN ( $L_{\text{bol}} \gtrsim 10^{44} \text{ erg s}^{-1}$ ) can drive large-scale multi-phase outflows and inject turbulence into the host galaxy’s interstellar medium (ISM; [Sturm et al., 2011](#); [Cicone et al., 2018](#); [Baron et al., 2018](#); [Baron & Netzer, 2019](#); [Veilleux et al., 2020](#); [Girdhar et al., 2022](#)). However, this evidence does not directly establish that luminous AGN have a significant, and lasting, impact on the global star formation or molecular gas content of the host galaxies ([Harrison, 2017](#)).

There has been a large array of observational studies over the last decade measuring the star formation and molecular gas properties of luminous AGN populations to search for evidence of an impact of AGN feedback. A variety of approaches have been taken; for example: (1) investigating the trends of star formation rates or  $H_2$  content, with the luminosity of the AGN (e.g. [Page et al., 2012](#); [Harrison et al., 2012](#); [Rosario et al., 2013](#); [Stanley et al., 2015](#); [Kakkad et al., 2017](#); [Shangguan & Ho, 2019](#); [Jarvis et al., 2020](#); [Zhuang et al., 2021](#); [Kim et al., 2022](#)); and (2) comparing the star formation rates and molecular

gas properties of AGN host galaxies to similar galaxies without an AGN, either by looking at averages properties, or by comparing distributions (e.g. [Bernhard et al., 2016](#); [Rosario et al., 2018](#); [Scholtz et al., 2018](#); [Schulze et al., 2019](#); [Kirkpatrick et al., 2019](#); [Florez et al., 2020](#); [Circosta et al., 2021](#); [Bischetti et al., 2021](#); [Valentino et al., 2021](#); [Scholtz et al., 2021](#)).

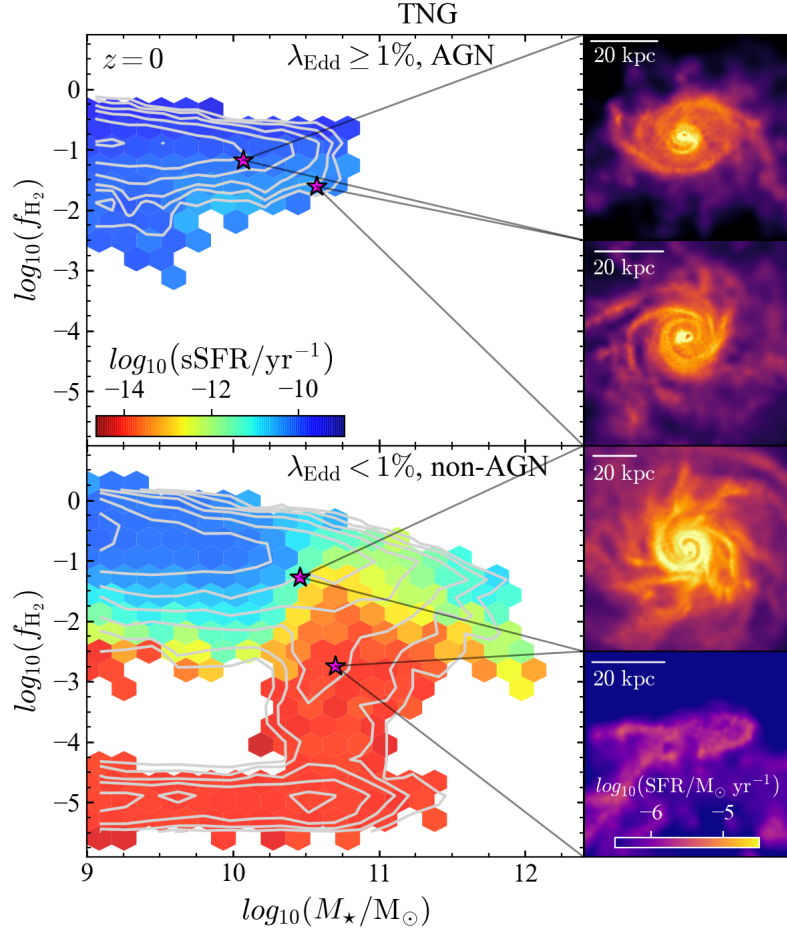
However, the observational studies taking these types of approaches have not provided definite evidence for strong AGN feedback. Across most studies, trends of star formation rates and molecular gas content with  $L_{\text{bol}}$  are either found to be flat or positive ([Lutz et al., 2010](#); [Mullaney et al., 2012](#); [Stanley et al., 2015](#); [Azadi et al., 2015](#); [Gürkan et al., 2015](#); [Shimizu et al., 2017](#); [Scholtz et al., 2018](#); [Kirkpatrick et al., 2019](#); [Ramasawmy et al., 2019](#); [Zhuang et al., 2021](#)) with some studies finding a positive correlation only in the most luminous ( $L_{\text{bol}} \gtrsim 10^{45} \text{ erg s}^{-1}$ ), although this may be due to underlying mass trends ([Rosario et al., 2012](#); [Stanley et al., 2017](#)). Additionally, once matched for galaxy morphology and stellar mass, AGN and non-AGN hosts are typically found to have similar levels of star formation and molecular gas ([Mainieri et al., 2011](#); [Husemann et al., 2017](#); [Stanley et al., 2017](#); [Rosario et al., 2018](#); [Shangguan et al., 2018](#); [Smirnova-Pinchukova et al., 2022](#); [Valentino et al., 2021](#); [Ramos Almeida et al., 2022](#)).

On the other hand, there are a few studies of high redshift AGN,  $z \gtrsim 1$ , that suggest there is reduced molecular gas in AGN hosts (e.g. [Kakkad et al., 2017](#); [Perna et al., 2018](#); [Circosta et al., 2021](#); [Bischetti et al., 2021](#)), and a possible negative connection between AGN-driven outflows and star formation rates at lower redshifts ([Wylezalek & Zakamska, 2016](#); [Chen et al., 2022](#)). The situation is further complicated by the use of different AGN selection methods (e.g. X-ray, IR or radio) which can bias the sample of galaxies taken, affecting the final results for the distributions of host galaxy properties (e.g. [Azadi et al., 2017](#); [Harrison, 2017](#); [Ji et al., 2022](#), and references within.) Nonetheless, the broad consensus is that luminous AGN are preferentially observed in gas-rich and highly star-forming galaxies (e.g. [Rosario et al., 2013](#); [Bernhard et al., 2016](#); [Jarvis et al., 2020](#); [Florez et al., 2020](#); [Xie et al., 2021](#); [Koss et al., 2021](#); [Vietri et al., 2022](#)). This apparent lack of unanimous evidence for reduced star formation or molecular gas in luminous AGN has led some recent studies to speculate that the evidence for feedback by luminous AGN is weak (e.g. [Trump et al., 2015](#); [Shangguan et al., 2018](#); [Ramasawmy et al., 2019](#); [Schulze et al., 2019](#); [Shangguan et al., 2020b](#); [Koss et al., 2021](#); [Valentino et al., 2021](#); [Ji et al., 2022](#)).

Despite the aforementioned conclusion of some studies, observationally connecting the observed AGN luminosity with the long-term impact of the AGN on the host galaxy molecular gas and star formation is a complex process. The high variability of AGN luminosity, the potential time delay between AGN activity and its effect, and the shared cold gas reservoir that fuels both the AGN and the star formation make it difficult to predict what trends, if any, would be seen in active galaxy populations (e.g. [Hickox et al., 2014](#); [Harrison, 2017](#); [Luo et al., 2021](#)). Therefore, a potentially better way to test proposed AGN feedback models is to extract predictions of galaxy properties from cosmological simulations directly.

In this study, we analyse three cosmological simulations using similar approaches to those taken by observers, described above, to establish the predicted relationships between AGN luminosity and star formation rates and molecular gas content. This approach enables us to: (1) determine if the observational results are indeed in tension with simulations where strong AGN feedback is present and (2) investigate how the different models of AGN feedback across the simulations result in different observational predictions. We will focus on two cosmological epochs in line with where observers have focussed their efforts: local galaxies at  $z \simeq 0$  (for which the data are easier to obtain) and cosmic noon,  $z \simeq 2$ , when star formation and AGN activity peaked ([Madau & Dickinson, 2014](#)).





**Figure 3.1:** An example presentation of the simulation data we use in this work. The logarithmic number density contours of galaxies is shown in the  $f_{H_2} - M_*$  plane for TNG at  $z = 0$ . Pixel bins are coloured by the mean sSFR of the galaxies enclosed. The upper panel shows an  $\lambda_{\text{Edd}}$ -selected AGN sample and the lower panel shows galaxies without an AGN. Star formation rate maps of four example galaxies are visualised in the right-hand columns. From the left panels we can see that AGN are hosted by galaxies with high  $f_{H_2}$  and sSFR. The images also reveal that AGN host galaxies have similar morphology to other star-forming galaxies that do not host AGN.

This paper is organised as follows: in Section 3.2 we explain our methodology and in Section 3.3 we present our analysis of the simulations following similar approaches to observational papers, namely looking at trends between galaxy properties and AGN luminosity (Section 3.3.1), comparing AGN host galaxies to inactive galaxies on the  $f_{H_2} - M_*$  plane (Section 3.3.2), investigating the gas fraction distribution for the highest  $L_{\text{bol}}$  systems (Section 3.3.3), and comparing the gas-depleted and quenched fractions of AGN host galaxies with a mass-matched non-AGN sample (Section 3.3.4). We discuss the implications of our results in the context of the observations in Section 3.4 and offer suggestions for how to improve both observational and simulation-based studies, before presenting our conclusions in Section 3.5.

In this work, we assume a flat,  $\Lambda$ CDM cosmology, using values from the [Planck Collaboration \(2016\)](#):  $H_0 = 67.7 \text{ kms}^{-1} \text{ Mpc}^{-1}$ ,  $\Omega_m = 0.3$ ,  $\Omega_\Lambda = 0.7$ ; in line with the cosmologies assumed by the three simulations we will study.

## 3.2 Methods

In this section we present our methods, including a brief introduction to the three cosmological simulations and the observational data used (Sections 3.2.1 & 3.2.2), the target galaxy quantities we will be considering (Section 3.2.3), our AGN-selection and quenching definitions (Sections 3.2.4 & 3.2.5) and finally our methods for calculating correlation coefficients (Section 3.2.6).

### 3.2.1 Simulations

In this study, we have selected three of the current generation of hydrodynamic, cosmological simulations: IllustrisTNG (Springel et al., 2018; Pillepich et al., 2018; Nelson et al., 2018; Marinacci et al., 2018; Naiman et al., 2018), EAGLE (Crain et al., 2015; Schaye et al., 2015) and SIMBA (Davé et al., 2019). These have comparable box sizes of  $L \simeq 100$  cMpc (comoving Mpc) and trace dark matter, gas, stellar masses and SMBH properties, allowing us to search for trends between SMBH accretion rates, gas fractions, star formation and stellar masses within the context of a large host galaxy population ( $\sim 10^4$  galaxies at  $z = 0$ ). As an illustration of the type of data available, Figure 4.2 shows the galaxy population in IllustrisTNG at  $z = 0$ , displaying specific star formation rates as a function of molecular gas fraction and stellar mass for Eddington-ratio selected AGN and non-AGN samples (see Section 3.2.4) as well as visualisations of the star formation rate distribution in four representative galaxies from these populations. These example galaxies give us a first hint about our main finding: bright AGN tend to reside in extended, star-forming discs (as seen in the first two panels on the right-hand side).

These simulations all model various crucial physical processes, including gas cooling, star formation, SMBH accretion, stellar feedback and AGN feedback. However, due to resolution limitations, these processes are not resolved from first principles, but rather modelled at a subgrid level. Of particular interest to this study are the prescriptions for AGN feedback included in each of the simulations. There are significant differences in how each simulation approaches feedback, which we summarise in Table 3.1, show in Figure 3.2, and explain in more detail below.

To study the impact of AGN on the host galaxy, many observational studies estimate  $H_2$  masses, using a tracer such as CO, and we would therefore like to test these observational results against the simulations' findings. However, due to computational limitations, the ISM is poorly resolved and the molecular gas phase is not directly traced in the simulations. Instead, we can estimate the molecular hydrogen by using 'post-processed' models such as the one presented in Gnedin & Kravtsov (2011) who ran a set of high-resolution zoom-in simulations which followed the detailed chemical evolution of the gas and used these to derive fitting functions for the  $H_2$  which can then be applied to cosmological simulations (a discussion of alternative post-processed models can be found in Appendix A of Lagos et al. 2015 and in Diemer et al. 2018). These models have been applied in a similar way across the three simulations we consider in this work (Lagos et al., 2015; Diemer et al., 2018, 2019; Davé et al., 2019), allowing us to consistently compare predictions for molecular hydrogen fractions between the simulations. Furthermore, the predictions from these models have been tested against observed galaxy populations and have been found to be in good agreement at  $z = 0$  (Lagos et al., 2015; Diemer et al., 2019). At  $z = 2$  there is a small discrepancy with the simulations predicting a lower molecular gas fraction than the observational constraints by a factor of 1.5 in EAGLE (Lagos et al., 2015) and a factor of 2 – 3 in IllustrisTNG (Popping et al., 2019). SIMBA also likely shows such a discrepancy but it has not been extensively discussed in the literature. We describe our correction for this in Section 3.2.5.

To ensure all galaxies in our simulation sample are sufficiently resolved, we take a stellar mass cut of  $M_{\star} \geq 10^9 M_{\odot}$ . This is raised slightly to  $M_{\star} \geq 10^{9.5} M_{\odot}$  in SIMBA as black holes are only seeded once a galaxy has reached this size. We keep the lower cut for TNG and EAGLE to increase the sample size, although we note that raising this mass cut yielded no qualitative differences to our results.

We now provide an overview of the three simulations, focusing specifically on the implementation of SMBH accretion and AGN feedback, and their molecular gas estimates. For more details, the reader is referred to the release papers listed.

### IllustrisTNG

IllustrisTNG<sup>1</sup> (Springel et al. 2018; Pillepich et al. 2018; Nelson et al. 2018; Marinacci et al. 2018; Naiman et al. 2018, data release: Nelson et al. 2019a, hereafter TNG) is a suite of simulations building on the original Illustris simulation (Vogelsberger et al., 2014a,b) and using the moving-mesh code AREPO (Springel, 2010b). TNG has been run with three box sizes with approximate side lengths of 50, 100 and 300 cMpc. In this work we use the mid-sized run, TNG100 (side length  $L = 110.7$  cMpc; baryonic mass resolution  $m_b = 1.39 \times 10^6 M_{\odot}$ ), as it provides the most direct comparison to EAGLE and SIMBA, although we also verify the convergence of our results by comparing them against the TNG300 run.

The calculated values for molecular gas masses are taken from the molecular and atomic hydrogen post-processed catalogues from Diemer et al. (2018, 2019) using the prescription by Gnedin & Kravtsov 2011.

A key feature of the simulations is the black hole accretion model and feedback processes (Weinberger et al., 2017, 2018). The SMBH accretion rate is set according to the Bondi-Hoyle model (Bondi, 1952) by performing a kernel-weighted average over nearest-neighbour gas cells around each SMBH particle. The accretion rate is also capped at the Eddington limit.

TNG features two modes of AGN feedback at low and high accretion rates which occur exclusively from each other. The mode a given SMBH is in at a particular time, is determined by its current Eddington ratio,  $\lambda_{\text{Edd}}$ , which is compared to a critical value,  $\chi$ , given by:

$$\chi(M_{\text{BH}}) = \min \left[ 0.002 \left( \frac{M_{\text{BH}}}{10^8 M_{\odot}} \right)^2, 0.1 \right], \quad (3.1)$$

For  $\lambda_{\text{Edd}} \geq \chi$ , the SMBH is in the high-accretion mode and for  $\lambda_{\text{Edd}} < \chi$ , the SMBH is in the low-accretion regime. This split between the two modes can be seen in Figure 3.2, in the accretion rate (scaled to bolometric luminosity; Equation 3.5) versus black hole mass plane.


The two modes are associated to two different feedback models, summarised here and also shown in Table 3.1:

- *Thermal mode*: in this high-accretion mode, feedback energy is continually injected isotropically into the local environment in thermal form, heating the surrounding gas cells. SMBHs in this mode lie in the top left of Figure 3.2 and generally have low masses and high accretion rates. This mode is designed to reflect a ‘wind’ or ‘quasar’ mode (e.g. Sijacki et al., 2007).
- *Kinetic mode*: the energy is injected as momentum into the neighbouring gas cells of the black hole, directed in a random orientation with a  $0^\circ$  opening angle, although this averages to an isotropic distribution after many episodes. A minimum energy

<sup>1</sup><https://www.tng-project.org>

Simulation	Mode Name	Criteria	Injection Energy Type	Direction	Description
TNG	Thermal	$\lambda_{\text{Edd}} \geq \chi(M_{\text{BH}})$	Thermal	Isotropic	Continuous thermal energy injection into small local environment
	Kinetic	$\lambda_{\text{Edd}} < \chi(M_{\text{BH}})$	Kinetic	Random (averages isotropic)	Pulsed momentum kick in random direction
EAGLE	Thermal (single mode)	Always active	Thermal	Isotropic	Pulsed thermal injection
SIMBA	Wind	Always active	Kinetic	Bipolar	Kinetic kick, $v \leq 1000 \text{ km s}^{-1}$
	Jet	$\lambda_{\text{Edd}} < 0.2$ $M_{\text{BH}} > 10^{7.5} M_{\odot}$	Kinetic (& few % thermal)	Bipolar	Kinetic kick, $v \leq 7000 \text{ km s}^{-1}$ , temperature raised to $T_{\text{halo}}$
SIMBA	X-ray	$\lambda_{\text{Edd}} < 0.02$ $f_{\text{gas}} < 0.2$	Thermal (non-ISM gas) or thermal & kinetic (ISM gas)	Isotropic	Local thermal heating

**Table 3.1:** A summary of the subgrid implementations for AGN feedback in the three simulations. We note that, for simplicity, we neglect the X-ray mode in SIMBA as it has negligible impact on the initial quenching of galaxies (Davé et al., 2019). Refer to Section 3.2.1 for the definitions of  $\chi(M_{\text{BH}})$  and  $\lambda_{\text{Edd}}$ .



**Figure 3.2:** Bolometric luminosity (calculated from SMBH accretion rate; Equation 3.5) versus black hole mass for the simulations. The dotted lines in TNG and SIMBA show the boundaries between the different feedback modes that can be active. We also show our two AGN definitions (Section 3.2.4): the solid black line shows a constant Eddington ratio of  $\lambda_{\text{Edd}} = 1\%$  and the blue dashed line shows a high luminosity cut of  $L_{\text{bol}} = 10^{44} \text{ erg s}^{-1}$ .

for each kick is required, so the energy injected is stored up until this threshold is reached and then released in a pulse. These sources lie in the bottom right of Figure 3.2 and have high masses and low accretion rates (Croton et al., 2006; Bower et al., 2006).

Additionally, radiative feedback from the AGN is modelled by modifying the cooling function of gas cells in the vicinity of the SMBH which acts to suppress cooling (Vogelsberger et al., 2013; Pillepich et al., 2017). However, this channel is only significant for sources with high Eddington ratios and is thus concurrent only with the thermal mode – for simplicity we will not distinguish it from the thermal mode in this study.

## EAGLE

EAGLE<sup>2</sup> (Schaye et al. 2015; Crain et al. 2015, data release: McAlpine et al. 2016) is a suite of cosmological simulations run with the smoothed particle hydrodynamics (SPH) code GADGET-3 (Springel, 2005). In this work, we use the largest box-size reference run, Ref-L100N1504, which has a side length of  $L = 100 \text{ cMpc}$  and a baryonic mass resolution of  $m_b = 1.81 \times 10^6 M_\odot$ .

Molecular gas masses are calculated in post-processing (Lagos et al., 2015) by first calculating the neutral gas fraction using the method of Rahmati et al. (2013) and then following the prescriptions of Gnedin & Kravtsov (2011) in a similar way to TNG. Previous results have noted that the SFRs in EAGLE are around 0.2 dex too low compared to observations (Furlong et al., 2015; McAlpine et al., 2017) so, following these studies, we also scale the SFR up by this value.

In EAGLE, SMBH growth is also modelled via Bondi-Hoyle accretion, modified by a factor of the ratio between the Bondi and viscous time-scales (Rosas-Guevara et al., 2016), and is capped at the Eddington limit. The coupling of the released energy from the black hole to its nearest neighbour particles is achieved using a single mode of feedback (Booth & Schaye, 2009) that operates at any Eddington ratio, in contrast to the dual modes of TNG and SIMBA. Feedback energy is stored until it is sufficient to heat the surrounding particles by  $\Delta T = 10^{7.5} \text{ K}$  and then stochastically injected as thermal energy. This ‘pulsed’ nature of the thermal feedback prevents the energy being immediately radiated away and offsets

<sup>2</sup><http://icc.dur.ac.uk/Eagle/index.php>;  
<https://eagle.strw.leidenuniv.nl>

cooling (Booth & Schaye, 2009). This makes it more efficient at quenching the galaxy than the corresponding thermal mode in TNG which is how EAGLE can successfully reproduce the galaxy stellar mass function without utilising a kinetic feedback mechanism (Schaye et al., 2015).

## SIMBA

SIMBA<sup>3</sup> (Davé et al., 2019) is based on the earlier MUFASA simulation (Davé et al., 2016) and features updated physics for the modelling of black hole growth and feedback. It is run using finite mass hydrodynamics from GIZMO (Hopkins, 2015). In this paper, we will use the fiducial SIMBA run (m100n1024) which has a box side length of  $L = 147$  cMpc and a baryonic mass resolution of  $m_b = 1.83 \times 10^7 M_\odot$ .

Unlike TNG and EAGLE, the molecular hydrogen content of the gas is calculated on-the-fly at each timestep (using the prescription of Gnedin & Kravtsov 2011), rather than calculated in post-processing, although we stress that this is still performed in subgrid fashion, as the molecular phase cannot be directly resolved. This allows SIMBA to utilise an  $H_2$ -based star formation model. In addition, the accretion of gas onto black holes depends on the gas phase: hot ( $T > 10^5$  K) gas is accreted via Bondi-Hoyle accretion and cold ( $T < 10^5$  K) gas using a torque-limited accretion model (Hopkins & Quataert, 2011; Anglés-Alcázar et al., 2017), which sets the black hole accretion model of SIMBA apart from the other two simulations we employ.

The black hole feedback model in SIMBA features two primary modes, plus an additional X-ray feedback mode. Unlike in TNG, the feedback modes occur concurrently. The modes are summarised in Table 3.1 and briefly described here:

- *‘Wind’ mode*: black holes with high Eddington ratios ( $\lambda_{\text{Edd}} > 0.2$ ) provide a kinetic kick to nearby gas with an outflow velocity proportional to  $\log M_{\text{BH}}$  which can reach up to  $1000 \text{ km s}^{-1}$ , based on the scaling relations of Fiore et al. (2017) and Ishibashi et al. (2018). For both this mode and the ‘jet’ mode (described below) the ejection is bipolar and parallel to the angular momentum vector of the inner disc used to calculate the SMBH accretion. Sources where only this mode is active are plotted in orange in Figure 3.2.
- *‘Jet’ mode*: for Eddington ratios  $\lambda_{\text{Edd}} < 0.2$ , the ‘jet’ mode switches on. Like the wind mode, this also ejects gas kinetically, although it can reach much higher velocities. Additionally, the temperature of this ejected gas is increased to the virial temperature of the halo although the thermal energy is typically only a few percent of the kinetic energy (Davé et al., 2019). The velocity of the outflow increases with decreasing Eddington ratio until it reaches a maximum speed of  $7000 \text{ km s}^{-1}$  at  $\lambda_{\text{Edd}} < 0.02$ . The mass of the black hole must also exceed  $10^{7.5} M_\odot$  for this mode to be activated. In Figure 3.2, sources in pink have  $\lambda_{\text{Edd}} = 0.02 - 0.2$  and are labelled as ‘transition jets’ as the maximum velocity has not yet been reached. Sources in blue have  $\lambda_{\text{Edd}} < 0.02$  and are labelled as ‘full jets’ as the Eddington ratio is now low enough for the ejection velocity to meet its maximum. For consistency with previous papers, we keep the naming convention of ‘jet’ for this mode, but we would like to stress that it does not seek to model a resolved jet structure as some high-resolution simulations have done (e.g. Mukherjee et al. 2016, 2018; Talbot et al. 2021; Bourne & Sijacki 2021; Mandal et al. 2021).
- *‘X-ray’ feedback*: the final mode aims to model X-ray feedback from the accretion disc as a spherical, mostly thermal feedback mode. This requires full-speed jets (i.e.  $\lambda_{\text{Edd}} < 0.02$  and  $M_{\text{BH}} > 10^{7.5} M_\odot$ ) and low molecular gas fractions ( $f_{H_2} < 0.2$ ) to

<sup>3</sup><http://simba.roe.ac.uk>

be activated. The X-ray heating is applied to gas within the SMBH particle kernel and is either entirely thermal (for non-ISM gas) or half thermal and half kinetic (for ISM gas). [Davé et al. \(2019\)](#) found that this mode has a minimal effect on the galaxy mass function, but does help to fully quench the most massive galaxies ([Cui et al., 2021](#); [Appleby et al., 2021](#)). For simplicity, we will generally neglect this mode in our discussion as it has little effect on the initial quenching of galaxies and only becomes significant once the galaxy has already been mostly cleared of its gas.

### 3.2.2 Observational Samples

In this section, we present the observational data we use in our analysis. We select example observational samples with large numbers of sources with published information on *all of*: (1) AGN luminosities, (2) stellar masses, (3) molecular gas fractions and (4) star formation rates. Meeting these criteria, we selected two of the most recent samples at  $z \simeq 0$  ([Koss et al., 2021](#); [Zhuang et al., 2021](#)) and one at  $z \simeq 2$  ([Bischetti et al., 2021](#)).

#### Koss et al. (2021)

Our first low redshift sample is from the *Swift*-BAT AGN Spectroscopic Survey<sup>4</sup> (BASS; [Koss et al. 2017](#)); an optical spectroscopic follow-up of AGN identified by the ultra-hard X-ray ( $> 10$  keV) *Swift*-BAT all-sky catalogue ([Baumgartner et al., 2013](#)). Stellar masses were calculated by combining near-IR data from 2MASS with mid-IR data from the AllWISE catalogue (see [Powell et al. 2018](#) for more details). Star formation rates were derived by decomposing spectral energy distributions (SEDs) of infrared data, mostly from *Herschel* ([Shimizu et al., 2017](#); [Ichikawa et al., 2019](#)).  $L_{\text{bol}}$  is calculated by applying a bolometric correction to the intrinsic hard X-ray luminosity from the AGN ([Marconi et al., 2004](#); [Rosas-Guevara et al., 2016](#)), given by:

$$\log\left(\frac{L_{\text{HX}}}{L_{\text{bol}}}\right) = -1.54 - 0.24L - 0.012L^2 + 0.0015L^3 \quad (3.2)$$

where  $L = \log(L_{\text{bol}}/L_{\odot}) - 12$ , which is then be solved for  $L_{\text{bol}}$ .

In this study, we use a subsample of nearby ( $0.01 < z < 0.05$ ) BAT AGN presented in [Koss et al. \(2021\)](#) for which there are additional CO(2-1) observations from JCMT and APEX, which are used by the authors to estimate measurements of molecular gas ( $H_2$ ) masses. This gives us a sample of 213 AGN host galaxies with  $H_2$  masses.

#### Zhuang et al. (2021)

Our other low redshift sample is presented in [Zhuang et al. \(2021\)](#) which is based on the catalogue formed by [Liu et al. \(2019\)](#) of Type I AGN from the Sloan Digital Sky Survey (SDSS; [York et al. 2000](#), DR7; [Abazajian et al. 2009](#)) which were selected primarily based on a detection of broad-line  $H\alpha$  emission. Galaxies in the redshift range  $0.3 < z < 0.35$  were chosen, with the requirement of sufficient signal-to-noise to allow AGN selection from emission-line diagnostics. SFRs are calculated from [O II]  $\lambda 3727$  and [O III]  $\lambda 5007$  following the method of [Zhuang & Ho \(2019\)](#). Black hole masses are calculated from broad  $H\alpha$  (following [Greene & Ho 2005](#)), allowing inference of stellar masses based on the empirical scaling relations of [Greene et al. \(2020\)](#).  $L_{\text{bol}}$  is calculated by applying a bolometric correction to the optical continuum luminosity at  $5100 \text{ \AA}$  ([McLure & Dunlop, 2004](#)).

<sup>4</sup><https://www.bass-survey.com>

	$z = 0$			$z = 2$		
	TNG	EAGLE	SIMBA	TNG	EAGLE	SIMBA
All galaxies	18991	12355	22389	11807	7856	5133
$\lambda_{\text{Edd}} \geq 1\%$	6280	245	5163	11078	717	4601
$L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$	1423	27	386	4361	228	2829
Combined selection	1263	18	170	4285	217	2729

**Table 3.2:** Sample sizes for our various AGN selection criteria at  $z = 0$  and  $z = 2$ . ‘All galaxies’ shows the number of sources at each redshift (after applying stellar mass cuts) and the ‘Combined selection’ requires both a high  $L_{\text{bol}}$  and high  $\lambda_{\text{Edd}}$ .

To estimate the molecular gas masses of their galaxies, the method presented in [Yesuf & Ho \(2019\)](#) is used which combines measurements of dust extinction, traced by the  $\text{H}\alpha/\text{H}\beta$  Balmer decrement, with gas-phase metallicities. This technique was calibrated against a sample of star-forming galaxies with CO measurements. [Zhuang et al. \(2021\)](#) show that this is a reliable method for Type I AGN as well being significantly less observationally intensive than CO measurements, resulting in estimated  $H_2$  masses in a large sample of 453 AGN host galaxies to be calculated.

### Bischetti et al. (2021)

For our high-redshift ( $z \sim 2$ ) sample, we use the data presented in [Bischetti et al. \(2021\)](#), which is a compilation of the relevant available host galaxy data on AGN at this epoch. This combines a subsample from the WISSH QSOs project ([Bischetti et al., 2017](#)), with additional data from [Perna et al. \(2018\)](#), X-ray selected QSOs from the SUPER survey ([Circosta et al., 2021](#)) and AGN-hosting SMGs ([Bothwell et al., 2013](#)). Measurements of the IR luminosity were obtained from AGN-corrected SED fitting and used to derive the SFR ([Kennicutt, 1998](#)). Molecular gas masses were calculated from measurements of various CO or [C II] transitions (see [Bischetti et al. 2021](#) and references therein for details).

Due to the time-consuming observations required, there are only a small number of high-redshift AGN with measurements of all the quantities we require ( $L_{\text{bol}}$ ,  $M_{H_2}$ , SFR and  $M_\star$ ). Therefore, to maximise the number of targets available, we select galaxies from this sample in the broad redshift range of  $z = 1 - 5$ , yielding 20 sources with suitable data, with a mean redshift of  $z = 2.2$ . However, we checked our results against the narrower redshift range of  $z = 2 - 3$  and found no qualitative differences. Additionally, the subsample we select has a slightly smaller range in  $f_{H_2}$  than the parent sample. We discuss this further when we compare these observational results to the simulation data in Section 3.3.3.

### 3.2.3 Key Quantities

In this study, we focus on two key global galaxy quantities and investigate how they depend on the luminosity of the AGN. The galaxy properties are the specific star formation rate;

$$\text{sSFR} = \frac{\text{SFR}}{M_\star}, \quad (3.3)$$

where  $M_\star$  is the stellar mass of the galaxy, and the molecular gas fraction;

$$f_{H_2} = \frac{M_{H_2}}{M_\star}, \quad (3.4)$$

where  $M_{H_2}$  is the mass of molecular hydrogen gas.



These quantities are commonly investigated in observational studies to probe the effect the AGN has on the molecular gas and how that influences the galaxy’s star formation (e.g. [Harrison 2017](#); [Shangguan et al. 2018](#); [Rosario et al. 2018](#); [Scholtz et al. 2018](#); [Zhuang et al. 2021](#)). Additionally, sSFR is frequently used to determine whether a galaxy is star-forming or quiescent (see Section 3.2.5 for our quenching/gas-depletion definitions).

Due to the resolution limits of the simulations, very low star formation rates and molecular gas masses values are not resolved (this is partly resolution-dependent and is also affected by the SFR averaging timescales, see Appendix A of [Donnari et al. 2019](#)). To track these galaxies, we artificially plots these sources in the figures with arbitrarily low values of  $\text{sSFR} = 10^{-14} \text{ yr}^{-1}$  and  $f_{H_2} = 10^{-5}$ , respectively. For presentation reasons we scatter these points around these values with a standard deviation of 0.2. (see [Weinberger et al. 2018](#) who follow a similar method). We note that these exact values are not used in any calculation in this study and they are only used to visually represent them as a quenched and gas-depleted populations in the figures. These sources represent around 10 – 20% of all galaxies across the three simulations in our sample at  $z = 0$ .

For each simulation, we chose an aperture of either 30 kpc or twice the stellar half-mass radius depending on the default choice of the original simulation teams. As the majority of the star formation of a galaxy falls well within both these definitions, there is a negligible difference between them. For a more detailed discussion of the effect of aperture choice, see [Donnari et al. \(2019\)](#) and Appendix C in [Weinberger et al. \(2018\)](#).

Another factor of consideration is the timescale over which we calculate the star formation rate in the simulations. [Donnari et al. \(2019\)](#) investigated the effect of changing the SFR averaging timescale, from instantaneous to 1000 Myr. By calculating the resulting star formation main sequence for each timescale, they showed that varying the timescale makes negligible difference to the slope of this sequence, although there is an offset of around 0.1 – 0.2 dex between the shortest (10 Myr) and longest (1000 Myr) timescales. These differences are smaller than the range of quenching definitions we use which are the dominant source of uncertainty in our results for correlation coefficients and quenched fraction (see Section 3.2.5). Therefore, we chose to use only the instantaneous SFRs in this study.

A final choice is whether to include both centrals and satellites in our simulation galaxy sample, with different approaches being taken in previous studies ([Donnari et al. 2019](#) make no distinction, whereas [Weinberger et al. 2018](#) only selects centrals). We include both centrals and satellites, with central galaxies making up 60 – 70% of our sample across the three simulations. We also test our results on a centrals-only selection and find no qualitative difference, except for a slight reduction in quenched and gas-depleted fractions in lower- $M_*$  galaxies, e.g. galaxies with  $M_* \lesssim 10^{10} M_\odot$  have a global gas-depleted fraction of 7 – 13 percentage points lower when we select only for centrals. We find that this change is even lower in AGN hosts, changing the gas-depleted fraction by  $\lesssim 5$  percentage points.

### 3.2.4 AGN Selection Criteria

Our goal is to broadly emulate the approach taken by observational studies to investigate the relationships between AGN activity and host galaxy properties (see Section 3.1). Therefore, we design two AGN-selection methods that roughly emulate current observational limits in separating luminous AGN hosting galaxies from non-AGN galaxies; one based on the bolometric luminosity of the AGN ( $L_{\text{bol}}$ ) and the other on the Eddington ratio ( $\lambda_{\text{Edd}}$ ).

In the simulations, the bolometric luminosity<sup>5</sup> is calculated from the instantaneous accretion rate of the central SMBH in each galaxy:

<sup>5</sup>Note, that the terms bolometric luminosity ( $L_{\text{bol}}$ ) and AGN luminosity ( $L_{\text{AGN}}$ ) are used interchangeably throughout this thesis.

$$L_{\text{bol}} = \epsilon_r \dot{M}_{\text{accr}} c^2, \quad (3.5)$$

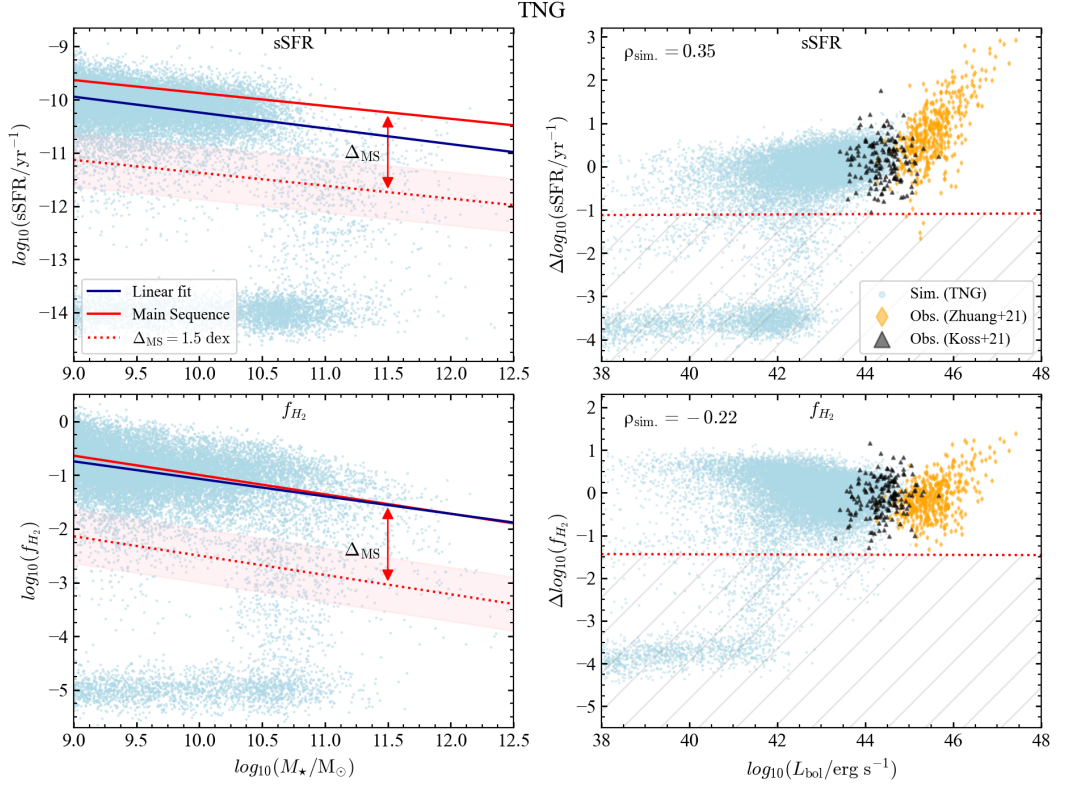
where  $\dot{M}_{\text{accr}}$  is the mass accretion rate flowing onto the black hole and  $\epsilon_r$  is the radiative efficiency of the AGN. EAGLE and SIMBA use the canonical value of  $\epsilon_r = 0.1$  (Crain et al., 2015; Davé et al., 2019) and TNG uses  $\epsilon_r = 0.2$  (Weinberger et al., 2017). A reasonable observational cut-off for studying high-luminosity AGN is  $L_{\text{bol}} \geq 10^{44}$  erg s<sup>-1</sup> which we use as our luminosity-based selection. This corresponds to the lower end of the bolometric luminosities in the Koss et al. (2021) comparison sample (see Figure 3.8).

However, this simple relation between the accretion rate and  $L_{\text{bol}}$  may only be valid for high-Eddington ratio, Shakura-Sunyaev discs (SSDs; Shakura & Sunyaev 1973). In low- $\lambda_{\text{Edd}}$  systems, the disc becomes geometrically thick and radiatively inefficient. The impact of the accretion efficiency on  $L_{\text{bol}}$  motivates a second AGN definition based on the Eddington ratio. Following Rosas-Guevara et al. (2016), we assume that sources with  $\lambda_{\text{Edd}} \geq 1\%$  are radiatively efficient, thin discs, with high X-ray luminosities and we define these as our high- $\lambda_{\text{Edd}}$  selection. We also define a ‘combined’ selection where we require the AGN to satisfy both criteria. There are also other methods to account for inefficient AGN accretion. For instance, in Habouzit et al. (2021, 2022a),  $L_{\text{bol}}$  is scaled down proportionally to the Eddington ratio if  $\lambda_{\text{Edd}} < 10\%$ . We tested our results using this model and found it yielded qualitatively similar results to our combined AGN definition, showing our conclusions are insensitive to the exact model used for inefficient AGN accretion.

Table 3.2 shows the number of sources for each of these selection criteria at  $z = 0$  and  $z = 2$ , across the three simulations. We note that, especially at  $z = 0$ , EAGLE shows far fewer high- $L_{\text{bol}}$  sources than the other two simulations. This can also be seen by the scarcity of points above the blue dashed line in Figure 3.2. Previous work has noted the lower AGN luminosity distribution function in EAGLE compared to TNG and SIMBA (Habouzit et al., 2022a) and it has been suggested that this is caused by both a lower BH mass function in the high- $M_{\text{BH}}$  regime in EAGLE (Habouzit et al., 2021) as well as lower accretion rates in these massive BHs, possibly due to strong supernovae feedback, the efficient thermal feedback model and the modified Bondi-Hoyle accretion model (Rosas-Guevara et al., 2016; Habouzit et al., 2022a). This discrepancy between the simulations does not effect our ability to investigate the trends across the three simulations; however, we do take care to account for the lower statistics of AGN host galaxies in EAGLE in our discussion of the results.

### 3.2.5 Quenching & Gas-Depletion Definitions

There are various ways to define a galaxy as having quenched star formation, including the distance from the main sequence (Weinberger et al., 2018), a cut in sSFR (Donnari et al., 2019) or based on a colour selection (e.g. UVJ diagrams, Donnari et al. 2019; Akins et al. 2022). In this study, we use a main sequence-based quenching definition. We prefer this over a colour-based definition as this would require assumptions and modelling of the dust content of galaxies (e.g. Trayford et al., 2016; Akins et al., 2022). We use the main-sequence model presented in Weinberger et al. (2018) for their analysis of the TNG simulation. This is a modification of the Ellison et al. (2015) model, with an additional redshift dependence based on Schreiber et al. (2015). We note that we repeated our analysis using other main sequence definitions (Speagle et al., 2014) including a non-linear model (Whitaker et al., 2014) and found it made a negligible difference to the results compared to the systematic uncertainties outlined below. For the molecular gas fraction, we use a similar method, utilising the scaling relations presented in Tacconi et al. (2018) as our ‘main-sequence’ model for this quantity. Both of these models are redshift-dependent, as the normalisation of both the star-forming and gas-fraction main sequences are higher at



**Figure 3.3:** Demonstration of how the  $M_*$  dependency is removed from the quantities of interest and the correlation coefficient between the resulting residuals and  $L_{\text{bol}}$  is calculated, using TNG at  $z = 0$  as an example. *Left Column:* sSFR (upper panel) and  $f_{H_2}$  (lower panel) against  $M_*$ . The red line shows the main sequence scaling relationship adopted for each quantity. Our quenched/gas-depleted definition is galaxies that lie some distance ( $\Delta_{\text{MS}}$ ) below these relations, which is allowed to vary between  $-1 \text{ dex} \leq \Delta_{\text{MS}} \leq -2 \text{ dex}$ , as shown by the red shaded region. The blue line shows a linear regression to the remaining points, after excluding the quenched/gas-depleted sources from the fit. *Right Column:* the resulting residuals’ relationship with  $L_{\text{bol}}$ , after removing the  $M_*$  dependence including comparison observational data. Calculated Spearman correlation coefficients are shown in the top left corner of the panels, revealing a weak positive correlation for sSFR and a weak negative correlation for  $f_{H_2}$ . The hatched region shows the quenched/gas-depleted sources that have been excluded from the correlation calculation.

$z = 2$  than  $z = 0$ . As noted before, the simulations tend to under-predict  $f_{H_2}$  at  $z = 2$  by a factor of 1.5 – 3 (Lagos et al., 2015; Popping et al., 2019), therefore we shift the scaling relation in Tacconi et al. (2018) down by 0.3 dex to better match the peak of the  $f_{H_2} - M_*$  plane in the simulations.

As illustrated in Figure 3.3, we then define a galaxy as ‘quenched’ if it lies more than a certain distance,  $\Delta_{\text{MS}}$ , below the star-forming main sequence (Weinberger et al., 2018). Likewise, a galaxy is defined as ‘gas-depleted’ if it lies  $> \Delta_{\text{MS}}$  below the gas fraction scaling relation (Tacconi et al., 2018). To check the sensitivity of our results to this cut, we vary this value between  $-1 \text{ dex} \leq \Delta_{\text{MS}} \leq -2 \text{ dex}$ ; a range similar to what has been presented in previous works. We show this range for sSFR on Figure 3.3 by the red shaded region. Where relevant, the uncertainty this adds to our results is represented by shaded violin plots, showing the posterior probability density for changing values of  $\Delta_{\text{MS}}$ .

### 3.2.6 Correlation Coefficients

Part of our analysis involves quantifying the trends seen between galaxy properties and AGN luminosity by calculating correlation coefficients between these quantities. However, as our galaxy properties of interest have a known dependence on stellar mass, which itself causes a correlation with  $L_{\text{bol}}$  (e.g. Stanley et al., 2017; Scholtz et al., 2018; Tacconi et al.,

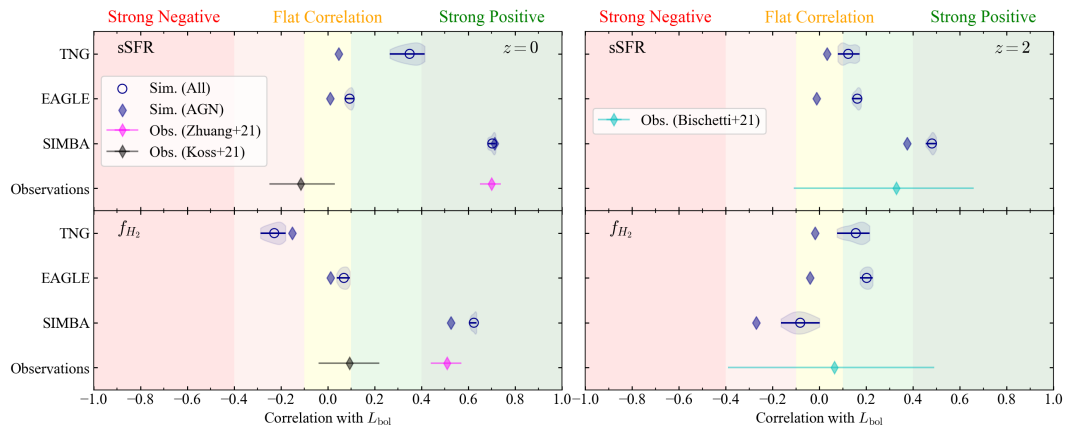
2020), we first need to correct for the effect of  $M_*$ . To do this, we use a residual method which is represented in Figure 3.3, using the TNG simulation as an example. The panels in the left-hand column show the sSFR- and  $f_{H_2}$ - $M_*$  planes.

For this correlation coefficient analysis we only consider ‘star-forming’ galaxies which are all galaxies that are not defined as quenched (for the sSFR analyses) or gas-depleted (for the  $f_{H_2}$  analyses). We discuss quenched and gas-depleted fractions in detail in Section 3.3.4, but, for a rough guide, the ‘star-forming’ fraction for these simulations at  $z = 0$  is 69 – 80% for sSFR and 79 – 86% for  $f_{H_2}$ . Once we have selected our star-forming sample, we characterise the stellar mass dependence. This approach avoids including the unresolved SFR and molecular gas mass values (see Section 3.2.3) in the correlation analysis which would have to be arbitrarily placed at some low value. Furthermore, this is well motivated because the majority of AGN host galaxies fall within the ‘star forming’ population (see Section 3.3.3 for a discussion on the quenched and gas-depleted fractions for all AGN-selected galaxies). However, we also tested the robustness of our results by including the quenched and unresolved systems and found it made a negligible difference to the qualitative trends.

After we have excluded the quenched or gas-depleted galaxies, we perform a linear regression on the remaining data points in the sSFR or  $f_{H_2}$ - $M_*$  plane and calculate the residual of this fit for each point. This is plotted against  $L_{\text{bol}}$  as can be seen in the top right panel of Figure 3.3. Points in the red hatched region show the quenched/depleted and thus excluded galaxies for an example value of  $\Delta_{\text{MS}} = -1.5$ . The Spearman correlation coefficient,  $\rho$ , is then calculated between the residuals and  $L_{\text{bol}}$ . This can take values between -1 and 1, where -1 is a strong negative correlation and 1 is a strong positive correlation with values around 0 showing no correlation between the quantities. In this study we consider the strength of the correlation as ‘strong’ for  $|\rho| > 0.4$ , ‘weak’ for  $0.1 \leq |\rho| < 0.4$  and ‘flat’ for  $|\rho| < 0.1$ . This method is equivalent to computing the semi-partial correlation between sSFR and  $M_*$  (this approach is similar to that in [Zhuang et al. 2021](#), although they control for  $M_{H_2}$  rather than  $M_*$ ). Figure 3.3 shows an example of a weak positive (sSFR, *top right*) and weak negative ( $f_{H_2}$ , *bottom right*) correlation in the TNG simulation, with the correlation coefficient values shown in the corner of each panel. The uncertainty caused by the varying quenching/depletion definition is propagated through to give a probability density region of correlation coefficients. These can be seen as the blue shaded ‘violin’ plots on Figure 3.4. We discuss this figure and present the results of this correlation analysis in Section 3.3.1.

### 3.3 Results

Here we present our results from an investigation of the relationships between AGN bolometric luminosities and host galaxies’ specific star formation rates (sSFR) and molecular gas fractions ( $f_{H_2}$ ), as predicted from three state-of-the-art cosmological simulations (TNG, EAGLE and SIMBA). We follow the broad approaches taken in many observational papers, which search for evidence of the impact of AGN on their host galaxies by investigating: (1) correlations of sSFR and  $f_{H_2}$  with  $L_{\text{bol}}$  (Section 3.3.1); (2) a comparison between AGN and non-AGN host galaxies in the  $f_{H_2}$ - $M_*$  or sSFR- $M_*$  plane (Section 3.3.2); and (3) gas fraction distributions for the highest  $L_{\text{bol}}$  sources (Section 3.3.3) as well as the fraction of AGN which are quenched or gas-depleted (Section 3.3.4). As motivated in Section 3.2.2 we investigate this at both  $z = 0$  and  $z = 2$ . In Sections 3.3.2, 3.3.3 & 3.3.4, we note that we obtain qualitatively similar results for sSFR as we do for  $f_{H_2}$ . Therefore, for brevity, we mostly show the results for  $f_{H_2}$  in the main body of the paper, referring interested readers to the equivalent plots for sSFR in Appendix 3.6.



**Figure 3.4:** Correlation coefficients between  $L_{\text{bol}}$  and host galaxy properties (upper panel: sSFR; lower panel:  $f_{H_2}$ ) after removing the dependence on  $M_*$  for the simulations and observational data investigated. The left column shows the results for  $z = 0$  and the right column shows  $z = 2$ . For the simulation data, the mean result for star-forming galaxies is shown by the blue circles, with the shaded regions showing the probability densities for the uncertainty on this value, based on varying the quenching or gas-depletion definition (see Section 3.2.6). The filled diamonds show the correlation of only the AGN (high- $\lambda_{\text{Edd}}$  selected sources). We also show the data from the observations as diamond-shaped points, with black points for Koss et al. (2021) and purple for Zhuang et al. (2021), where the error bars on these values show the 95% confidence interval on the calculation of the correlation. At  $z = 2$ , the cyan points show the sample from Bischetti et al. (2021). Although there are clear differences in the predictions, it can be seen that none of the simulations predict strong negative correlations ( $\rho < -0.4$ ) between  $f_{H_2}$  or sSFR with  $L_{\text{bol}}$ , qualitatively consistent with observational results.

### 3.3.1 Correlations with $L_{\text{bol}}$

For each of TNG, EAGLE and SIMBA, we compute the strength and sign of the Spearman correlation coefficient,  $\rho$ , between  $L_{\text{bol}}$  and our key galaxy properties of interest: sSFR and  $f_{H_2}$  (see Section 3.2.3). For this analysis we focus only on ‘star-forming’ galaxies, which we define as those which remain after excluding quenched galaxies (for sSFR analyses) or gas-depleted galaxies (for  $f_{H_2}$  analyses) following Section 3.2.5. For this star-forming population, we then remove the intrinsic correlation with stellar mass following Section 3.2.6. We note that we return to consider the quenched and gas-depleted galaxies in Sections 3.3.3 & 3.3.4.

The results from the correlation analysis are presented in Figure 3.4, which shows the calculated Spearman correlation coefficients between  $L_{\text{bol}}$  and sSFR (*top row*) and between  $L_{\text{bol}}$  and  $f_{H_2}$  (*bottom row*) at both  $z = 0$  (*left column*) and  $z = 2$  (*right column*). The open circles show the star-forming galaxy population and the diamonds show the  $\lambda_{\text{Edd}}$ -selected AGN. Several observational papers have investigated the trends between star formation rates or molecular gas fractions and  $L_{\text{bol}}$  as evidence for AGN feedback (Lutz et al., 2010; Mainieri et al., 2011; Page et al., 2012; Harrison et al., 2012; Stanley et al., 2015; Kakkad et al., 2017; Shangguan & Ho, 2019; Ramasawmy et al., 2019; Jarvis et al., 2020; Circosta et al., 2021). However, we find that the simulations do not predict any strong negative correlations between  $L_{\text{bol}}$  and host galaxy properties.

Out of the 12 coefficients calculated, eight show positive trends (populating the right hand side of Figure 3.4), three show flat trends and only one is (weakly) negatively correlated, with  $\rho = -0.23$ . Even for this negative trend, which is seen in TNG, there is only a reduction in the mean gas fraction of around 1 dex, over  $\sim 6$  orders of magnitude in AGN luminosity (see lower left panel of Figure 3.3).

Looking at Figure 3.4 in more detail, there are clear differences between the three simulations for their predictions of the trends between sSFR or  $f_{H_2}$  and  $L_{\text{bol}}$ . If we first consider the sSFR of the entire star-forming galaxy population with no AGN selection

(open circles) at  $z = 0$  (*top left*): TNG is weakly positive with  $\rho = 0.35$ , EAGLE shows a negligible correlation (i.e. flat trend) of  $\rho = 0.09$  and SIMBA shows a strong positive relation of  $\rho = 0.70$ . Looking at  $f_{H_2}$  at  $z = 0$ , we see that SIMBA is again strongly positively correlated ( $\rho = 0.62$ ) and EAGLE again shows negligible correlation ( $\rho = 0.07$ ), whilst TNG shows a weak negative correlation ( $\rho = -0.23$ ). If we also consider the results at  $z = 2$  (*right column*), there is generally more agreement across the three simulations. However, there are clearly different predictions of the evolution of these relationships. At  $z = 2$ , TNG now shows weakly positive correlations for both quantities, whereas there is a weak negative correlation at  $z=0$  for gas fraction. EAGLE shows very little evolution across the two epochs, although there is a weakly positive trend with gas fraction at  $z = 2$  ( $\rho = 0.20$ ). SIMBA is still strongly correlated in sSFR at  $z = 2$ , however, it shows a flat trend with gas fraction ( $\rho = -0.08$ ), which is a significant change from the strong positive trend at  $z = 0$ .

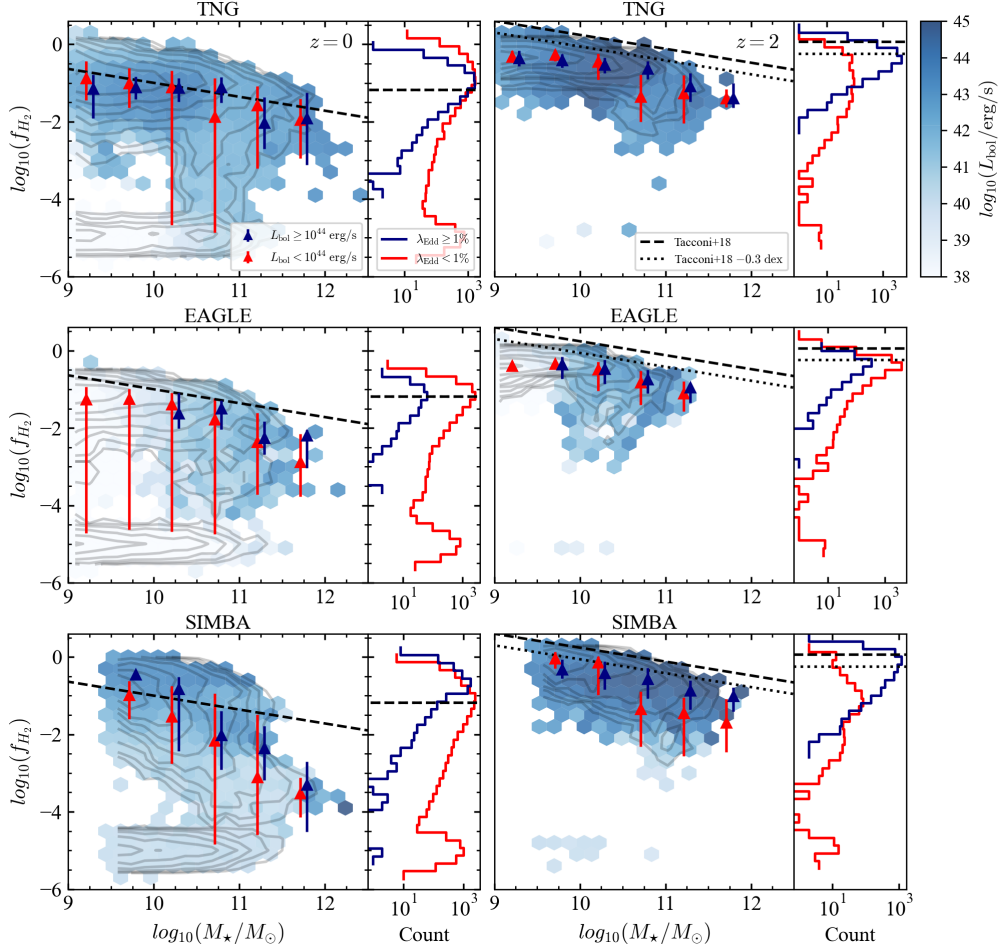
We next consider the correlation coefficients for only star-forming galaxies in the simulations that are also selected to be AGN based on their Eddington ratios ( $\lambda_{\text{Edd}} > 1\%$ , see Section 3.2.4 and shown as filled diamonds here). We find that selecting only high- $\lambda_{\text{Edd}}$  AGN for this correlation analysis has either little effect or reduces the strength of the trends observed, i.e.  $\rho$  values remain unchanged, or in most cases move closer to zero. This latter effect is likely due to a smaller dynamic range in  $L_{\text{bol}}$  being sampled. In no case does only considering high- $\lambda_{\text{Edd}}$  AGN convert a positive correlation to a negative one, showing that they generally behave in a similar way to the overall star-forming galaxy population in the simulations. The only exception here is in  $f_{H_2}$  in SIMBA at  $z = 2$ , where the trend moves from uncorrelated ( $\rho = -0.05$ ) to weakly negative ( $\rho = -0.27$ ).

The comparison observational data, as described in Section 3.2.2, are shown at the bottom of each panel of Figure 3.4. Of our two primary comparison samples at  $z = 0$ , the [Koss et al. \(2021\)](#) sample overlaps more with the  $L_{\text{bol}}$  parameter space of the simulations (see Figure 3.3). This shows a mostly flat trend which is qualitatively consistent with the results of EAGLE and TNG (for sSFR) for AGN-selected samples. In contrast, the [Zhuang et al. \(2021\)](#) data predicts very strong positive trends. This tension between the two observational samples could be caused by the different selection criteria used by the two studies, or be because the sources in [Zhuang et al. \(2021\)](#) probe higher values of  $L_{\text{bol}}$ , possibly corroborating the work of [Rosario et al. \(2012\)](#) who found that positive correlations are only seen once the AGN luminosity exceeds  $L_{\text{bol}} > 10^{44.7} \text{ erg s}^{-1}$ .

At  $z = 2$ , the observational data from [Bischetti et al. \(2021\)](#) suggest a positive correlation for sSFR and a flat correlation for  $f_{H_2}$ . Due to the small number of sources, the uncertainties are large, but we note that these error bars cover a similar range to the spread we find between the simulations.

It is beyond the scope of this work to perform a more quantitative assessment to compare the observations to the simulations, which would require careful considerations of sample selection effects and controlling for uncertainties in conversion factors between observables and derived quantities, etc.. Instead, we use these results to demonstrate in a *qualitative* sense that the broad behaviour seen in the simulations in these trends is consistent with observations, in that no strong negative correlations are observed between gas fractions and star formation rates.

Although there is agreement between the simulations in not predicting strong negative correlations, we have shown that the differences in their predictions are still quite stark. However, by simply looking at trends or correlation coefficients, it is difficult to understand the details of the investigated relationships and to infer whether this is due to their differing AGN feedback subgrid models or because of other factors. Therefore, in Sections 3.3.2, 3.3.3 & 3.3.4 we consider the *distribution* of galaxy properties, also considering the different feedback modes used in the simulations, to help understand the differences seen here.



**Figure 3.5:** The molecular gas fraction against stellar mass at  $z = 0$  (left column) and  $z = 2$  (right column) for the three simulations. Contour lines show logarithmic number density and the colouring of the pixels shows the mean bolometric luminosity of the sources within each bin. Galaxies with unresolved  $f_{H_2}$  have been scattered around  $f_{H_2} = 10^{-5}$ . The dashed line shows the observed  $f_{H_2}$  main sequence from Tacconi et al. (2018). At  $z = 2$ , we shift this line down by 0.3 dex, shown as the dotted line. The triangular points in the left-hand subpanels show the median and 16<sup>th</sup>-84<sup>th</sup> percentiles for AGN (blue) and non-AGN (red) galaxies based on an  $L_{\text{bol}}$  selection, grouped in stellar mass bins of 0.5 dex. The histograms in the right-hand subpanels show the logarithm of the number of sources classified as AGN or non-AGN based on an  $\lambda_{\text{Edd}}$  definition, with the dotted line showing the expected gas fraction for a fiducial  $M_* = 10^{10.5} M_\odot$  galaxy on the main sequence. We can see that both the  $L_{\text{bol}}$ - and  $\lambda_{\text{Edd}}$ -selected sample of AGN reside preferentially in gas-rich galaxies.

### 3.3.2 AGN and non-AGN in the $f_{H_2} - M_*$ plane

Another common method that observers have used to look for evidence of the effect of AGN feedback, is to compare the host galaxy properties of active galaxies to those without known AGN (Rosario et al., 2018; Bischetti et al., 2021; Circosta et al., 2021). This is often done by comparing the location of AGN host galaxies with mass-matched galaxies without AGN, or to established scaling relations of sSFR or  $f_{H_2}$  with stellar mass for ‘main sequence’ star-forming galaxies. In this section, we aim to emulate this method using the simulations, by considering our two ways of categorising AGN, bolometric luminosities and Eddington ratios (Section 3.2.4), and comparing their host galaxy’s  $f_{H_2}$  to that of inactive galaxies in the simulations and also to the gas fraction scaling relations of Tacconi et al. (2018). Although we focus on  $f_{H_2}$ , the equivalent figures for sSFR are shown in Appendix 3.6, for which we obtain broadly consistent conclusions. Unlike for the correlation coefficient analysis, which only considered ‘star-forming galaxies’, for the following analysis we now consider the whole galaxy sample.

Figure 3.5 presents the molecular gas fractions as a function of stellar mass at  $z = 0$  (*left column*) and  $z = 2$  (*right column*) for galaxies above  $M_\star \geq 10^9 M_\odot$  in the three simulations. The contour lines in the main panels show logarithmic number density of sources and the galaxies have also been grouped into hexagonal bins, which are then coloured by the mean bolometric luminosity of the AGN within. We also plot the median and 16<sup>th</sup>-84<sup>th</sup> percentiles for a luminosity-selected ( $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$ ) sample of AGN (blue triangles) and for the corresponding non-AGN sample (red), grouped in stellar mass bins of 0.5 dex. The histograms to the side show logarithmic number counts of sources with high  $\lambda_{\text{Edd}}$  ratios ( $> 1\%$ ) in blue and with low  $\lambda_{\text{Edd}}$  in red. As an observational comparison, we show the molecular gas scaling relation from [Tacconi et al. \(2018\)](#) as our ‘main-sequence’ by the black dashed line in the  $f_{H_2} - M_\star$  plane. On the histograms, this is represented by another black line showing the main-sequence gas fraction for a fiducial galaxy of  $M_\star = 10^{10.5} M_\odot$ .

Considering the distribution of all galaxies in this  $f_{H_2} - M_\star$  plane (grey contours), we can see that all three simulations reproduce the observed galaxy bimodality with star forming systems clustered around the main sequence at  $f_{H_2} \sim 10^{-1}$  and a large gas-depleted fraction below (i.e. 15 – 25% of the galaxies lie  $> 1$  dex below the gas fraction main sequence at  $z = 0$ ). At  $z = 0$ , the simulations are generally successful at reproducing the gradient and normalisation of the observed gas scaling relations ([Lagos et al., 2015](#); [Diemer et al., 2018, 2019](#); [Davé et al., 2019](#)). At  $z = 2$ , all three simulations slightly under-predict the gas fraction, as previously noted in [Lagos et al. \(2015\)](#) and [Popping et al. \(2019\)](#). Figure 3.9 shows a similar result for the sSFR.

In Figure 3.5, the darker blue pixels show the brightest AGN (highest accretion rates) and the lighter pixels show the lowest accretion rate systems. In TNG at  $z = 0$ , we can see that the most luminous AGN lie around the gas fraction main sequence and at lower stellar masses. We see a strong downwards ‘plume’ of quenching galaxies all with similar luminosities of  $L_{\text{bol}} = 10^{42} \text{ erg s}^{-1}$  (i.e. these low luminosity systems would not be detected as ‘AGN’ in most observational work). As we further discuss in Section 3.4.2 this dramatic plume of low- $f_{H_2}$  galaxies over a narrow mass range ( $M_\star = 10^{10.2-11} M_\odot$ ) is associated with a switch to the low-accretion (‘kinetic’) feedback mode in TNG ([Weinberger et al., 2018](#); [Terrazas et al., 2020](#)). At  $z = 2$ , there are very few gas-depleted systems and the mean  $L_{\text{bol}}$  is high all along the main sequence, peaking at  $M_\star \simeq 10^{10.5} M_\odot$ .

The brightest  $L_{\text{bol}}$  sources in EAGLE are in high-mass galaxies ( $M_\star \gtrsim 10^{10.5} M_\odot$ ) on the main sequence and the mean  $L_{\text{bol}}$  increases smoothly from low- to high- $M_\star$ ; a trend seen at both redshifts.

The mean  $L_{\text{bol}}$  in SIMBA is high ( $L_{\text{bol}} \gtrsim 10^{43} \text{ erg s}^{-1}$ ) around the main sequence line, but at around 1 – 2 dex below this, the mean  $L_{\text{bol}}$  drops significantly in both low- and high- $M_\star$  galaxies (although there are a few high- $L_{\text{bol}}$  sources here as we explore in Section 3.3.3). Like the other two simulations, there are very few galaxies with unresolved  $f_{H_2}$  at  $z = 2$ , although the mean  $L_{\text{bol}}$  of the galaxies towards the bottom of the main-sequence grouping is much lower than similar galaxies in the other two simulations.

Considering a luminosity selection of  $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$ , the triangular points, representing the median gas fraction of AGN and non-AGN across different mass bins, show that  $L_{\text{bol}}$ -selected AGN reside on or near the peak of the gas-rich galaxies and the observational main sequence in all three simulations. The median gas fraction of the AGN are always close to the non-AGN points across the three simulations, except they show a much narrower range: the 16<sup>th</sup>-84<sup>th</sup> percentiles shown in Figure 3.5 never extend down to the lowest ‘unresolved’ gas fractions.

As motivated in Section 3.2.4, we can also make an AGN selection based on the Eddington ratio. Looking at the right-hand subplots, we can see that in all three simulations the highly-accreting AGN (shown by the blue line) reside in high- $f_{H_2}$  galaxies on the main



sequence with the peak of the AGN population matching or even exceeding the peak of the non-AGN population (in red). We explore this further in Section 3.3.3 by considering the gas-depletion fractions of these high- $\lambda_{\text{Edd}}$  systems.

In Figure 3.9 we present the same analysis for the  $sSFR - M_*$  plane. The results are generally consistent with those for  $f_{H_2}$ , but there are some small differences, especially in TNG. For example, at  $z = 2$  in TNG there is a larger population of quenched galaxies than there is gas-depleted. These quenched galaxies are mostly associated with the kinetic feedback mode and lie in the narrow stellar mass range  $M_* = 10^{10.2-11} M_\odot$  (we discuss this further in Section 3.4.2). This suggests that, in the TNG model, a galaxy does not need to be fully depleted of molecular gas before the star formation has quenched at  $z = 2$ . We also see that the non-AGN sample in this mass range at  $z = 0$  has a much lower median sSFR compared to the  $L_{\text{bol}}$ -selected AGN sample. This is in contrast to the result for the molecular gas, where the two samples have similar median values for  $f_{H_2}$  (Figure 3.5). Additionally, the highest mass bin ( $M_* \geq 10^{11.5} M_\odot$ ) in SIMBA at  $z = 0$  is entirely quenched, whereas the galaxies had non-negligible gas fractions in the  $f_{H_2}$  plane. Nevertheless, the median value for the AGN sample is still higher than the non-AGN for both sSFR and  $f_{H_2}$  in the high-mass ( $M_* \geq 10^{11} M_\odot$ ) regime in SIMBA.

From Figure 3.5 (and 3.9), we can conclude that the simulations predict that AGN preferentially live in high-sSFR, high- $f_{H_2}$  galaxies. This result holds for both an  $L_{\text{bol}}$ - and  $\lambda_{\text{Edd}}$ -based AGN selection, showing that the simulations do not predict large differences between active and inactive galaxy populations, despite the implementation of strong AGN feedback models. This is in qualitative agreement with observational work at  $z = 0$  (e.g. Rosario et al., 2018; Jarvis et al., 2020). However, the observations have less of a coherent picture at  $z \simeq 2$  where some studies find AGN might tend to live in lower  $f_{H_2}$  sources (e.g. Kakkad et al., 2017; Perna et al., 2018; Circosta et al., 2021), potentially in tension with these results from the simulations.

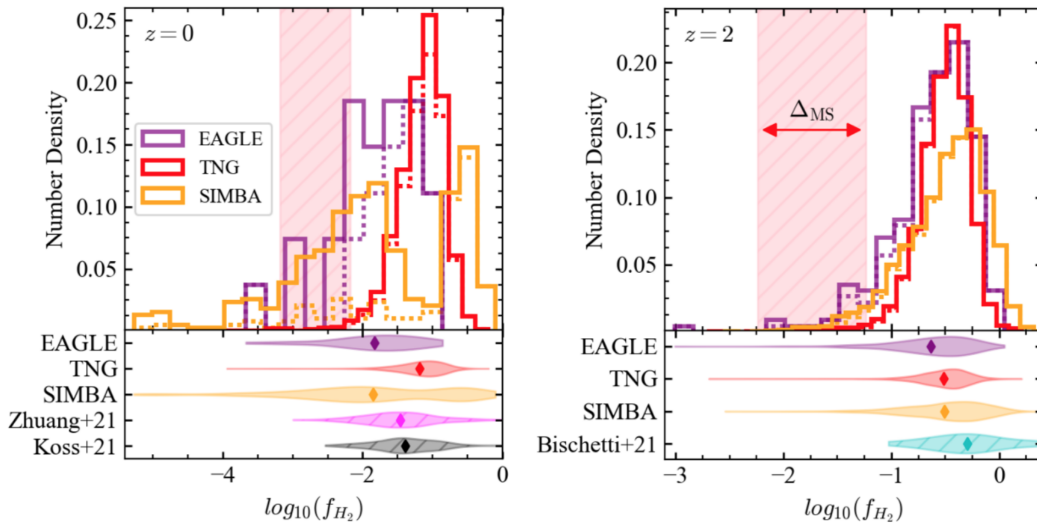
Despite a broadly consistent picture across the three simulations, we find that there are differences in the predictions between the three simulations, especially in the exact location of the brightest AGN in the  $f_{H_2} - M_*$  plane. In Section 3.3.3 we explore these brightest systems and investigate the effects of the feedback and accretion modes.

### 3.3.3 $f_{H_2}$ Distributions in AGN

To investigate the effect of powerful AGN on their host galaxies, many observational studies look for evidence of feedback in the brightest ( $L_{\text{bol}} \gtrsim 10^{44} \text{ ergs}^{-1}$ ) AGN (e.g. Schulze et al., 2019; Scholtz et al., 2021; Zhuang et al., 2021; Bischetti et al., 2021). In this section, we look at the simulation predictions for the distribution of  $f_{H_2}$  among the brightest AGN in the simulations.

The results for the gas fraction distribution in  $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$  AGN hosts are shown in Figure 3.6 for both  $z = 0$  (*top panel*) and  $z = 2$  (*bottom panel*). The top part of each panel shows the number density distributions in  $f_{H_2}$  for the three simulations and the bottom part shows violin density plots, with the observations as a comparison. Gas-depleted systems lie to the left of the plot.

At  $z = 2$  (*bottom panel* of Figure 3.6), all three simulations show very similar behaviour. There is a roughly Gaussian distribution at a gas fraction of  $f_{H_2} \simeq 10^{-0.5}$  with a slightly longer tail extending into the depleted region to the left. The shapes of the distributions compare quite well with the observations from the Bischetti et al. (2021) sample as can be seen in the violin plots. We note that the full Bischetti et al. (2021) sample does contain some sources with lower  $f_{H_2}$  (down to  $f_{H_2} \simeq 10^{-1.4}$ ), but as we require values for all of  $f_{H_2}$ ,  $M_*$ , SFR and  $L_{\text{bol}}$  we had to select a subsample from their catalogue (see Section 3.2.2). However, the mean  $f_{H_2}$  value from their full sample is comparable to the mean of our subsample, even if our range in  $f_{H_2}$  is slightly narrower.



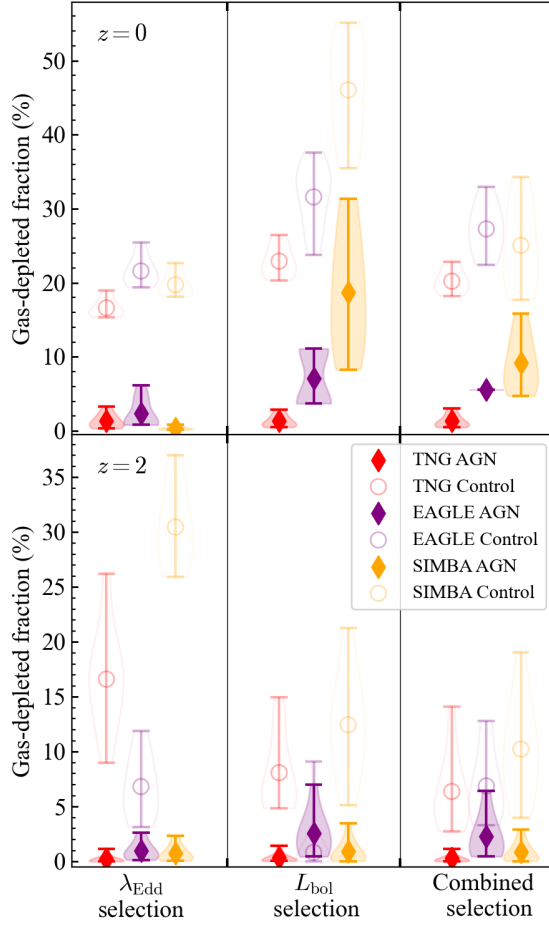
**Figure 3.6:** Gas fraction distributions for bright ( $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$ ) AGN. The top part of each panel shows the binned density distribution in  $f_{H_2}$  for the three simulations. The red shaded region to the left shows a range of definitions for gas-depletion. The dotted lines show the result if the additional AGN criteria that  $\lambda_{\text{Edd}} \geq 1\%$  is required. The lower part shows violin plots of the density distributions for both the simulation data and observational samples (Zhuang et al., 2021; Koss et al., 2021; Bischetti et al., 2021), where the mean value of  $f_{H_2}$  is shown by the diamond-shaped point. Although the simulations are in good agreement at  $z = 2$ , the predictions at  $z = 0$  are starkly different for our  $L_{\text{bol}}$ -selection.

However, at  $z = 0$ , there is much less agreement between the simulations. TNG still shows a similar shape to  $z = 2$  with a peak at just below  $f_{H_2} \simeq 10^{-1}$  and a tail extending down towards the gas-depleted region. SIMBA, however, shows a bimodal structure with a higher peak at  $f_{H_2} \simeq 10^{-0.7}$  and a second peak at  $f_{H_2} \simeq 10^{-2}$  with a tail which extends significantly into the depleted region. This includes a few sources with ‘zero’ gas fractions which have been artificially scattered at  $f_{H_2} \simeq 10^{-3.5}$ . However, this lower peak almost disappears if we include the additional requirement of  $\lambda_{\text{Edd}} \geq 1\%$ , shown by the orange dotted line. This bimodality may be caused by the low-accretion feedback mode in SIMBA, which we discuss more in the Section 3.4.2. Due to the poorer statistics of EAGLE (see Section 3.2.4), it is difficult to precisely describe the distribution, but there appears to be a broad peak at  $f_{H_2} \simeq 10^{-2}$  and some sources in the gas-depleted region.

If we consider the observational samples at  $z = 0$ , we can see they both show a roughly symmetrical distribution around  $f_{H_2} \simeq 10^{-1.5}$  and are broadly consistent with these high- $L_{\text{bol}}$  AGN residing in gas-rich galaxies in the simulations. A similar result is found if we look at sSFR, as seen in Figure 3.10 in the Appendix. The observational data at  $z = 0$  seems to be consistent with a single high- $f_{H_2}$  peak, although we caution that the samples may not be complete for low- $f_{H_2}$  and there is limited overlap in  $L_{\text{bol}}$  with the sample from Zhuang et al. (2021) (see Section 3.4.2).

### 3.3.4 Gas-Depleted Fractions of AGN

To quantify the distributions seen in Figure 3.6 in more detail, we calculate the gas-depleted fraction,  $f_{\text{GD}}$ , of the AGN host galaxies across the whole galaxy sample based on the offset to the gas fraction main sequence in Tacconi et al. (2018). A galaxy is defined as quenched if it lies more than some distance,  $\Delta_{\text{MS}}$ , below this main sequence which is both  $M_{\star}$ - and  $z$ -dependent (see Section 3.2.5). We vary this value in the range  $\Delta_{\text{MS}} \in [1, 2]$  dex to test the sensitivity of our results to the quenching definition. The red shaded region in Figure 3.6 shows the variation in this quenching cut for an example galaxy of  $M_{\star} = 10^{10.5} M_{\odot}$ .



**Figure 3.7:** Gas-depleted fraction ( $f_{\text{GD}}$ ) for three AGN selection criteria: an Eddington ratio selection ( $\lambda_{\text{Edd}} > 1\%$ ), a luminosity cut ( $L_{\text{bol}} > 10^{44} \text{ ergs}^{-1}$ ), and the result if we combine the two selections. Each selection is also compared to a stellar-mass-matched control sample of non-AGN from each simulation, shown as the fainter points. The uncertainty regions are calculated from varying the offset from the main sequence ( $\Delta_{\text{MS}}$ ) for the gas-depleted definition (Section 3.2.5). The points represent the mean value for  $f_{\text{GD}}$ . At  $z = 0$ , all three simulations show very low gas depleted fractions compared to the control galaxies for the high- $\lambda_{\text{Edd}}$  selection, but SIMBA has a significantly higher depleted fraction for the high- $L_{\text{bol}}$  selection.

We present the results of this analysis in Figure 3.7. The range of gas-depletion definitions from  $\Delta_{\text{MS}}$  is represented by violin plots showing the probability density and the diamond points show the mean value. The depleted fraction is investigated for three AGN selections: an Eddington ratio selection ( $\lambda_{\text{Edd}} > 1\%$ ), a luminosity cut ( $L_{\text{bol}} > 10^{44} \text{ ergs}^{-1}$ ), and finally the result if we combine the two selections (high- $L_{\text{bol}}$  and high- $\lambda_{\text{Edd}}$ ). We also create a control sample for each AGN selection criteria by randomly sampling from the non-AGN population to match the stellar-mass distribution of AGN. This is repeated for each simulation with each AGN selection and this mass-matched control sample is plotted as the faint circular point in Figure 3.7 to provide a reference sample.

At  $z = 0$  we can see clear differences between the simulations. Taking a high- $\lambda_{\text{Edd}}$  cut, all three simulations show low gas-depleted fractions for the AGN ( $f_{\text{GD}} \lesssim 6\%$ ) which are always much lower than the control galaxy samples ( $18\% < f_{\text{GD}} < 24\%$ ). This is also demonstrated by the blue histograms in the right-hand panels of Figure 3.5 which show that very few high- $\lambda_{\text{Edd}}$  AGN reside in gas-depleted galaxies. However, for the high- $L_{\text{bol}}$  selection, we see a higher depletion fraction in SIMBA ( $8\% < f_{\text{GD}} < 31\%$ ) – due to the lower peak of the bimodal distribution seen in Figure 3.7 – and much lower depleted fractions in TNG and EAGLE. Combining these two selections has the effect of reducing

the depletion fraction in SIMBA to  $4\% < f_{\text{GD}} < 16\%$  (although this is still the highest of the simulations) but has little impact on EAGLE or TNG. The reasons for this dramatic change in the distribution in SIMBA are discussed further in Section 3.4.2.

At  $z = 2$  there are fewer differences between the simulations with all three showing  $f_{\text{GD}} < 7\%$  regardless of AGN selection. This is lower than the control samples although these show large variation between the simulations, especially in the  $\lambda_{\text{Edd}}$ -selection criteria. This is because in SIMBA and TNG, the high- $\lambda_{\text{Edd}}$  sources are more heavily biased towards lower  $M_*$  galaxies which have a higher quenched fraction at  $z = 2$ .

The results for sSFR (Figure 3.11) are again similar, with generally low quenched fractions across all the simulations regardless of AGN selection, except for the high fraction at  $z = 0$  for  $L_{\text{bol}}$ -selected AGN in SIMBA which is also seen in  $f_{\text{H}_2}$ . Again, the AGN-selected galaxies always have significantly lower quenched fractions than their mass-matched non-AGN counterparts.

In this section, we have shown that gas-depleted and quenched fractions in the simulations are lower for AGN than for a mass-matched control sample of galaxies, showing that AGN are preferentially found in gas-rich and star-forming galaxies. However, we have also shown that at  $z = 0$  there are differences between the simulations in their predictions for the gas fraction distributions when selecting using a luminosity cut ( $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$ ). In the Section 3.4.2, we will explore whether these differences are linked to the subgrid SMBH feedback models.

## 3.4 Discussion

Many observational studies seek evidence of AGN feedback, as invoked by simulations, by examining trends between star formation rates or molecular gas content and the AGN luminosity, or by comparing active to inactive galaxies (e.g. Page et al., 2012; Harrison et al., 2012; Scholtz et al., 2018; Rosario et al., 2018; Schulze et al., 2019; Shangguan & Ho, 2019; Florez et al., 2020; Circosta et al., 2021; Zhuang et al., 2021; Koss et al., 2021; Ji et al., 2022; Scholtz et al., 2021). However, these observational results have not led to a clear picture of negative feedback by AGN. Therefore, to help understand this, we have presented results on the relationship between AGN activity and host galaxy molecular gas fractions and specific star formation rates as directly predicted by three state-of-the-art cosmological simulations (TNG, EAGLE and SIMBA). Here we discuss these results in the context of the observations and possible future directions to test the theoretical framework of AGN feedback.

### 3.4.1 Cosmological models predict AGN preferentially live in gas rich, star-forming galaxies

The overall picture we have shown is that the simulations do not predict strong negative correlations between AGN luminosity and host galaxy properties of  $f_{\text{H}_2}$  or sSFR (Figure 3.4) in local galaxies ( $z = 0$ ) or at cosmic noon ( $z = 2$ ). Furthermore, by considering either an  $L_{\text{bol}}$ - or a  $\lambda_{\text{Edd}}$ -based AGN definition, we found that AGN are predicted to almost exclusively reside in gas-rich, star-forming galaxies (Figure 3.5 and Figure 3.9). More quantitatively, we found that the simulations generally show low gas-depleted fractions (mean fraction,  $f_{\text{GD}} < 10\%$  at  $z = 0$  for a combined AGN selection) in AGN hosts compared to a mass-matched control sample (Figure 3.6; Figure 3.7). None of these tests have shown what might be considered a clear sign of immediate quenching by luminous AGN.

We note that a small number of observational studies, mostly of high redshift AGN, have suggested luminous AGN may have lower gas fractions than matched non-active galaxies (e.g. Perna et al., 2018; Circosta et al., 2021; Bischetti et al., 2021). At face value, the

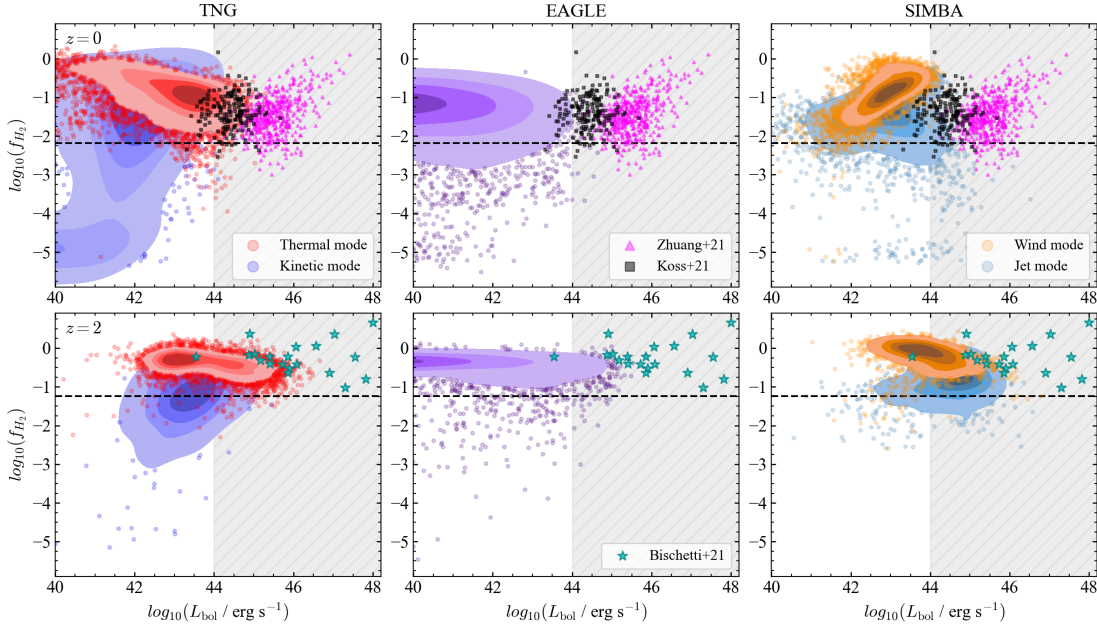
simulation predictions would not agree with these observations. However, throughout this study we have not attempted to account for such effects as the selection criteria of individual studies (e.g. using different wavelengths to select AGN), or different methods to calculate SFRs. These are all important for robust quantitative comparisons and may introduce some systematic differences, (e.g. Harrison, 2017; Ji et al., 2022), which may also explain the different correlation coefficients seen in the observational samples investigated here (see Figure 3.4). A thorough investigation is beyond the scope of this work. Nonetheless, in a broad qualitative sense, the simulations agree with most observational studies that find flat or positive correlations between  $L_{\text{bol}}$  and sSFR and  $f_{H_2}$ , and that AGN are preferentially hosted in gas-rich (Rosario et al., 2018; Kirkpatrick et al., 2019; Jarvis et al., 2020; Shangguan et al., 2020b; Zhuang et al., 2021; Koss et al., 2021; Valentino et al., 2021; Salvestrini et al., 2022) and star-forming (Rosario et al., 2012; Zhuang et al., 2019; Ji et al., 2022; Kim et al., 2022) galaxies. This result is also in agreement with previous studies of these simulations which mostly investigated star formation rates (e.g. McAlpine et al., 2017; Scholtz et al., 2018).

These results show that the observational result of finding luminous AGN located preferentially in star-forming or gas-rich galaxies is not in tension with negative AGN feedback, as implemented in these cosmological simulations (see also e.g. Thacker et al., 2014; Scholtz et al., 2018; Jackson et al., 2020). This is likely not surprising when noting that AGN and host galaxy star formation both require a gas reservoir. Indeed, if the gas is heated, destroyed or ejected by the feedback mechanism, this will not only affect the star formation, causing the galaxy to quench, but will also reduce the gas available to be accreted by the SMBH, limiting the accretion rate and the luminosity. This suggests that a realistic feedback process is one in which an observable impact on the galaxy-wide molecular gas and star formation properties is not observed in-situ with a highly luminous phase of AGN activity or, alternatively, a feedback process powered by low accretion rates, not associated with bolometrically luminous AGN (Harrison et al., 2012; Scholtz et al., 2018; Florez et al., 2020; Luo et al., 2021). Indeed, the simulations do invoke different prescriptions of AGN feedback (see Table 3.1) and observationally testing how realistic these are is an important step as we discuss in Section 3.4.2. Additionally, this does not rule out a more localised immediate effect of the AGN, as we discuss in Section 3.4.3.

Furthermore, even if there is no direct, rapid impact on cold, dense gas, outflows still deposit large amounts of energy into diffuse halo gas (e.g. Costa et al., 2014a, 2018b) or the CGM (e.g. Zinger et al., 2020), increasing the cooling time and preventing gas accretion back onto the galaxy, preventing future star formation. This effect operates over Gyr timescales so will be shaped by the cumulative energy injected by many outflows.

Another factor to consider when observationally searching for the impact of luminous AGN on their host galaxies, is the high variability of accretion rate and corresponding bolometric luminosity. The observed luminosity of an AGN – whether through optical, IR or X-ray measurements – is the key way we categorise galaxies into active or inactive. However, it only captures the *instantaneous* accretion rate of the SMBH and therefore is only a weak indicator of the cumulative energy injected into the host galaxy during feedback. It has been shown both observationally and in simulations that AGN accretion rates can vary by multiple orders of magnitude over timescales much shorter than typical star formation episodes (e.g. Novak et al., 2011; Hickox et al., 2014; Schawinski et al., 2015). Consequently, very little information on the total energy input by a SMBH on timescales relevant to star formation can be inferred from a single measurement of  $L_{\text{bol}}$  (Harrison, 2017).

Rather than looking at the instantaneous quantity of  $L_{\text{bol}}$ , we could use the mass of the black hole ( $M_{\text{BH}}$ ) as an integrated quantity of the SMBH’s accretion history. For example, Piotrowska et al. (2022) study the predictive power of different galaxy properties



**Figure 3.8:** The distribution of gas fraction with  $L_{\text{bol}}$  for each of the three simulations at  $z = 0$  (top row) and  $z = 2$  (bottom row). Shaded contours show the number density containing 90% of the subhalos and the extreme systems that lie outside are plotted as single points. The observational samples are also shown as coloured points. The hatched grey region to the right corresponds to a cut of  $L_{\text{bol}} > 10^{44} \text{ erg s}^{-1}$  and the black dashed line shows a rough gas-depletion definition for a galaxy with  $M_{\star} = 10^{10.5} M_{\odot}$  to guide the eye. Galaxies with unresolved gas fraction are scattered around  $f_{\text{H}_2} = 10^{-5}$ . The simulation data is coloured according to the accretion mode the AGN is in. In contrast to the other two simulations, at  $z = 0$  SIMBA predicts a significant fraction of high- $L_{\text{bol}}$  systems with depleted gas reservoirs, which are all in the ‘jet mode’ (grey data points).

on whether or not a galaxy is quenched, in the EAGLE, Illustris and TNG simulations, and in observations. They find that it is the mass of the SMBH, not its accretion rate, that is the strongest predictor for quenching. This corroborates the study of Terrazas et al. (2017), who found that quiescence correlates strongly with black hole mass both in an observational sample and the Illustris simulation, finding that the sSFR was a smoothly declining function of the specific black hole mass ( $M_{\text{BH}}/M_{\star}$ ), and the conclusion of Thomas et al. (2019) who found black hole mass to be a key indicator of quenching in the SIMBA simulation. Therefore, these authors conclude that black hole mass is a good indicator of quiescence as it provides a robust, model-independent feedback tracer which is insensitive to the mode of feedback. However, this still tells us little about the details of how or when the energy was deposited.

To test our understanding of the physical processes behind AGN feedback, we need to look for *different* predictions across the simulations, which all have different subgrid implementations of feedback mechanisms (Section 3.2.1; Table 3.1). Indeed our results have shown different predicted behaviours across the three simulations of galaxy properties as a function of AGN luminosity (Figures 3.4, 3.5, 3.6 and 3.7). We now discuss how these differences may relate to the subgrid feedback models and whether this could allow us to evaluate how realistic the implementations of AGN feedback are.

### 3.4.2 Different feedback modes lead to different simulation predictions at $z = 0$

The main differences between the simulations predictions are: (1) the strength of the correlations between sSFR and  $f_{\text{H}_2}$  with  $L_{\text{bol}}$  (Figure 3.4); (2) the distribution of AGN in the  $f_{\text{H}_2}$  and sSFR vs  $M_{\star}$  plane (Figures 3.5 and 3.9) and; (3) the varying gas-depleted

and quenched fractions for the highest  $L_{\text{bol}}$  and  $\lambda_{\text{Edd}}$  AGN (Figures 3.6 & 3.7 and 3.10 & 3.11). Interestingly, these different predictions show much greater variation at  $z = 0$  than at  $z = 2$ , which might seem surprising because free parameters in cosmological simulations are generally calibrated on the local universe (Schaye et al., 2015; Springel et al., 2018; Davé et al., 2019).

To explore these differences further, Figure 3.8 shows the gas fraction against the bolometric luminosity for the whole galaxy sample in the three simulations. We also split the data from TNG and SIMBA into their two primary feedback modes. The grey hatched region to the right of each panel shows our  $L_{\text{bol}} \geq 10^{44}$  erg s $^{-1}$  selection. To guide the eye, the black dotted line shows a rough gas-depleted definition, assuming a fiducial  $M_{\star} = 10^{10.5}$   $M_{\odot}$  galaxy (typical of our AGN hosts) and a tight cut of  $\Delta_{\text{MS}} = -1$ . The equivalent plot for sSFR can be found in the appendix (Figure 3.12).

At  $z = 0$ , almost all TNG galaxies in the thermal mode are found above the dashed line in Figure 3.8, i.e. they are not gas depleted, although they do seem to show a slight negative trend with  $L_{\text{bol}}$  (see also Figure 3.4). In contrast, gas depleted galaxies are almost exclusively found in the kinetic mode. There are very few high- $L_{\text{bol}}$ , low- $f_{\text{H}_2}$  sources and they are heavily outnumbered by the large number of bright quasars at  $f_{\text{H}_2} \simeq 10^{-1}$ , causing very low gas-depleted fractions, regardless of the AGN selection (Figure 3.7). We can see this switch in accretion mode more clearly in Figure 3.2: low mass SMBHs generally have high accretion rates and are in the thermal mode, growing in size until they reach the invoked accretion mode boundary (black dotted line) at which point they drop in luminosity and become kinetic mode sources. This creates a large population of kinetic mode sources at a black hole mass of  $M_{\text{BH}} \simeq 10^{8.1-8.4}$   $M_{\odot}$  and luminosity of  $L_{\text{bol}} \simeq 10^{41-43}$  erg s $^{-1}$ . This can also be seen in the downward ‘plume’ in Figure 3.5 at a stellar mass of  $M_{\star} \simeq 10^{10.5-11}$   $M_{\odot}$  and is linked to the onset of quenching in these massive galaxies.

Our findings for TNG corroborate the work of Terrazas et al. (2020) and Zinger et al. (2020) who both find a critical black hole mass threshold for quiescence of  $M_{\text{BH}} \simeq 10^{8.2}$   $M_{\odot}$ , above which the switch to the kinetic feedback mode efficiently pushed cold gas out of the galaxy, causing a sharp decrease in the amount of material available to cool and form stars (see also, Weinberger et al., 2018). From Figure 3.8 we can see that this gas-depleted population in the kinetic mode is already beginning to be established at  $z = 2$ , although the relatively low gas-depleted fraction in the overall galaxy population at this redshift suggests it has not had time to fully establish a large gas-depleted population. Interestingly, Figure 3.12 shows that, for sSFR, there is a quenched population in the kinetic mode at  $z = 2$ , although these sources make up a much lower proportion of the total galaxy population than at  $z = 0$ .

For both redshifts investigated, the results for EAGLE in Figure 3.8 show a similar picture; a flat trend in  $f_{\text{H}_2}$  with  $L_{\text{bol}}$ , as also seen in Figure 3.4. This similarity between the epochs is likely to be due to the single feedback mode; the dominant feedback mode is the same at both  $z = 0$  and  $z = 2$  and there isn’t a low- $\lambda_{\text{Edd}}$  mode kicking in at later times unlike in TNG and SIMBA. The consistently flat trends seen in EAGLE suggest that the instantaneous  $L_{\text{bol}}$  is not driving an immediate gas-depleting impact. This could be due to the fact that the energy input is ‘pulsed’ (following Booth & Schaye, 2009), which separates the accretion rate at any given time from its later cumulative impact on the galaxy. This is in contrast to the otherwise similar thermal mode in TNG where the energy input is continuous. Therefore, this may explain the weak negative correlations between  $f_{\text{H}_2}$  and  $L_{\text{bol}}$  seen in Figure 3.8 for the thermal mode in TNG, but which are flat in EAGLE.

If we consider the correlations seen in SIMBA (Figure 3.4), we notice that the direction of the trend between  $f_{H_2}$  and  $L_{\text{bol}}$  reverses between  $z = 2$  and  $z = 0$ . Figure 3.8 suggests that this is driven by the ‘wind’-mode sources (in orange) which dominate the correlations. However, the biggest difference from TNG is in the location of the ‘jet’-mode sources (in grey) which can have high accretion rates, and thus be bolometrically luminous, despite having relatively low Eddington ratios (see Figure 3.2). Because the ‘jet’ mode is more efficient at removing gas from the host galaxy, due to more powerful kinetic outflows and thermal heating (Thomas et al., 2019), we thus find high- $L_{\text{bol}}$  AGN in low- $f_{H_2}$  galaxies in SIMBA which we do not see in EAGLE or TNG. This population is the source of the lower peak seen in Figure 3.6 and causes the high quenched fraction in Figure 3.7 when we select AGN based on their  $L_{\text{bol}}$ . This also explains why this peak and the gas depleted fraction is reduced when we select for both luminosity and Eddington ratio – these sources making up this depleted population have low  $\lambda_{\text{Edd}}$ , despite their high  $L_{\text{bol}}$ . This could present an observational challenge as low- $\lambda_{\text{Edd}}$  AGN may be accreting inefficiently and thus be less bolometrically luminous than their high accretion rate may imply (see Section 3.2.4).

We have shown that the different predicted trends with bolometric luminosity (Figure 3.4) and the different distribution of galaxy’s gas fractions and sSFRs as a function of AGN luminosity (Figure 3.5 and 3.6) can be attributed, at least in part, to the different feedback modes invoked by the three different simulations. These differences are greater at lower redshifts, possibly because the low-accretion modes in TNG and SIMBA only become dominant at late cosmic times and it is these modes that are designed to have the greatest impact in galaxy quenching (Weinberger et al., 2017). EAGLE does not have such a mode-switch and thus is the most similar across the two epochs. This suggests that any observational work attempting to test these simulations with the types of experiments discussed would have more diagnostic power if performed on more local galaxies. However there may be other factors that act from  $z = 2$  to  $z = 0$  apart from AGN feedback that create such differences, for example, the switch from cold to hot halo gas that happens around cosmic noon (Dekel & Birnboim, 2006).

The varied subgrid implementations of feedback do produce some quantifiable differences between the simulations which may help to constrain the models, for example, the narrow  $L_{\text{bol}}$  and stellar mass range of the quenching kinetic-mode systems in TNG (Figures 3.5 and 3.8) around a characteristic black hole mass of  $M_{\text{BH}} \simeq 10^{8.1-8.4} M_{\odot}$  stands in contrast to EAGLE where there is no preferred  $L_{\text{bol}}$  for galaxies undergoing quenching. Additionally, SIMBA is the only simulation to predict quenched galaxies with high SMBH accretion rates.

### 3.4.3 Limitations and Outlook

One of the limitations of this study is the lack of comparable high- $L_{\text{bol}}$  AGN in these cosmological simulations with respect to the observational samples. As we can see in Figure 3.8, at  $z = 0$  there is only a small overlap with the data from Koss et al. (2021) ( $L_{\text{bol}} \simeq 10^{43.5-45} \text{ erg s}^{-1}$ ) and almost no overlap with the data from Zhuang et al. (2021) ( $L_{\text{bol}} \gtrsim 10^{44.5} \text{ erg s}^{-1}$ ). This is problematic, as some studies have suggested that quasar feedback is only effective in this extremely luminous regime (Costa et al., 2018b; Valentino et al., 2021). Such systems may exist in these simulations, but they are too short-lived to be captured in the limited box-size and coarsely time-sampled output snap-shot data of the current generation of cosmological simulations. To improve the sampling of high- $L_{\text{bol}}$  AGN, we could increase the volume of these big-box cosmological simulations, or run zoom-in simulations on massive halos where we can specify high accretion rates to test the effectiveness of quasar feedback in this regime. The larger volume of TNG300 may offer some improvements here; it contains 273 galaxies at  $z = 0$  with an AGN luminosity



of  $L_{\text{bol}} \geq 10^{45}$  erg s $^{-1}$  compared with only 10 in TNG100, but the brightest AGN is only  $L_{\text{bol}} = 10^{46.3}$  erg s $^{-1}$  which still leads to the extremely luminous end of the parameter space being too sparsely populated.

Additionally, improvements to the numerical treatment of the molecular phase of the ISM are required for a more realistic comparison to observations. This demands a combined treatment of complex chemistry and non-equilibrium low-temperature cooling (e.g. Richings et al., 2014a,b), dust, and radiative transfer (e.g. Rosdahl & Teyssier, 2015), which have not been modelled in the current generation of cosmological simulations. Although the post-processing models used in this work are useful for estimating the molecular gas mass in a galaxy, by not directly simulating the ISM we may be missing important feedback channels related to the direct effect of AGN on molecular gas (e.g. Bieri et al., 2017; Costa et al., 2018b; Costa et al., 2020). Furthermore, applying next-generation models of quasar winds (Costa et al., 2020), jets (Mukherjee et al., 2016, 2018; Talbot et al., 2021; Bourne & Sijacki, 2021; Mandal et al., 2021; Tanner & Weaver, 2022) and radiation pressure (Bieri et al., 2017; Costa et al., 2018a,b; Ishibashi et al., 2018) to cosmological simulations will make it possible to test more realistic AGN physical mechanisms by comparing to observations and consequently provide new insights into the true physical mechanisms that quench massive galaxies.

In this study, we have only looked at global galaxy properties and broadly found no connection between AGN luminosity and reduced gas content or star formation rates. However, building on the improved spatial resolution offered by modern instruments, many observers are looking at individual regions of galaxies to see where quenching or star formation might be occurring through high spatial-resolution and multi-wavelength observations (e.g. Al Yazeedi et al., 2021; Girdhar et al., 2022). Some studies have found that central regions are gas-depleted which could be linked to AGN activity (Ellison et al., 2021), although others have found no difference in the central molecular gas content of active and inactive galaxies (Rosario et al., 2018). There have also been studies looking at the radial profiles of galaxies in the simulations which found that quenching galaxies in EAGLE and Illustris have too centrally-concentrated star formation (Starkenburg et al., 2019), although Appleby et al. (2020) found that SIMBA did sufficiently suppress central star formation to match their observational sample. Investigating the timescale on which such central suppression occurs and seeing if this could be linked to a specific AGN feedback event could be an interesting avenue for future analysis of these simulations.

Finally, we note that there are several interesting populations of active galaxies, beyond just simply focusing on the most luminous, that may be interesting for understanding feedback and for testing simulations. In particular, galaxies undergoing the transition between star-forming and quiescent could display some unique properties in how the AGN is interacting with the host. Post-Starburst Galaxies (PSBs; also known as E+A galaxies) have been proposed as recently quenched galaxies, with evidence of AGN winds in some cases (Wild et al., 2016; Baron et al., 2018, 2022). Another population is Red Geyser galaxies (Cheung et al., 2016; Roy et al., 2021). These host low-luminosity AGN, but exhibit powerful outflows which could be responsible for their quiescence. Such systems could share many properties with the quenching population seen in TNG, where the SMBH has low accretion rates ( $L_{\text{bol}} \simeq 10^{42}$  erg s $^{-1}$ ) but through the kinetic mode feedback is able to launch outflows that act to quench the galaxy. Studying these complex and nuanced systems may provide another way of furthering our understanding of the intricate interaction between the AGN and host galaxies, allowing us to include ever more physically-motivated models in our theories of galaxy evolution.

### 3.5 Conclusions

We have investigated the predicted relationships between AGN activity and the host galaxy properties of molecular gas fraction ( $f_{H_2}$ ) and specific star formation rate (sSFR) in three contemporary cosmological simulations (TNG, EAGLE and SIMBA) for epochs of  $z = 0$  and  $z = 2$ . We were motivated by observational studies that have searched for signatures of AGN feedback by looking at trends between  $L_{\text{bol}}$  and galaxy properties, or that compared AGN to non-AGN. Several observational studies have suggested that the lack of negative trends between sSFR or  $f_{H_2}$  and AGN luminosity could indicate that AGN feedback may be ineffective. To investigate this, we followed the broad approaches taken in observational studies by looking at: (1) the correlation coefficients between the AGN luminosity and the host galaxy properties, after controlling for stellar mass; (2) the location of the AGN host galaxies compared to non-AGN hosts in the sSFR– and  $f_{H_2}–M_*$  plane; and (3) the gas-depleted and quenched fractions in AGN host galaxies. For all of these methods, we explore the effect of two different AGN selection definitions: a luminosity-based cut of  $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$ , and an Eddington ratio criterion of  $\lambda_{\text{Edd}} \geq 1\%$ . Our main results are:

1. In all three simulations we see no strong negative correlations between  $L_{\text{bol}}$  and sSFR or  $f_{H_2}$  at  $z = 0$  or  $z = 2$  for star-forming galaxies (Figure 3.4). Out of the 12 correlations investigated, only one shows even a weak negative correlation with the rest either being flat or positively correlated. This result is consistent whether we consider the whole star-forming population or only those with an Eddington ratio-selected AGN.
2. AGN are predicted to preferentially reside in gas-rich, high-sSFR galaxies in all three simulations at both redshifts investigated. This result holds whether we consider an  $L_{\text{bol}}$ - or an  $\lambda_{\text{Edd}}$ -based AGN selection (Figure 3.5).
3. Overall, the distributions in  $f_{H_2}$  for the highest  $L_{\text{bol}}$  and  $\lambda_{\text{Edd}}$  sources (Figure 3.6) show a low ( $\lesssim 7\%$ ) gas depleted fraction (Figure 3.7) which is always lower than for a mass-matched control sample of inactive galaxies. However, unlike the other simulations, SIMBA shows a bimodal distribution with a unique population of high- $L_{\text{bol}}$  galaxies located in gas-depleted and quenched galaxies.

All three simulations are heavily reliant on negative AGN feedback in order to reproduce realistic galaxies and yet we have shown that: (1) despite this negative feedback, we do not see any predictions of strong negative correlations between  $L_{\text{bol}}$  and sSFR or  $f_{H_2}$  and (2) AGN host galaxies are not predicted to contain significantly lower gas fractions or star formation activity compared to mass-matched non-AGN host galaxies. Therefore, this demonstrates that the result of finding bright AGN in star-forming, gas rich galaxies is not necessarily in contradiction with the presence of negative feedback. This could be due to various reasons, including timescale effects and the relative contribution of high or low accretion rate AGN to causing the quenching (Section 3.4.1).

However, the simulations do produce some different predictions for the trends of global galaxy properties with  $L_{\text{bol}}$  (Figure 3.4) and the distribution of galaxies in the  $L_{\text{bol}}–f_{H_2}$  plane (Figure 3.8). These can be attributed to the different subgrid accretion and feedback models implemented (Figures 3.2 & 3.8) which produce unique populations; for example, the highly-accreting, but low- $\lambda_{\text{Edd}}$  population in SIMBA, or the downwards ‘plume’ of kinetic-mode galaxies at low- $L_{\text{bol}}$  in TNG, which are both associated with gas-depleted galaxies. We also note that these differences are greater at  $z = 0$  compared to  $z = 2$ , perhaps suggesting that observations in local Universe will provide more diagnostic power for testing the different simulations.

Based on the results from this work, we suggest that focusing on population trends of only the most luminous AGN does not provide strong diagnostic power for testing AGN feedback. Instead, we recommend that observational efforts focus on testing some of the unusual quenching/quenched populations seen in the simulations and try to link these to observational samples. Complementary to this is observationally investigating AGN host galaxies in detail to understand the spatially-resolved physical processes of AGN feedback. Together, these kinds of studies could help to constrain which feedback models are the most realistic and help to inform future theoretical models of galaxy evolution.

### 3.6 Chapter Appendix: sSFR Plots

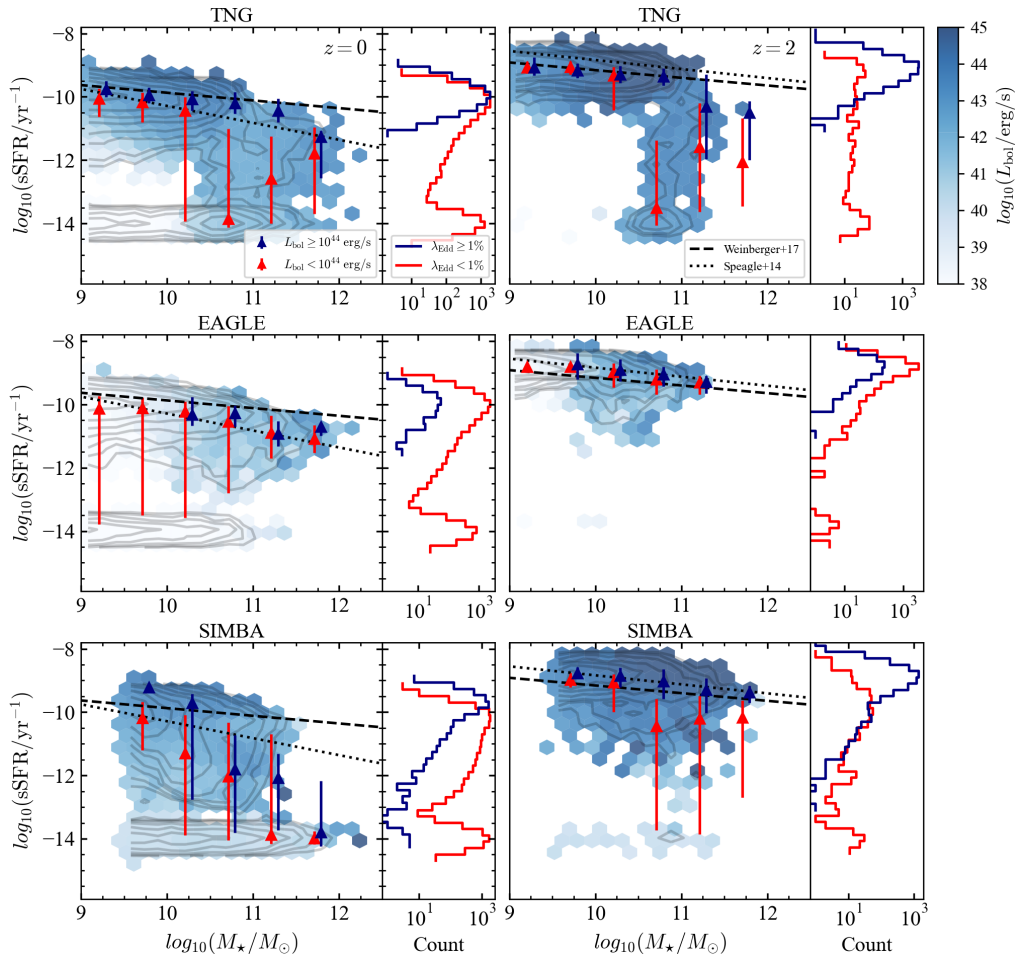
In this work we have mostly focussed on the simulations' predictions for the molecular gas as this is a common quantity observers study when looking at the impact of AGN feedback on galaxy quenching. For completeness, here we present the equivalent plots for the specific star formation rate.

Figure 3.9 shows the number density of galaxies in the  $\text{sSFR}-M_*$  plane with pixels coloured by the mean  $L_{\text{bol}}$ . We show the observational main sequence from Weinberger et al. (2017) and Speagle et al. (2014) as dashed and dotted lines respectively (note, we use Weinberger et al. 2017 as our model throughout the paper, although we find no qualitative difference between the two). The sSFR bimodality is clear at  $z = 0$  and is beginning to be established at  $z = 2$  in TNG and SIMBA. Both high- $L_{\text{bol}}$  and high- $\lambda_{\text{Edd}}$  AGN are preferentially located in high-sSFR galaxies.

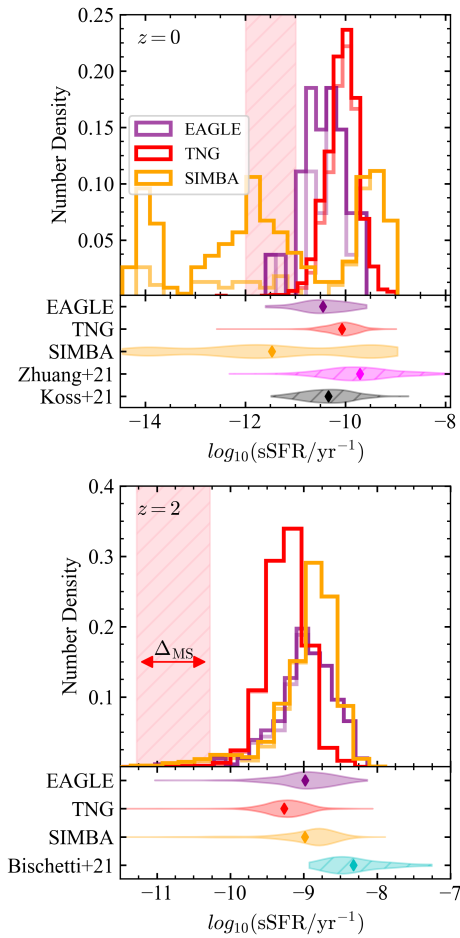
In Figure 3.10 we show the distribution in sSFR. At  $z = 0$  we see some clear differences between the simulations, with SIMBA showing a bimodal distribution. The lowest peak at  $\text{sSFR} = 10^{-14} \text{ yr}^{-1}$  is due to the galaxies which have unresolved star formation rates and should be considered part of the quiescent population at  $\text{sSFR} < 10^{-11} \text{ yr}^{-1}$ . At  $z = 2$ , the simulations all show high-sSFR with the peaks of all three simulations falling within 0.5 dex of each other.

Figure 3.11 shows the quenched fractions for three different AGN selections and their corresponding mass-matched control sample. We can see that the quenched fraction for EAGLE and TNG is low for all three AGN selections at  $z = 0$ , however, SIMBA shows a significantly higher value for the  $L_{\text{bol}}$  selection. This is due to the fact that SIMBA's 'jet' mode feedback can operate at high accretion rates and is efficient at quenching galaxies whilst still retaining such accretion rates (see Section 3.4.2).

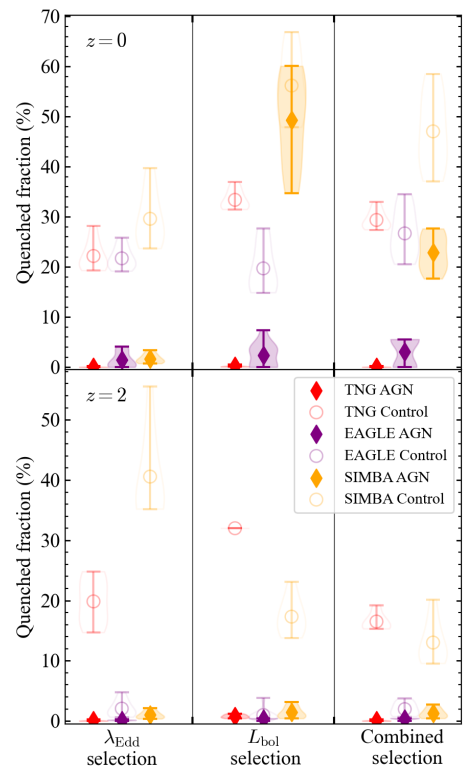
In Figure 3.12 we show the  $\text{sSFR}-L_{\text{bol}}$  plane for all three simulations. We also split TNG and SIMBA into their two primary feedback modes. The black dotted line shows a quenching definition for a  $M_* = 10^{10.5} M_\odot$  lying  $\Delta_{\text{MS}} = -1$  below the main sequence. We can see that in TNG, almost all of the thermal mode (in red) sources lie above this line and that the quenched population is exclusively composed of kinetic mode sources (blue). Neither TNG or EAGLE have significant numbers of AGN in both the quenched region and the high- $L_{\text{bol}}$  region. In SIMBA, however, some quenched galaxies do have high luminosity and these are all in the low- $\lambda_{\text{Edd}}$  'jet' mode.



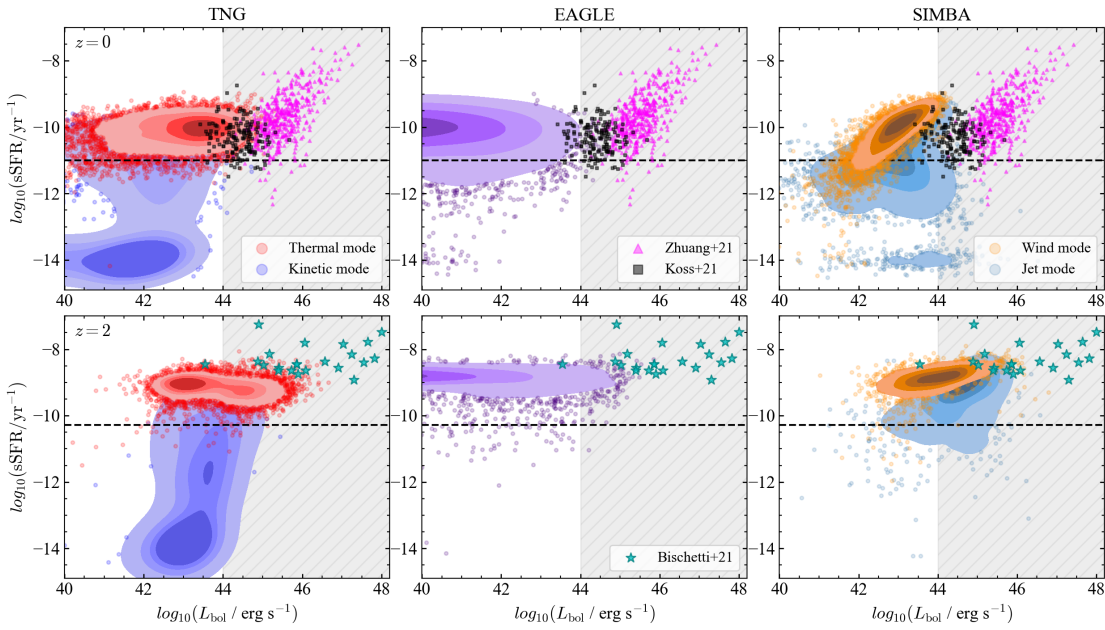
**Figure 3.9:** Galaxies in the  $\text{sSFR}-M_*$  plane coloured according to mean  $L_{\text{bol}}$  within each pixel. The triangular points show the median and 16<sup>th</sup>-84<sup>th</sup> percentiles of the  $\text{sSFR}$  in mass bins of 0.5 dex, split by  $L_{\text{bol}}$  in AGN (blue) and non-AGN (red). The panels on the right show the logarithmic number density for AGN selected by Eddington ratio - AGN are shown in blue and non-AGN in red.



**Figure 3.10:** The distribution of sSFR for luminous AGN in the three simulations. Solid lines show all AGN with  $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$  and dotted lines show the additional requirement of having  $\lambda_{\text{Edd}} \geq 10^{1\%}$ . The bottom panels show the smoothed density plot for each simulation sample compared to the comparison observational sample.



**Figure 3.11:** Quenched fraction for sSFR. We use three AGN selections:  $L_{\text{bol}} \geq 10^{44} \text{ erg s}^{-1}$ ,  $\lambda_{\text{Edd}} \geq 1\%$ , and a selection combining the two. We also compare to a stellar-mass-matched control sample taken for each AGN selection criteria.



**Figure 3.12:** sSFR against  $L_{\text{bol}}$  for the simulations at  $z = 0$  and  $z = 2$ . We split TNG and SIMBA into their two main modes of feedback and also plot the observational sample. The blue shaded region represents our high- $L_{\text{bol}}$  selection and the black dotted line shows a rough quenching definition based on a fiducial galaxy of  $M_{\star} = 10^{10.5} M_{\odot}$  and a quenching distance from the main sequence of  $\Delta_{\text{MS}} = -1$ .

4

# Multiphase Outflows from AGN Winds

AGN-driven outflows in clumpy media: multiphase  
structure and scaling relations

**Samuel Ruthven Ward**; T.Costa; C.M.Harrison; V.Mainieri

Accepted MNRAS (2024)

---

Small-scale winds driven from accretion discs surrounding active galactic nuclei (AGN) are expected to launch kpc-scale outflows into their host galaxies. However, the ways in which the structure of the interstellar medium (ISM) affects the multiphase content and impact of the outflow remains uncertain. We present a series of numerical experiments featuring a realistic small-scale AGN wind with velocity  $5 \times 10^3 - 10^4 \text{ km s}^{-1}$  interacting with an isolated galaxy disc with a manually-controlled clumpy ISM, followed at sub-pc resolution. Our simulations are performed with AREPO and probe a wide range of AGN luminosities ( $L_{\text{bol}} = 10^{43-47} \text{ erg s}^{-1}$ ) and ISM substructures. In homogeneous discs, the AGN wind sweeps up an outflowing, cooling shell, where the emerging cold phase dominates the mass and kinetic energy budgets, reaching a momentum flux  $\dot{p} \approx 7 L/c$ . However, when the ISM is clumpy, outflow properties are profoundly different. They contain small, long-lived ( $\gtrsim 5 \text{ Myr}$ ), cold ( $T \lesssim 10^{4.5} \text{ K}$ ) cloudlets entrained in the faster, hot outflow phase, which are only present in the outflow if radiative cooling is included in the simulation. While the cold phase dominates the mass of the outflow, most of the kinetic luminosity is now carried by a tenuous, hot phase with  $T \gtrsim 10^7 \text{ K}$ . While the hot phase reaches momentum fluxes  $\dot{p} \approx (1 - 5) L/c$ , energy-driven bubbles couple to the cold phase inefficiently, producing modest momentum fluxes  $\dot{p} \lesssim L/c$  in the fast-outflowing cold gas. These low momentum fluxes could lead to the outflows being misclassified as momentum-driven using common observational diagnostics. We also show predictions for scaling relations between outflow properties and AGN luminosity and discuss the challenges in constraining outflow driving mechanisms and kinetic coupling efficiencies using observed quantities.

---



## 4.1 Introduction

At the centres of massive galaxies lie supermassive black holes (Kormendy & Ho, 2013). These objects have masses ranging from  $\sim 10^6 - 10^{10} M_{\odot}$ , and during periods of intense gas accretion can ‘light up’ to become Active Galactic Nuclei (AGN). Over the lifetime of a supermassive black hole, more than the binding energy of the galaxy can be released, giving AGN the potential to influence the fate of their host galaxies; a process known as AGN feedback (e.g. Fabian, 2012; King & Pounds, 2015). AGN feedback is an essential component in all contemporary theoretical models and simulations of galaxy formation (e.g., Hirschmann et al., 2014; Schaye et al., 2015; Weinberger et al., 2018; Davé et al., 2019; Dubois et al., 2021; Wellons et al., 2023). However, direct observational evidence of this feedback on a population scale is lacking, with studies finding that AGN-hosting galaxies are no more quenched or gas-depleted than their inactive counterparts (Rosario et al. 2013, 2018; Kirkpatrick et al. 2019; Jarvis et al. 2020; Shangguan et al. 2020a; Valentino et al. 2021; Koss et al. 2021; Zhuang et al. 2021; Ji et al. 2022; Kim et al. 2022; Frias Castillo et al. 2024, but see some counterexamples at higher redshift: Perna et al. 2018; Bischetti et al. 2021; Circosta et al. 2021; Bertola et al. 2024). As shown in previous work (Ward et al. 2022, see also Scholtz et al. 2018; Piotrowska et al. 2022), this apparent tension is not in contradiction with models that rely on AGN feedback to quench galaxies. It instead highlights the difficulty in studying this problem due to the vast range in timescales and distances involved (Harrison, 2017), and motivates further work on the impact of AGN feedback on the host galaxy.

The power of an AGN can be coupled through jets (e.g. Mukherjee et al., 2016; Bourne & Yang, 2023), accretion disc winds (Silk & Rees, 1998; King, 2003; Faucher-Giguère & Quataert, 2012; Almeida et al., 2023) and/or direct radiation pressure (Thompson et al., 2015; Ishibashi et al., 2018; Costa et al., 2018a), all of which can drive kpc-scale outflows in the host galaxy, transferring mass, momentum and energy to the interstellar medium (ISM) and circumgalactic medium (CGM). Studying the interaction between these AGN-driven outflows and the multiphase ISM in galaxies is therefore critical in understanding the role AGN play in galaxy evolution (Harrison & Ramos Almeida, 2024). To this end, the last decade has seen an explosion of observational work studying the multiphase nature of AGN outflows. For example, radio and sub-mm observatories have allowed measurements of the coldest phase of the gas (often using CO, [C II] or H I transitions; e.g., Morganti et al. 2005; Cicone et al. 2014; Fluetsch et al. 2019; Veilleux et al. 2020; Lamperti et al. 2022; Ramos Almeida et al. 2022; Girdhar et al. 2024); and X-ray–NIR spectroscopic data has enabled measurements of high velocity disk winds (e.g., King & Pounds, 2015; Gofford et al., 2015; Chartas et al., 2021; Matzeu et al., 2023), the ionised or atomic phases in the ISM (e.g., using [O III] or Na I D; Rupke & Veilleux 2013; Harrison et al. 2014; Zakamska & Greene 2014; Molyneux et al. 2019; Musiimenta et al. 2023), and X-ray imaging/spectroscopic studies for the hottest gas on large scales (e.g., Lansbury et al., 2018; Longinotti et al., 2023). Such studies often seek to quantify the mass outflow rate, momentum rate and kinetic coupling efficiency of the outflows. These quantities are often used to infer if the outflows are momentum- or energy-driven, and if the outflows are considered energetic enough to impact the host (e.g., Cicone et al., 2014; Lamperti et al., 2022; Riffel et al., 2023).

Some studies have identified scaling relations between outflow properties, such as the mass outflow rate or kinetic energy coupling efficiency, and the AGN luminosity (e.g., Fiore et al., 2017; Leung et al., 2019; Davies et al., 2020; Musiimenta et al., 2023). However, these trends may partly be driven by selection effects, such as the choice to target CO-bright systems with known outflows in other phases (see discussion in Ramos Almeida et al. 2022; Harrison & Ramos Almeida 2024). These efforts are further hampered by

observational challenges; for example, the spatial extent of the outflow can be challenging to measure, especially if the resolution of the observations is only marginally better than the galaxy size; the velocity of the outflow is difficult to disentangle due to systemic galaxy motion and projection effects; and poorly constrained conversion factors and electron densities add large uncertainties to the outflow mass estimates (Husemann et al., 2016; Tadhunter et al., 2018; Harrison et al., 2018; Veilleux et al., 2020). Furthermore, if only one gas phase is measured, a significant amount of the outflow mass or energy will be missed, potentially changing the resulting conclusion on the outflow driving mechanism. Therefore, an unbiased, multiwavelength approach is essential (Cicone et al., 2018).

To help interpret the results of these observations, comparisons to predictions from simulations need to be made (e.g., Meenakshi et al., 2022a,b). In particular, an outstanding task is to establish the theoretical expectations of scaling relations and examine how these compare to the observed trends. However, accurately simulating the interaction between the ISM and AGN-driven outflows is a complex numerical task as the dynamic range of the problem is vast, from outflows that can reach 100s kpc to sub-pc-scale structures in the ISM. One of the major challenges is self-consistently modelling realistic ISM conditions. The complex structure of the ISM is maintained by a range of physical processes, such as supernova and star formation feedback (Gent et al., 2013), dust formation and destruction (Hirashita & Murga, 2020; Kirchschrager et al., 2022), molecular chemistry and cooling (Richings & Faucher-Giguère, 2018a,b), and cosmic ray feedback (Ruszkowski & Pfrommer, 2023). Simulating all these effects, especially on a galaxy scale, is computationally challenging, both in terms of the complex physical processes involved and the resolution requirements. Therefore, most galaxy and cosmological-scale simulations use an effective equation of state model (eEOS; e.g., Springel & Hernquist 2003). This packages up the unresolved ISM physics such as star formation and feedback into a subgrid model, allowing galaxies to be simulated without resolving all the physical processes involved (e.g., Vogelsberger et al., 2013). However, this simplified model does not accurately reproduce the spatial or multiphase structure of the ISM, which may play a significant role in the propagation of outflows through the galaxy. To overcome the inability to resolve the ISM substructure, some studies manually set the spatial distribution of the gas, creating a clumpy, two-phase media composed of fractally-distributed cold gas clouds surrounded by a tenuous hot phase. This setup also allows the effect of wind-ISM interactions to be studied in inhomogeneous environments, unlike analytic models, which often assume spherically-symmetric media, and eEOS models, which do not capture small-scale ISM structure. This method has been used by previous studies to investigate the effect of jets (Wagner & Bicknell, 2011; Mukherjee et al., 2016; Tanner & Weaver, 2022), winds from AGN (Wagner et al., 2013) and starbursts (Cooper et al., 2008), radiation pressure (Bieri et al., 2017), and shocks (Banda-Barragán et al., 2020, 2021) in clumpy media.

Another open question is how cold gas becomes entrained in a hot outflow. The crushing time of cold clouds is shorter than their entrainment timescale (Klein et al., 1994; Zhang et al., 2017). Yet cold gas clouds are observed in outflows travelling at 100s km s<sup>-1</sup> (Di Teodoro et al., 2019; Veilleux et al., 2020). Recent simulations of the interaction between single clouds and a hot wind have shown that the clouds can survive, and even increase in mass, due to efficient radiative cooling at the mixing boundary between the cloud and the wind (Gronke & Oh, 2018, 2020a; Fielding et al., 2020), although this effect is also sensitive to the density structure of the clouds themselves (Banda-Barragán et al., 2019; Mandal et al., 2024). Therefore, understanding how these small clouds form, and resolving their interaction with the hot wind, is crucial in understanding the cold outflows seen in observations.

In this work, we set up a series of computational experiments to study the interaction between a clumpy ISM and a fast, small-scale AGN accretion-disc wind, modelled following [Costa et al. \(2020\)](#). This paper is the first in our new project "AGN in Clumpy DisCs" (ACDC). The main questions we set out to answer are:

1. What are the properties of outflows formed from wind-ISM interactions?
2. How does having a clumpy ISM affect the properties of multiphase outflows compared to a smooth medium?
3. How can observed outflow measurements, such as momentum rate and kinetic energy coupling efficiency, constrain AGN feedback mechanisms?
4. Do we expect scaling relations between outflow properties and AGN luminosity?

We structure our paper as follows: in Section 4.2 we present our experimental setup, including our method for creating a clumpy ISM (Section 4.2.3) and our AGN wind model (Section 4.2.4). We split our results into two parts; in Section 4.3 we show the formation of multiphase outflows in our simulations and quantify their properties. In Section 4.4 we then discuss the implications of our findings for observational studies. We finish by presenting our conclusions and discussing directions for future studies in Section 4.5.


## 4.2 Numerical Simulations

To investigate the effect of a clumpy ISM on the propagation of AGN wind-driven outflows, we set up a series of controlled experiments. These feature an idealised galaxy disc within a uniform gaseous halo where the structure of gas in the disc has been manually set to mimic the observed fractal substructure of the ISM (similar to [Wagner & Bicknell 2011](#); [Wagner et al. 2012, 2013](#); [Mukherjee et al. 2016](#); [Bieri et al. 2017](#); [Mukherjee et al. 2018](#); [Tanner & Weaver 2022](#)). An AGN is placed at the centre of the disc and a spherical wind solution is set up via BOLA ([Costa et al. 2020](#); see Section 4.2.4). Using this setup, we can investigate the interaction between an AGN wind and a spatially resolved ISM. By varying the ISM structure and the properties of the AGN wind, we can also evaluate how these affect the resulting outflow. This is motivated both by previous numerical studies, which found that the initial sizes of clouds in the ISM was a critical parameter in determining feedback efficiency ([Wagner et al., 2012](#)), and by observations which suggest that outflow properties correlate strongly with AGN luminosity (e.g., [Fiore et al., 2017](#)).

In Section 4.2.1 we present the hydrodynamic code used, before explaining how the galaxy disc (Section 4.2.2), ISM structure (Section 4.2.3) and AGN wind (Section 4.2.4) were set up. Finally, we describe the suite of simulations explored in Section 4.2.5 and we explain the calculations used to extract outflow properties in Section 4.2.6.

### 4.2.1 Hydrodynamic Code

Our simulations are performed using the hydrodynamic code AREPO ([Springel, 2010b](#)). This uses an unstructured, Voronoi mesh that moves, in a quasi-Lagrangian fashion, with the fluid. This maintains Galilean invariance whilst also providing excellent shock capturing without the need for artificial viscosity, as in particle-based codes. Primitive variables are linearly reconstructed within each cell, providing second-order spatial accuracy. Extrapolated values at cell boundaries are used to compute hydrodynamic fluxes using an exact Riemann solver ([Pakmor et al., 2016](#)). Cells are refined or de-refined according to a pre-determined refinement criterion. By default, this involves (de)refining to keep cells within a factor of two from a target mass. This results in high spatial resolution in regions



**Figure 4.1:** The initial conditions for our different simulations, showing a face-on slice through our disc. We explore a range of largest cloud sizes of  $\lambda_{\max}=40$  pc,  $\lambda_{\max}=170$  pc and  $\lambda_{\max}=330$  pc, and a smooth disc for comparison. The disc gas mass is  $M_{\text{disc}} = 1.4 \times 10^9 M_{\odot}$  and the ambient background gas is in pressure equilibrium with the clumps, and has a temperature of  $T=10^7$  K. Presented on the left is a schematic of the BOLA boundary structure (Costa et al., 2020), showing the two layers of AREPO cells used to launch the AGN wind.

of high density, allowing us to resolve small-scale structures such as the clumpy distribution of the ISM on pc-scales. For our fiducial-resolution simulations, we use a target gas mass of  $M_{\text{target}} = 100 M_{\odot}$  which gives us a maximum spatial resolution of around  $d_{\text{cell}} = 1$  pc (see Figure 4.12). We discuss the numerical convergence of our results in Appendix 4.6. We also extend the refinement criteria to increase the resolution in the AGN wind (see Section 4.2.4).

### 4.2.2 Galaxy Setup

In this study, we analyse an idealised galaxy disc located within a hot halo. We use a box size of  $L_{\text{box}} = 20$  kpc with periodic boundary conditions. The simulations are performed for a period of time shorter than the outflow crossing time ( $t=5$  Myr for our fiducial simulations). The disc has a diameter of 4 kpc and a thickness of 1 kpc (following similar setups in Mukherjee et al. 2018; Tanner & Weaver 2022) and tapers at the edges to smooth the interface between the background. The mean gas number density in the disc is set to  $\langle n_0 \rangle = 5 \text{ cm}^{-3}$ , resulting in a disc gas mass of  $M_{\text{disc}} = 1.4 \times 10^9 M_{\odot}$ . The background is set in pressure equilibrium with the disc, giving temperatures of  $T_{0,\text{disc}} = 10^4$  K and  $T_{0,\text{bkg}} = 10^7$  K and a number density of  $n_{\text{bkg}} = 10^{-2} \text{ cm}^{-3}$ . In our fiducial simulations, we also include the standard AREPO prescription for primordial cooling, excluding a UV background. We investigate the effect of cooling in Section 4.3.2. We also investigate the effect of altering the disc mass by reducing the disc height (to  $h = 0.5$  kpc) and reducing the mean gas density (to  $\langle n_0 \rangle = 2.5 \text{ cm}^{-3}$ ). The sensitivity of our results to these changes is explored in Section 4.3.3.

### 4.2.3 Setting the ISM Structure

In this study, we take a controlled experiment approach to investigating the interaction between an AGN wind and a clumpy ISM. We manually create a two-phase ISM of cold clumps, arranged in a fractal distribution, surrounded by a hot, diffuse phase. Thanks to our high resolution, this allows us to spatially resolve the ISM structure to investigate what effect this has on the AGN wind moving through the galaxy, and vice-versa.

To initialise our ISM structure, we make use of the PyFC<sup>1</sup> Python package (Wagner et al., 2012). This generates a random, three-dimensional scalar field from a given probability distribution function with a fractal spatial correlation, a method introduced by Sutherland & Bicknell (2007) based on a scheme developed for terrestrial clouds by Lewis & Austin (2002). We use a log-normal distribution function, with width  $\sigma = \sqrt{5}$ , and a Kolmogorov power-law spectrum ( $\beta = -\frac{5}{3}$ ), motivated by observed ISM density

<sup>1</sup><https://pypi.org/project/pyFC/>

distributions (Fischera et al., 2003). The resulting fractal structure is parameterised by the lower wavenumber cutoff,  $k_{\min}$ , which represents the largest correlated spatial scale. This can be related to the average largest cloud size,  $\lambda_{\max}$ , by

$$\lambda_{\max} = \frac{L}{2k_{\min}}, \quad (4.1)$$

where  $L$  is the box length of the fractal cube being generated. Low values of  $k_{\min}$  represent large initial cloud sizes, and high values represent small initial cloud sizes. The maximum value of  $k_{\min}$  that can be used is set by the Nyquist limit which depends on the simulation resolution. In our study, we use a range of minimum wavenumber values, from  $6 \text{ kpc}^{-1} \leq k_{\min} \leq 50 \text{ kpc}^{-1}$  which corresponds to average largest cloud sizes of  $40 \text{ pc} \leq \lambda_{\max} \leq 333 \text{ pc}$ , giving a wide range of ISM conditions. The resulting fractal cube is then cropped into the desired disc shape and the densities scaled to the mean disc density. To generate porosity in the ISM, cells above a temperature threshold of  $T_{\text{crit}} = 3 \times 10^4 \text{ K}$  are considered thermally unstable and replaced by gas of the same temperature as the background, in pressure equilibrium with the cold clumps (following Sutherland & Bicknell, 2007; Cooper et al., 2008). This generates a two-phase medium of fractally-distributed cold clumps and hot, diffuse gas. The edges of the disc ( $> 400 \text{ pc}$  above the midplane) are then tapered with a *tanh* profile (Tanner & Weaver, 2022) to create a smoother boundary with the background. We note that using this method maintains a constant disc mass regardless of the value of  $k_{\min}$ .

As this method was designed for grid-based (constant cell volume) codes, a further step is required to set-up the disc for use with AREPO which uses the constant mass approach. The Voronoi grid is first evolved with hydro fluxes disabled, which allows the grid to regularise (Springel, 2010b) and to refine (de-refine) in regions of high (low) density. Once the total number of cells has converged, the disc retains the desired fractal properties, but now with a quasi-Lagrangian setup. The resulting initial conditions are shown in Figure 4.1 for a range of clumping factors, alongside our smooth disc case. We can see that the initial clumps have a range of sizes and densities, from  $n \approx 10^{1-3} \text{ cm}^{-3}$ , roughly in line with what is expected for cold H I clouds (Cox, 2005).

The advantage of using this manually-set ISM structure is it allows us to conduct controlled experiments of the interaction between an outflow and a spatially-resolved ISM without attempting to create an ISM structure from first-principles, which would be highly dependent on a range of loosely-constrained subgrid models. It also allows very high spatial resolution to be achieved to study the small-scale structure of the ISM and its resulting fragmentation. However, this method means our simulations are very idealised and only run for short timescales ( $\leq 5 \text{ Myr}$ ), both due to boxsize constraints and cooling losses, which cause the fractal structure to diffuse on longer timescales. We also neglect effects such as self-gravity, turbulence and star formation. We note that some studies using this method do include a static gravitational potential (e.g., Mukherjee et al., 2016; Tanner & Weaver, 2022) and an initial gas velocity dispersion (Mukherjee et al., 2018). This has the effect of changing the initial clumps to a filamentary structure and smoothing out the density contrast between the clouds and the porous gaps to create conditions more similar to a turbulent medium. However, this makes it harder to control the porosity of the ISM, so we follow Wagner et al. (2012); Bieri et al. (2017) and neglect gravity to focus solely on the interaction between the quasar wind and our manually-set ISM structure. We note that the outflow crossing time of the fastest-moving gas is shorter than the free-fall time of the galaxy ( $\approx 50 \text{ Myr}$ ) meaning it is unlikely to be significantly affected by gravity (see Wagner et al., 2012). However, the slower-moving tail of the outflow is most likely to be affected by the gravity of the galaxy, so these velocities should be seen as an upper limit.

#### 4.2.4 AGN Wind Model

To model an AGN wind, we employ the BOLA (BOundary Layer for AGN) model introduced in [Costa et al. \(2020\)](#). A sphere of cells are fixed in place at the location of the black hole. This is composed of two layers: an ‘inner layer’, which is excluded from the hydrodynamic calculations and used only to define the fixed boundary, and an ‘outer layer’, which is free to evolve hydrodynamically. Both layers are discretised following a HEALPIX tessellation into  $12n_{\text{side}}^2$  pixels of equal surface area ([Górski et al., 2005](#)). In this study, we use a value of  $n_{\text{side}} = 8$ . We set the radius of the inner spherical boundary as  $r_{\text{sp}} = 10$  pc. The mass, momentum and energy flux is set at the interface between the two layers and then communicated from the outer layer into the surrounding gas. BOLA makes it possible to model feedback processes operating below the resolution scale through appropriate choice of boundary conditions. Here, we adopt boundary conditions for a spherical, ultra-fast outflow (UFO) as in [Costa et al. \(2020\)](#). These boundary conditions directly reproduce the UFO conditions investigated in detail in [King \(2003\)](#); [Faucher-Giguère & Quataert \(2012\)](#); [Costa et al. \(2014b\)](#) which lead to energy-driven bubbles that can produce strong large-scale feedback. In addition, a passive scalar is injected across the boundary which is then advected along with the injected wind. This allows BOLA to modify the AREPO refinement scheme to increase the resolution of wind cells by boosting the resolution in cells with a high density of wind fluid. This helps to reduce the problem that the wind itself is poorly-resolved due to its low density. We decrease the target mass of wind cells by a factor of 10.

The main free parameters for this model are the AGN luminosity ( $L_{\text{AGN}}$ ), wind velocity ( $v_{\text{AGN}}$ ), temperature ( $T_{\text{AGN}}$ ) and momentum boost factor  $\tau = \dot{p}/(L_{\text{AGN}}/c)$  (all at injection scale). We use a momentum boost factor of  $\tau = 1$  and an initial wind temperature of  $T_{\text{AGN}} = 10^6$  K. These parameters ensure that the wind remains highly supersonic out to the free-expansion radius, and that the pressure contribution to the kinetic luminosity and momentum flux is marginal (see [King 2003](#); [Faucher-Giguère & Quataert 2012](#)). Furthermore, the choice of initial wind temperature plays a negligible role, as the wind is quickly shocked to higher temperatures. We take fiducial values for the AGN luminosity and wind velocity of  $L_{\text{bol}}=10^{45}$  erg s $^{-1}$  and  $v_{\text{AGN}}=10,000$  km s $^{-1}$  but explore a range of parameters in our simulation, as shown in Table 4.1. The initial kinetic luminosity injected is given by

$$\dot{E}_w = \frac{\tau\beta}{2}L_{\text{AGN}}, \quad (4.2)$$

where  $\beta = v_{\text{AGN}}/c$ . For our chosen values of  $\tau = 1$  and  $v_{\text{AGN}}=10,000$  km s $^{-1}$ , this gives an initial energy coupling efficiency of  $\dot{E}_w/L_{\text{AGN}} = 1.7\%$ , in the fiducial case. These values are consistent with studies of ultra-fast outflows detected at small-scales in X-rays (e.g., [Gofford et al., 2015](#); [Matzeu et al., 2023](#)). It is important to note that these injected values are on the scale of the AGN accretion disc (sub-pc) and are therefore not necessarily comparable to the resulting large-scale outflows on kpc-scales ([Harrison et al., 2018](#); [Costa et al., 2020](#)). We discuss the derived momentum boost and kinetic coupling efficiency of the large-scale outflow in Section 4.4.2. In this study we don’t explore gas accretion onto the AGN, and our AGN maintains a constant luminosity throughout.

#### 4.2.5 Simulation Suite

Table 4.1 shows the range of simulation parameters investigated in this study. For clarity, we name the ones most commonly discussed in the paper. In particular, we highlight ‘Medium clumps’ as our fiducial simulation: this has a clumping parameter of  $k_{\text{min}}=12$  ( $\lambda_{\text{max}}=170$  pc), an AGN luminosity of  $L_{\text{bol}}=10^{45}$  erg s $^{-1}$  and includes radiative cooling. The simulation with the same parameters, but without cooling, is called ‘No Cooling’. For

Simulation Name	$M_{\text{target}}$ [ $M_{\odot}$ ]	$k_{\text{min}}$ [ $\text{kpc}^{-1}$ ]	$\lambda_{\text{max}}$ [pc]	$h$ [kpc]	$\langle n_0 \rangle$ [ $\text{cm}^{-3}$ ]	Cooling	$L_{\text{AGN}}$ [ $\text{erg s}^{-1}$ ]	$v_{\text{AGN}}$ [ $\text{km s}^{-1}$ ]
Small clumps	100	50	40	1	5	Y	$10^{45}$	10000
<b>Medium clumps (Fiducial)</b>	<b>100</b>	<b>12</b>	<b>170</b>	<b>1</b>	<b>5</b>	<b>Y</b>	<b><math>10^{45}</math></b>	<b>10000</b>
Large clumps	100	6	330	1	5	Y	$10^{45}$	10000
Smooth	100	-	-	1	5	Y	$10^{46}$	10000
No cooling	100	12	170	1	5	N	$10^{45}$	10000
Slow wind	100	12	170	1	5	Y	$10^{45}$	5000
Thin disc	100	12	170	0.5	5	Y	$10^{45}$	10000
Low density	100	12	170	1	2.5	Y	$10^{46}$	10000
L43	100	12	170	1	5	Y	$10^{43}$	10000
L44	100	12	170	1	5	Y	$10^{44}$	10000
L46	100	12	170	1	5	Y	$10^{46}$	10000
L47	100	12	170	1	5	Y	$10^{47}$	10000
No AGN	100	12	170	1	5	Y	-	-
Low-res	1000	12	170	1	5	Y	$10^{45}$	10000
Low-res, no cooling	1000	12	170	1	5	N	$10^{45}$	10000
High-res	10	12	170	1	5	Y	$10^{45}$	10000

**Table 4.1:** The suite of simulations explored in this work. In bold is our fiducial simulation. The columns represent: 1) The nickname for each run; 2) The target mass resolution; 3) The wavenumber of the initial clump sizes; 4) The average maximum cloud size of the clumps; 5) The height of the disc; 6) The initial mean number density of the clumps; 7) Whether radiative cooling is included; 8) The luminosity of the AGN; and 9) The velocity of the AGN wind at injection.

comparison, we also analyse the resulting outflow from an AGN place in a homogeneous disc, with a constant density of  $n_0 = 5 \text{ cm}^{-3}$  ('Smooth'). We show the results for the smooth disc in Section 4.3.1 and compare the phase content and energetics of the resulting outflow to the cooling and no-cooling cases in Section 4.3.2. In Section 4.3.3, we investigate the effect of increasing ('Large clumps':  $\lambda_{\text{max}}=330 \text{ pc}$ ) and decreasing ('Small clumps':  $\lambda_{\text{max}}=40 \text{ pc}$ ) the average sizes of the initial clouds (see Figure 4.1). In Section 4.3.3 we investigate the sensitivity of our results to varying the initial parameters, in particular, by reducing the height of the disc ('Thin disc'), decreasing the initial mean clump density ('Low density') and reducing the initial AGN wind speed ('Slow wind'). In Section 4.4.3, we investigate a range of AGN luminosities ( $L_{\text{bol}}=10^{43-47} \text{ erg s}^{-1}$ ) to ascertain if we predict scaling relations between the luminosity, and the outflow mass or kinetic energy coupling efficiency. Finally, we also simulate our fiducial setup at lower ( $M_{\text{target}}=1000 M_{\odot}$ ) and higher ( $M_{\text{target}}=10 M_{\odot}$ ) resolutions to investigate the convergence of our simulations, which we show in Appendix 4.6.

## 4.2.6 Calculating Outflow Properties

A key aspect of this work is to evaluate the mass  $M_{\text{OF}}$ , outflow rate  $\dot{M}_{\text{OF}}$ , momentum flux  $\dot{p}$  and kinetic luminosity  $\dot{E}_k$  of the outflow, which are commonly used in observational studies to characterise AGN-driven outflows (e.g., [Cicone et al., 2014](#); [Fiore et al., 2017](#); [González-Alfonso et al., 2017](#); [Bischetti et al., 2019](#); [Musiimenta et al., 2023](#)). There are a diversity of approaches used in observational studies to derive these values (see reviews in [Harrison et al., 2018](#); [Veilleux et al., 2020](#); [Harrison & Ramos Almeida, 2024](#)), which we will discuss in Section 4.4.1. As a baseline method, we take a time-averaged approach to calculate these global quantities from our simulations. Additionally, calculating the mass contained in the outflow is challenging observationally and most approaches have some sensitivity to the minimum velocity of gas that can be detected as part of the outflow (see discussion in Section 4.4.1). Therefore, to emulate the limitation of observations, we use a minimum radial velocity cut ( $v_{\text{min}}$ ) to mimic isolating the outflow gas from that attributed to host galaxy dynamics. Because our simulations have an idealised, initially isobaric setup with no rotation or velocity dispersion, any gas that is moving should be due to the AGN wind. In practice, there is some slight collapse of the initial clouds due to cooling, but selecting a definition of  $v_{\text{min}}=10 \text{ km s}^{-1}$  yields a contamination rate of  $< 1\%$  based on a comparison to our non-AGN simulation. This value is also of the same order as the sound speed in cold clouds ( $c_s \approx 15 \text{ km/s}$ ). Therefore we take a 'theoretical' cut of  $v_{\text{min}}=10 \text{ km s}^{-1}$  to represent the outflow directly caused by the AGN. However, in an observed galaxy, gas will also be in motion due to rotation, turbulence, star formation-driven outflows, etc. Such a low cut of  $v_{\text{min}}=10 \text{ km s}^{-1}$  for an AGN-driven outflow may often not be clearly separable from other gas motions within the host galaxy (e.g., [Marconcini et al., 2023](#); [Harrison & Ramos Almeida, 2024](#)). Therefore, we also investigate a cut of  $v_{\text{min}}=100 \text{ km s}^{-1}$  as a representative example of a more 'observational' limit. The effect of  $v_{\text{min}}$  on our results is discussed in Section 4.4.1.

To calculate the global, time-averaged outflow mass, we take a cell-by-cell sum of all the gas moving at radial velocities  $> v_{\text{min}}$ ,

$$M_{\text{OF}} = \sum_i^{\text{for } v_i > v_{\text{min}}} m_i. \quad (4.3)$$

Mass outflow rates are then calculated as time-averages by dividing by the simulation time ( $t_{\text{sim}}$ ),

$$\dot{M}_{\text{OF}} = \frac{1}{t_{\text{sim}}} \sum_i^{\text{for } v_i > v_{\text{min}}} m_i, \quad (4.4)$$



while for the momentum rates, this is

$$\dot{p}_{\text{OF}} = \frac{1}{t_{\text{sim}}} \sum_i^{\text{for } v_i > v_{\text{min}}} m_i v_i, \quad (4.5)$$

and the kinetic energy flux reads

$$\dot{E}_{\text{OF}} = \frac{1}{2t_{\text{sim}}} \sum_i^{\text{for } v_i > v_{\text{min}}} m_i v_i^2. \quad (4.6)$$

We apply these equations to the simulations presented in Section 4.2.5 to investigate the outflow properties of an AGN wind interacting with a clumpy ISM structure.

### 4.3 Results I: Multiphase Outflows from Small-Scale Winds

In this section, we present the results of our simulations of an AGN wind in an idealised, clumpy galaxy disc. We first introduce our fiducial simulation, showing the multiscale structure of the outflows generated (Section 4.3.1), before analysing the phase structure and energetics (Section 4.3.2). We then discuss how changing the parameters of our galaxy setup affects our results (Section 4.3.3).

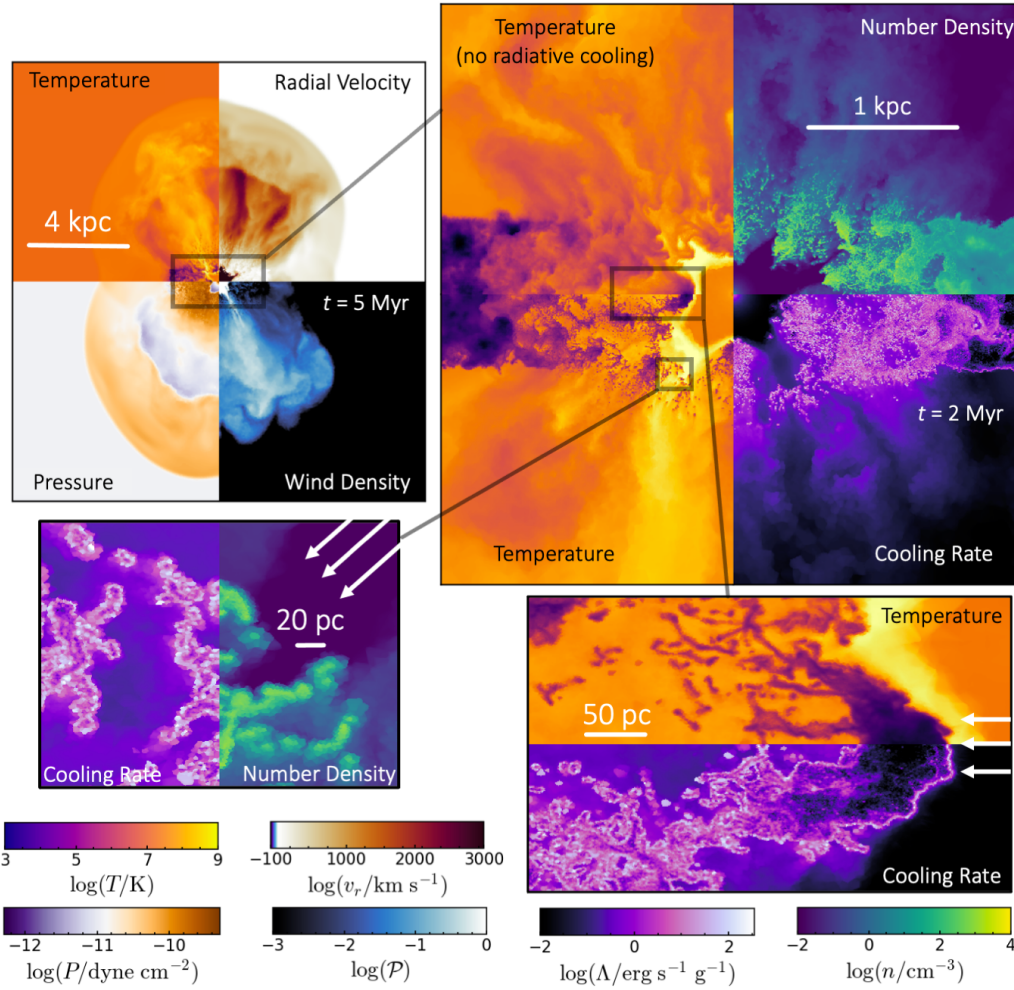
#### 4.3.1 Multiscale Outflow Structure

Figure 4.2 shows a qualitative overview of our simulations, showing the formation of a multiscale, multiphase outflow caused by the AGN wind. We use our fiducial simulation, with intermediate initial clumpiness<sup>2</sup> (corresponding to a largest clump size of  $\lambda_{\text{max}}=170$  pc) and mass resolution  $M_{\text{target}}=100 M_{\odot}$ . The AGN has a luminosity of  $L_{\text{AGN}}=10^{45}$  erg s<sup>-1</sup> and an initial wind velocity of  $v_{\text{AGN}}=10\,000$  km s<sup>-1</sup>. We show the large-scale structure of the outflow in the top left at  $t=5$  Myr, showing the maximum outflow extent reached in the simulation. The other panels show the galaxy at  $t=2$  Myr, to demonstrate the interactions between the wind and the initial cold clumps.

**Halo scale:** In the top left panel, we show an edge-on slice of the galaxy and halo, presenting (clockwise from top left quadrant) the temperature, radial velocity, wind tracer density and pressure. At this large scale, the outflow is biconical, with a clear forward shock propagating into the halo at a radius of around 6 kpc. Although the wind is injected spherically on small scales, it encounters more mass and slows down to lower velocities in the equatorial direction, leading to a bipolar outflow emerging from the top and bottom of the disc (e.g., [Costa et al., 2014a](#); [Nelson et al., 2019b](#)). The outskirts of the region populated by the wind tracer (lower-right quadrant) correspond to a contact surface separating shocked wind and shocked ambient medium fluid ([Costa et al., 2020](#)). Due to the initially inhomogeneous disc medium, we also see several high velocity ( $v_r \gtrsim 2000$  km s<sup>-1</sup>) ‘chimneys’ where the wind escapes the quickest. This can be most clearly seen in the radial velocity and wind tracer panels as narrow streams of the fastest-moving gas which create elongated structures a few kpc long, before running into the back of the shocked ambient medium at 4 kpc. The overall structure of the wind-driven outflow is in line with previous theoretical work (e.g., [King 2003](#); [Faucher-Giguère & Quataert 2012](#); [Zubovas & King 2012](#); [Costa et al. 2014b](#); [Nims et al. 2015](#); [Costa et al. 2020](#), see also our Figure 4.3) with the addition of the high-velocity chimneys caused by the initial clumpiness.


**Disc scale:** The top right panel of Figure 4.2 shows a closer view of the disc itself, showing (clockwise from top left quadrant) the temperature (in a simulation without radiative cooling) then the number density, cooling rate, and temperature, all in the

<sup>2</sup>Note: we will use the term ‘clumps’ to refer to initial overdensities in the density field, and the term ‘cloud’ or ‘cloudlet’ to refer to the resulting fragments entrained in the outflow



**Figure 4.2:** Summary figure showing the results of our fiducial simulation at a range of spatial scales, as detailed in Section 4.3.1. We show our fiducial model, with medium-sized clumps ( $\lambda_{\max}=170$  pc), an AGN luminosity of  $L_{\text{bol}}=10^{45}$  erg s $^{-1}$  and a wind injection velocity of  $v_{\text{AGN}}=10,000$  km s $^{-1}$ . The top left panel shows the large-scale biconical outflow propagating into the halo at  $t=5$  Myr and the top right panel shows the central region of the disc at  $t=2$  Myr, showing the AGN wind structure and the effect of neglecting radiative cooling (top left quadrant). The bottom two panels show the small scale fragmentation and entrainment of cold gas at  $t=2$  Myr at different spatial scales. Scale bars are shown for reference in each panel, and the white arrows demonstrate the direction of the AGN wind in the bottom two panels. All plots share the same colour scales, with corresponding values shown by the colourbars at the bottom.

fiducial simulation including cooling. In the centre, we can see the freely-expanding wind, and the forward and reverse shock fronts. We also see clear bow shocks forming around the initial clumps and free streaming of the wind in low-density regions, leading to the aforementioned ‘chimneys’. In the left-hand half of the panel, we compare the effect of including and excluding radiative cooling in the simulations. The non-cooling simulation (top left quadrant) shows a more continuous billowing of gas as it is blown away by the wind. The outflowing gas is hotter than in the cooling case, resulting in very little cold ( $T < 10^{4.5}$  K) outflowing gas. However, in the cooling case (bottom left quadrant in the top right panel), we see that the clumps cool, collapse and fragment into much smaller cloudlets than the initial clump size. These clouds become entrained in the hotter, faster wind, creating an outflow of relatively cold ( $T < 10^{4.5}$  K) gas contained in many small, high density cloudlets. We note that the exact size of these fragments is likely resolution-dependent; we explore the numerical convergence of our simulations in Appendix 4.6.



**Figure 4.3:** Our simulation with an initially smooth disc. Time increases clockwise from the top-left. The AGN wind sweeps up material forming a thin shell of outflowing material. Once this reaches a critical density, it cools suddenly via post-shock cooling, resulting in a cold ( $T < 10^{4.5}$  K) outflow at  $t \geq 0.8$  Myr. The resulting morphology and energetics are starkly different from a wind launched in an inhomogeneous medium.

**Cloud scale:** The bottom two panels of Figure 4.2 show the cloudlet features in more detail. The bottom right panel shows the temperature (top half) and cooling rate per unit mass (bottom half) of one high-density clump. The AGN wind approaches from the right (white arrows), creating a strong bow shock around the clump and creating a tail of cold fragments and filaments (see also [Cooper et al., 2008](#)). The mixing layer between the cold and the hot gas cools rapidly. This cooling ‘skin’ around the clouds very likely leads to cold cloud formation in the outflow ([Gronke & Oh, 2018](#); [Schneider et al., 2018](#); [Gronke & Oh, 2020a](#); [Fielding et al., 2020](#)), although a full exploration of this is left to future work. The bottom left panel of Figure 4.2 shows some of the entrained cloudlets, with the cooling rate on the left and the number density on the right. These cloudlets are small ( $d_{\text{cloud}} \approx 10 - 20$  pc), dense ( $n \approx 10^3 \text{ cm}^{-3}$ ; generally more dense than the initial clumps), and rapidly cooling, especially at their surface. The cooling time for these clumps is on the order of a few Myr, which is similar to their cloud crushing time. The fact that these cloudlets can survive in the wind for  $t \approx 5$  Myr adds weight to the argument that radiative cooling can sustain cold cloud growth even under strong ablation driven by ram pressure from the hot phase (see also [Gronke & Oh 2018](#)). This is in contrast to the no-cooling case where we find no cold gas in the outflow (see Section 4.3.2).

### Comparison to a smooth disc

Figure 4.3 shows our simulations with an initially smooth disc for comparison with the clumpy setup. Many previous analytic and numerical studies consider a spherically-symmetric, homogeneous setup to study the propagation of AGN winds (e.g. [King, 2003](#);

Faucher-Giguère & Quataert, 2012; Nims et al., 2015; Richings & Faucher-Giguère, 2018a). Before disc break-out, our setup mirrors this regime. The temperature of the outflow is shown, with the time increasing in 0.1 Myr increments clockwise from the top left. The wind sweeps up the gas in the disc into a thin shell of shocked ambient gas which, at  $t=0.7$  Myr, is travelling at a velocity of  $v_r \approx 400 \text{ km s}^{-1}$  with a temperature of  $T \approx 10^6 \text{ K}$  (top-left quadrant). By  $t=0.9$  Myr however, we see the emergence of a thin shell of cold ( $T \approx 10^4 \text{ K}$ ) gas, at a radius of  $R_{\text{cool}} \approx 500 \text{ pc}$  (bottom-right). This is in accordance with the expected analytic result for an ambient density of  $n_0 \approx 5 \text{ cm}^{-3}$  (see Figure 1 in Costa et al. 2020). The high density of this shell results in a shorter cooling time, creating a multiphase outflow where the hot and cold components are expanding co-spatially with the same velocity. The sudden emergence of this cold shell introduces two outflow regimes: before post-shock cooling ( $t \lesssim 0.8 \text{ Myr}$ ) and after post-shock cooling ( $t \gtrsim 0.8 \text{ Myr}$ ). In the following analysis, we will often evaluate our results at  $t=0.5 \text{ Myr}$  and  $t \geq 1 \text{ Myr}$  to represent both regimes. The outflow starts breaking out of the disc along the polar directions at  $t \approx 0.5 \text{ Myr}$  which allows some hot gas to vent out at higher velocities and turbulently mix with the cold outflow (seen at the bottom of Figure 4.3). However, the overall outflow mass is still dominated by the spherical, expanding shell (see Section 4.3.2). This shell-like outflow is starkly different to the one seen in Figure 4.2 where the cold gas is contained in small clouds entrained in a hot wind.

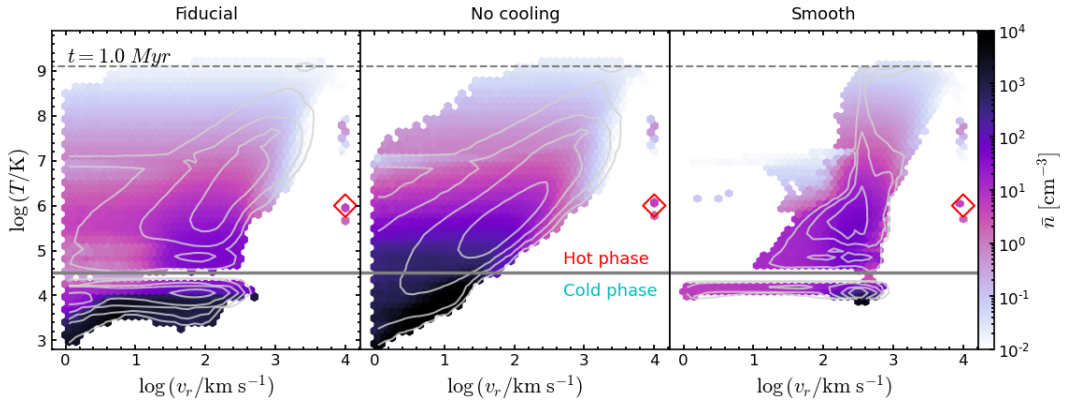
In this section, we have shown that a small-scale AGN wind model can produce kpc-scale, multiphase outflows. However, the morphology of this outflow depends strongly on the initial conditions: a smooth disc results in a thin shell of cold gas, formed by post-shock cooling, outflowing at a similar velocity as the shocked ambient medium. An initially clumpy setup results in high-velocity chimneys of hot gas, as the wind punches through regions of low density. The addition of cooling further results in the initial clumps fragmenting into small (10 – 20 pc) cloudlets which are then entrained in the hot outflow and can survive, or grow, on Myr timescales due to efficient cooling at their surface. These two setups, and their different cold gas formation channels, result in very different outflow energetics, which we quantify in Section 4.3.2 and discuss the implication of for observational estimates of outflow properties in Section 4.4.1.

### 4.3.2 Multiphase Gas Energetics

In this section, we analyse the multiphase structure of the outflows generated by our model (Section 4.3.2) before investigating how the outflow momentum and energy is distributed among the different gas phases (Section 4.3.2).

#### Phase diagrams

Figure 4.4 shows the temperature of the gas as a function of radial velocity at  $t=1 \text{ Myr}$ . The pixel colour shows the mean number density,  $\bar{n} [\text{cm}^{-3}]$ , within each pixel, and the logarithmic contours show the total mass of the outflow. We show these phase diagrams for three simulations (see Table 4.1): ‘Medium clumps’ (the fiducial model; left panel), ‘No cooling’ (everything the same as fiducial, but without cooling; middle panel), and ‘Smooth’ (homogeneous medium, with cooling; right panel). In our fiducial case (left panel) we see a clear two-phase outflow, with a cold, dense component at  $T \lesssim 10^{4.5} \text{ K}$  and a faster moving hot component centred on  $T \approx 10^{5.5-7.5} \text{ K}$ . This clear division in the phase structure motivates our later cut between the cold and hot outflow components at  $T=10^{4.5} \text{ K}$  (horizontal grey line). Looking first at the hot component, we can see the injected wind at  $v_{\text{AGN}}=10^4 \text{ km s}^{-1}$  and  $T=10^6 \text{ K}$  is quickly shock-heated to high temperatures ( $T \approx 10^9 \text{ K}$ ), which is in accordance with the expected post-shock temperature of the injected wind (see Equation 17 in Costa et al. 2020, shown here as a dashed line).




**Figure 4.4:** Phase diagrams showing the gas temperature against the radial velocity. Bins are coloured by the mean number density in each bin, with the contours showing lines of constant mass, increasing logarithmically. The red diamond marks the initial wind injection, which quickly shocks to a post-shock temperature of  $T=10^{9.1}$  K, matching the analytic expectation (dashed line; Costa et al. 2020). The left panel shows the fiducial simulation, with medium clumps and radiative cooling; the middle panel shows the same initial clumps, but without cooling; and the right panel shows a smooth disc, with cooling. The horizontal grey line shows the cut between our ‘hot’ and ‘cold’ outflow definitions at  $T=10^{4.5}$  K. We can see that cooling is essential to forming a significant cold outflow, and that the velocity structure of both outflow phases is strongly dependent on whether the initial medium was clumpy or smooth.

The bulk of the hot component (highest mass contour) is moving at a radial velocity of  $v_r \approx 500 - 1000$  km s<sup>-1</sup> and has a low density ( $\bar{n} \approx 0.01 - 1$  cm<sup>-3</sup>), especially at high temperatures ( $T \gtrsim 10^7$  K). Intriguingly at  $T \approx 10^5$  K, we see a clear enhancement in density at a velocity of  $v_r \approx 100$  km s<sup>-1</sup>, to around  $\bar{n} \approx 100$  cm<sup>-3</sup>, which is much more dense than any of the hot gas in the initial setup. The cold ( $T \leq 10^{4.5}$  K) outflow moves more slowly than the hot phase, up to a radial velocity of  $v_r \approx 400$  km s<sup>-1</sup>. The number density of this component is high, up to  $\bar{n} = 10^4$  cm<sup>-3</sup> which is similar to colder gas phases traced in observations (discussed further in Section 4.4.1). The cold gas has a wide range of outflow velocities and is moving slower than the hot phase, as it is the densest gas, which is difficult for the wind to accelerate.

The middle panel of Figure 4.4 shows the phase structure for the same initial conditions as the fiducial simulation, but without cooling. It can be seen in this case that the AGN outflow does not contain a significant cold outflow phase – there is no cold gas moving at any velocities higher than  $v_r \gtrsim 40$  km s<sup>-1</sup>. This highlights that radiative cooling of the cold outflow is essential for its survival, as the AGN wind cannot simply push out the existing cold clumps without destroying them. This matches literature results which struggle to generate a cold outflow that can survive on Myr timescales without efficient radiative cooling (e.g., Klein et al., 1994; Costa et al., 2015).

The third panel of Figure 4.4 shows the phase structure for a smooth disc. As mentioned in Section 4.3.1, the formation and evolution of the multiphase outflow differs markedly from the clumpy case. The time shown ( $t=1$  Myr) is after the post-shock cooling regime, so we have both a cold ( $T < 10^{4.5}$  K) and hot phase. There is a clear velocity peak in both phases at  $v_r \approx 400$  km s<sup>-1</sup>. This differs from the clumpy case where both phases have a broader range of velocities, and the hot phase moves significantly faster than the entrained cold clouds. In the smooth case, we do see some hot gas moving at speeds  $v_r \approx 1000$  km s<sup>-1</sup>, but this is due to the small amount of gas leaking out of the edge of the disc (bottom of Figure 4.3) and the fast small-scale wind in the centre. Nevertheless, the bulk of the outflowing mass remains moving at the characteristic velocity of  $v_r \approx 400$  km s<sup>-1</sup>. The peak density of the outflow is in the cold shell, and is similar to the peak density in the clumpy case of  $\bar{n} \approx 10^4$  cm<sup>-3</sup>.



**Figure 4.5:** Cumulative momentum rate (top panel) and kinetic energy coupling efficiency (bottom panel) as a function of temperature. The dotted lines show the simulation at  $t=0.5$  Myr and the solid lines at  $t=1$  Myr; the colours represent initial conditions of medium clumps with radiative cooling (pink) and without cooling (green), and a smooth disc (blue). The vertical grey line shows our hot/cold phase split at  $T=10^{4.5}$  K and the horizontal dashed line shows  $\dot{p}/(L/c) = 1$  which is commonly used to differentiate the energy- and momentum-driven regimes. In the fiducial case, the hot outflow dominates the energy budget due to its higher velocity, whereas in a smooth disc, the cold outflow contains a significant fraction of the momentum and energy.

### Energetics distribution

We now investigate the proportion of the radial (scalar) momentum and kinetic energy of the outflow contained in the different phases for the same simulations as in the previous section. Figure 4.5 shows the cumulative momentum flux ( $\dot{p}/(L/c)$ ; top panel) and kinetic coupling efficiency ( $\dot{E}_k/L$ ; bottom panel) as a function of temperature. We show the global outflow quantities, calculated based on a velocity cut of  $v_{\min}=10 \text{ km s}^{-1}$  (Section 4.2.6). The dotted lines show a time of  $t=0.5$  Myr and the solid lines show  $t=1$  Myr (shown to represent before and after the post-shock cooling in the smooth disc case). The vertical line shows our cold/hot phase split of  $T=10^{4.5}$  K. There is negligible difference in the momentum or kinetic energy fluxes carried in the range  $T=10^{4-5}$  K, demonstrating that our results are insensitive to the exact choice of temperature cut.

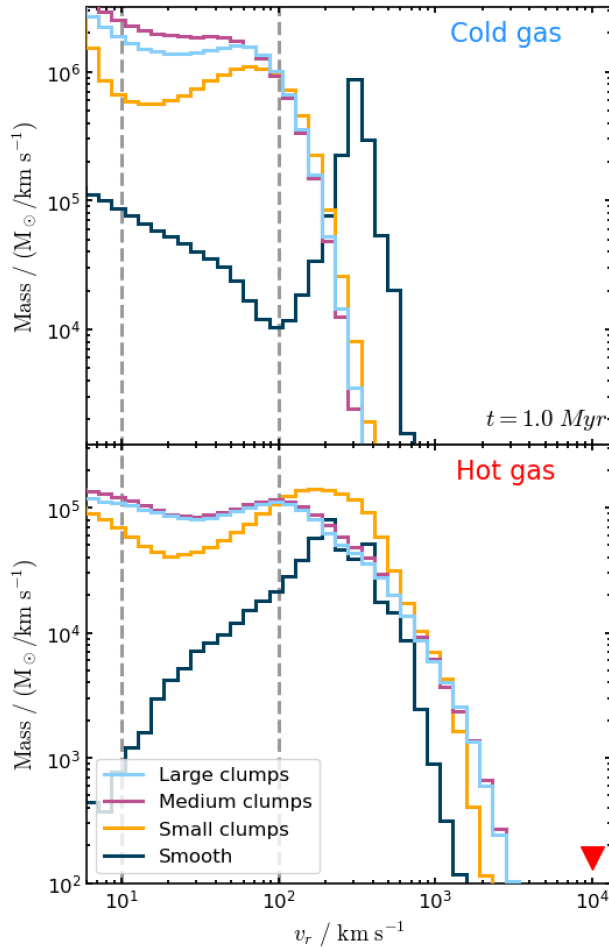
Results for our fiducial simulation (medium clumps and cooling) are shown in pink in Figure 4.5. The time evolution of this system is slight, with an increase in the momentum flux of  $\approx 0.5$  in both the cold phase and hottest ( $T \gtrsim 10^7$  K) gas, and an increase in the energy coupling of 0.1 per cent in the hottest gas. We can see that, at  $t=1$  Myr, the cold phase ( $T < 10^{4.5}$  K) has a momentum boost of  $\dot{p}/(L/c) \approx 2.4$ , which is in the energy-driving regime ( $\dot{p}/(L/c) > 1$ , horizontal dashed line). The total momentum boost rises to  $\dot{p}/(L/c) \approx 5$

when we include the hot gas, showing that the momentum is split roughly evenly between the two phases. However, when we look at the kinetic coupling efficiency, we can see that the hot phase is dominant: the cold phase only has a coupling of  $\dot{E}_k/L=0.04\%$ , compared to the total of  $\dot{E}_k/L=0.4\%$ , i.e. the hot gas contains 90% of the kinetic energy of the system. The reason for this difference in energy balance between the phases when compared to the momentum boost is the 5-10 times higher velocity of the hot outflow with respect to the entrained cold gas clouds.

For comparison with our fiducial simulation, the no-cooling case is shown in Figure 4.5 in green. As seen in Figure 4.4, there is no cold outflow without cooling, so the energy and momentum are all contained in gas with  $T \gtrsim 10^{5.5}$  K. Again, there is a mild time dependence between the two times shown, with the increase mostly occurring in the hottest gas. Comparing the no-cooling to the cooling simulation, we can see that the inclusion of radiative cooling reduces the momentum boost from  $\dot{p}/(L/c)=6.5$  to  $\dot{p}/(L/c)=5$ , representing a 23% loss, and the kinetic energy coupling rate from  $\dot{E}_k/L=0.5\%$  to  $\dot{E}_k/L=0.4\%$ , representing a 20% loss, at  $t=1$  Myr. This reduction in energy rate is both due to radiative losses and is also the result of the outflow having done different amounts of PdV work, due to the different structure of the outflows in the cooling/no cooling cases (see Figure 4.2). However, this reduction in momentum boost in the cooling simulation is not enough to change the outflow solution from energy- to the momentum-conserving, as  $\dot{p}/(L/c) > 1$  still holds, although more efficient cooling (i.e., via metal- or molecular-lines) could reduce these values further. The exact reduction in momentum boost and energy rate when cooling is included may also be dependent on the ISM structure, with different sized clouds affecting how much ablation and mixing can occur between the wind and the ISM. The effect of initial clump size is discussed further in Section 4.3.3.

The dark blue line in Figure 4.5 shows the resulting outflow from an initially smooth disc. As discussed in Section 4.3.2 the cold gas formation mechanism in outflows propagating in a homogeneous medium is post-shock cooling, which results in a phase transition at the cooling time (in our case, at  $t_{\text{cool}} \approx 0.8$  Myr, see Figure 4.3). Thus the time evolution of this system is much more significant than the clumpy case. At  $t=0.5$  Myr there is no cold outflow whereas at  $t=1$  Myr the cold outflow dominates both the momentum and energy coupling (containing 85% of the momentum flux and 75% of the kinetic luminosity). This is in stark contrast to the clumpy case where the kinetic energy in the cold gas is negligible (10% of the total energy). This is because most mass in both the cold and hot outflow in the smooth case is confined in a thin shell (Figure 4.3), expanding at  $v_r \approx 400$  km s<sup>-1</sup> (Figure 4.4). However, in the clumpy case, the hot gas is able to vent much faster (up to  $v_r \approx 1000$  km s<sup>-1</sup>) past the entrained cold clouds (Figure 4.2) which are generally travelling much slower (around  $v_r \approx 10 - 100$  km s<sup>-1</sup>; Figure 4.4). This results in much less efficient transfer of momentum and energy from the hot AGN wind to the cold component of the outflow. The overall momentum of the smooth case is also higher, at  $\dot{p}/(L/c)=7$ , however, the total kinetic energy is similar to the cooling clumpy case at  $\dot{E}_k/L=0.44\%$ .

To summarise, in this section we have shown that cooling creates a two-phase outflow, but the share of the momentum and energy carried by each phase strongly depends on the initial conditions of the disc. The outflows are all in the energy-conserving regime ( $\dot{p}/(L/c) > 1$ ), however, in the clumpy case there is significant mixing between wind fluid and ambient gas. The associated cooling gives rise to a cold outflowing phase, but does not cause sufficient cooling losses to drive the outflow solution into the momentum-driven regime. Furthermore, the clumpy media confines the energy-driven bubble less efficiently, resulting in lower momentum fluxes than in a smooth disc.



**Figure 4.6:** Histograms showing the outflowing mass at different radial velocities at  $t=1$  Myr. Each bin is normalised by its width to remove dependence on bin size. Gas is split by phase into cold gas (top panel) and hot gas (bottom panel). The colour of the lines represents the initial conditions of the simulation: large clumps (light blue); medium clumps (pink); small clumps (yellow) and smooth medium (dark blue). The red arrow marks the velocity of the injected AGN wind and the dashed grey vertical lines show the two velocity cuts used ( $v_{\min}$ ). The smooth case shows a clear characteristic velocity peak at  $v_r \approx 400$  km s $^{-1}$  in both the cold and hot phases, whereas the clumpy setup shows a range of outflow velocities, and a clear difference in velocities between the phases.

### 4.3.3 Parameter Variation

In Section 4.3.2, we explored how the outflow produced from interactions between a clumpy ISM and an AGN wind differs from that in a smooth disc. We will now briefly analyse the sensitivity of our results to changes in the initial conditions of the simulations to assess the robustness of our conclusions. We investigate the impact of: the choice of distribution of clump sizes in the ISM initial conditions ( $\lambda_{\max}=40, 170, 330$  pc; see Figure 4.1); a different choice for the density and thickness of the disc (and hence the total mass); and the speed of the wind at injection. Variations in AGN luminosity are explored later, in Section 4.4.3. In Section 4.3.3 we focus first on the initial ISM conditions and the impact this has on the radial velocity distribution of the outflow, and then in Section 4.3.3 we explore the change in mass outflow rate for the other simulation variations shown in Table 4.1. We also test the numerical convergence of these global properties, finding that they are well-converged at our fiducial resolution of  $M_{\text{target}}=100 M_{\odot}$  (see Appendix 4.6).



### Outflow variation due to initial clumpiness

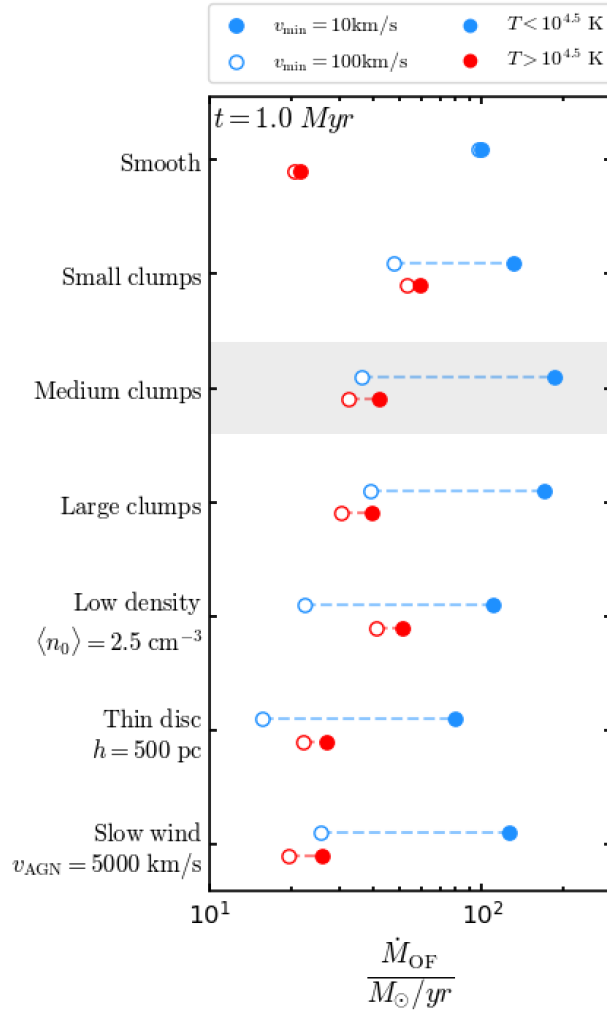
Figure 4.6 shows the distribution of the outflowing mass as a function of radial velocity at  $t = 1$  Myr for the three different initial clump sizes. The top panel shows the cold ( $T < 10^{4.5}$  K) phase and the bottom shows the hot gas. Note that the y-scale on the bottom panel is smaller, demonstrating that the hot phase carries less mass in all but the highest velocity bins ( $v_r \gtrsim 500$  km s $^{-1}$ ). The red arrow marks the injection velocity of the AGN wind (not seen in the plot due to its very low density).

We can clearly see, in both gas phases, that the differences between the clumpy cases (colourful lines) and the homogeneous medium (dark blue line) is greater than the variation between the initial clump sizes. As we are in the post-shock-cooling regime ( $t_{\text{cool}} \geq 0.8$  Myr; see Section 4.3.1), the smooth case has produced a cold outflow which shows a narrow peak around  $v_r = 400$  km s $^{-1}$ . The hot gas also peaks around a similar characteristic velocity, as it traces the shocked ambient gas shell. The shocked wind at small scales is also hot, but has very little mass. In contrast, the clumpy simulations show no such characteristic velocity, instead showing a range of speeds that gradually drop off after  $v_r \gtrsim 100$  km s $^{-1}$ , for the cold phase, and  $v_r \gtrsim 500$  km s $^{-1}$ , for the hot phase. This further demonstrates our earlier finding (Section 4.3.2) that there is a significant velocity differential between the phases when the initial medium is clumped, which is not seen in the homogeneous case.

We also find variations between the different initial clump size distributions (as also seen in [Wagner et al. 2012](#); [Bieri et al. 2017](#)). Our medium ( $\lambda_{\text{max}} = 170$  pc; purple line) and large ( $\lambda_{\text{max}} = 330$  pc; blue) clumps show only small differences, despite the size of the average largest clumps being twice that of those in the medium case. The small clumps ( $\lambda_{\text{max}} = 40$  pc) show the greatest difference: there is a weak peak in velocity at  $v_r \approx 80$  km s $^{-1}$  in the cold phase and  $v_r \approx 200$  km s $^{-1}$  in the hot phase. There is also the lowest mass of slow-moving ( $v_r < 100$  km s $^{-1}$ ) cold gas compared to the other two clumpy simulations, with  $\sim 0.6$  dex less outflowing mass above  $v_r = 10$  km s $^{-1}$ . However, at radial velocities  $v_r > 100$  km s $^{-1}$ , the small clumps case has the highest outflowing mass in all velocity bins. A similar picture is seen in the hot phase, with the small clumps having the lowest outflowing mass at  $v_r < 100$  km s $^{-1}$ , but the most in the higher velocity gas ( $v_r = 100 - 1000$  km s $^{-1}$ ). However, in the fastest moving gas ( $v_r > 1000$  km s $^{-1}$ ), the small clumps case again drops to the lowest. An interesting point to note is that the lines for the small clumps (yellow) start trending towards the smooth case (dark blue) compared to the other two clump sizes, such as showing a more peaked distribution. This could suggest that many small clumps ( $\lambda_{\text{max}} = 40$  pc) start acting like a smooth medium in trapping the outflow and could contain a mix of post-shock-cooling shell and mixing layer cooling mechanisms. This restricts the fastest-moving hot gas, but is more effective at accelerating the cold cloudlets to velocities  $v_r > 100$  km s $^{-1}$ . Furthermore, the resulting morphology of the cold clouds produced may still be dependent on the initial clumps from which they form; i.e. we may see smaller average cloudlets being formed from smaller initial clumps. These more detailed investigations are left to future work ([Almeida et al., prep](#)).

### Sensitivity to disc setup

Figure 4.7 shows the mass outflow rate at  $t = 1$  Myr for a range of setup variations, as an example of the broad quantitative effect on outflow properties of our setup choices (this analysis was also performed for  $\dot{E}_k/L$  with similar results). In this Section, we restrict ourselves to only analysing simulations with  $L_{\text{bol}} = 10^{45}$  erg s $^{-1}$  and examine differences due to AGN luminosity to Section 4.4.3. We use our two-phase cut with blue and red representing the cold and hot gas respectively. As discussed in the previous section, when the initial medium is clumpy, the outflow does not have a characteristic velocity, showing a range of velocities (Figure 4.6). For this reason, we show two outflow definitions using



**Figure 4.7:** The mass outflow rates at  $t=1$  Myr for the range of simulation parameters explored, with our fiducial simulation shaded for reference. The cold phase ( $T < 10^{4.5}$  K) is shown in blue and the hot phase ( $T > 10^{4.5}$  K) in red. Two different outflow definitions are shown based on a radial velocity cut:  $v_{\min}=10$  km  $s^{-1}$  is shown in solid points and  $v_{\min}=100$  km  $s^{-1}$  are shown as empty circles. Despite the wide range in initial conditions shown, the outflow masses generally vary by less than a factor of two. However, the choice of  $v_{\min}$  has a much greater effect, with the cold phase outflow rates varying by around a factor of eight for all setups except the smooth disc.

differing radial velocity cuts:  $v_{\min}=10$  km  $s^{-1}$  as solid points and  $v_{\min}=100$  km  $s^{-1}$  as empty circles. The observational motivation and implications of these two cuts is discussed further in Section 4.4.1. The following discussion will mostly focus on the  $v_{\min}=10$  km  $s^{-1}$  definition (solid points).

Our fiducial simulations (medium clumps) has a global mass outflow rate (following Section 4.2.6) in the cold gas of  $185 M_{\odot}/\text{yr}$  and a hot mass outflow rate of  $\approx 40 M_{\odot}/\text{yr}$  (both at a velocity cut of  $v_{\min}=10$  km  $s^{-1}$ ). This represents a mass outflow rate of almost a factor of two higher than the respective phases in the smooth case. This reiterates the point that, although the smooth case has greater momentum and energy fluxes (Figure 4.5), the clumpy setup creates a greater mass in the outflow, albeit moving at a slower velocity (Figure 4.6). The large clump setup ( $\lambda_{\max}=330$  pc) has similar values to the fiducial case, with outflow rates lower by  $\lesssim 10 M_{\odot}/\text{yr}$ . The small clumps case ( $\lambda_{\max}=40$  pc) has slightly elevated hot gas ( $60 M_{\odot}/\text{yr}$ ) but reduced cold gas ( $130 M_{\odot}/\text{yr}$ ) in its outflow, suggesting

that smaller initial clouds lead to a higher trapping efficiency and/or an increase in mixing and cooling between the phases. Overall, varying the initial clumpiness of the disc changes the outflow rate by up to  $\lesssim 30\%$ .

We also varied the initial mass of the galaxy, by decreasing both the initial mean density and the height of the disc by a factor of two compared to the fiducial case. The low density setup ( $\langle n_0 \rangle = 2.5 \text{ cm}^{-3}$ ) has a slightly reduced mass outflow rate compared to the fiducial case ( $110 \text{ M}_\odot/\text{yr}$ ) and shows a marginal ( $10 \text{ M}_\odot/\text{yr}$ ) increase in the hot phase. However, lowering the initial density also increases the porosity of the disc due to more regions falling above the  $T_{\text{crit}}$  cut (see Section 4.2.3), which slightly complicates the comparison here (see [Wagner et al. 2012](#)). The thinner disc has just under half the outflow rate in the cold phase compared to the fiducial case, and the hot phase is reduced by  $\approx 10 \text{ M}_\odot/\text{yr}$ . This reduction is mostly in the polar direction as there is less gas for the wind to interact with before it breaks through the top and bottom of the disc. These results for the varying the disc height are consistent with the work of [Wagner et al. \(2013\)](#) when they moved from a (thin) disc to a (thick) bulge distribution for the gas clumps. Finally, a slower AGN wind ( $v_{\text{AGN}}=5000 \text{ km s}^{-1}$ ) results in a mass outflow rate which is lower by  $\approx 30\%$ , due to the injection energy also being reduced by half (see Equation 4.2). Such difference in galaxy disc structure and ISM conditions would add some modest scatter to observationally-derived scaling relations between the AGN power and outflow properties.

In conclusion, we have shown that changes in the ISM structure and galaxy setup change the mass outflow rate by up to a factor of two. We performed the same analysis on the momentum and energy rates and found similar trends. In Section 4.4.3, we place these differences in the context of global scaling relations to further investigate how great an impact they have on our final results. Furthermore, the phase and multiscale structure shown in Figures 4.2 & 4.4 are qualitatively consistent across all the clumpy simulations, and in contrast with the smooth disc setup (Figure 4.3). However, a large factor in determining the quantitative outflow rates is the choice of  $v_{\text{min}}$  cut due to the relative distribution between phases. An increase in this value from  $v_{\text{min}}=10 \text{ km s}^{-1}$  to  $v_{\text{min}}=100 \text{ km s}^{-1}$  can lead to a reduction of the measured cold mass outflow rate by a factor of eight, as seen by the empty circles, which we will explore in more detail in Section 4.4.1.

## 4.4 Results II: Implications for Observed Outflow Properties

In the previous section, we explored how the interaction between an AGN wind and an initially clumpy medium creates a multiphase, multiscale outflow, which is qualitatively distinct from the spherical shell-like outflow generated in the homogeneous disc. In this section, we discuss the implications of our results for observational studies of AGN outflows. We review how the nature of these clumpy outflows poses challenges for measuring outflow properties such as outflow rates (Section 4.4.1) and also explore how this affects whether the outflow is inferred to be energy- or momentum-driven (Section 4.4.2). Finally, in Section 4.4.3, we investigate scaling relations between outflow properties and AGN luminosity and compare these to recent observational attempts.

### 4.4.1 Challenges for Measuring Outflow Properties

When assessing the potential impact of AGN on the ISM and host galaxy evolution, observers take multiple approaches (see review in [Harrison & Ramos Almeida, 2024](#)). One common approach is to calculate the time-averaged mass outflow rate, which, in its

simplest form, is given by

$$\dot{M}_{\text{OF}} = B \cdot \frac{M_{\text{OF}} \cdot v_{\text{OF}}}{R_{\text{OF}}}, \quad (4.7)$$

where, usually,  $B = 1$  or  $B = 3$ , depending on the assumed geometry of the outflow (see e.g., [González-Alfonso et al., 2017](#)). The momentum flux and kinetic luminosity<sup>3</sup> of the outflows are then calculated with increasing powers of velocity, respectively as

$$\dot{p} = \dot{M}_{\text{OF}} \cdot v_{\text{OF}}, \quad (4.8)$$

and

$$\dot{E}_k = \frac{1}{2} \cdot \dot{M}_{\text{OF}} \cdot v_{\text{OF}}^2. \quad (4.9)$$

These equations are then used to assess the potential impact of outflows, and underlying driving mechanisms, by normalising by the AGN luminosity to create the dimensionless quantities of momentum flux,  $\dot{p}/(L/c)$ , and outflow kinetic coupling efficiency,  $\dot{E}_k/L$ . These equations are commonly used to explore multiphase outflow properties across different observed galaxies populations, with the benefit that the energetics of the outflow can be measured using just the three quantities of  $M_{\text{OF}}$ ,  $v_{\text{OF}}$  and  $R_{\text{OF}}$  (as well as  $L_{\text{bol}}$  for normalisation). However, it is important to assess their validity (see derivations and discussion of challenges in e.g., [Rupke et al., 2005](#); [González-Alfonso et al., 2017](#); [Harrison et al., 2018](#); [Lutz et al., 2020](#); [Veilleux et al., 2020](#); [Kakkad et al., 2020](#)), and detailed observations and modelling are challenging some of their assumptions (e.g., [Crenshaw & Kraemer, 2000](#); [Meena et al., 2023](#)). In this Section, we therefore explore using Equations 4.7, 4.8, & 4.9 to compute outflow properties in line with these observational approaches (see Section 4.4.2).

### Outflow velocity & mass

The bulk outflow velocity ( $v_{\text{OF}}$ ) is a critical measurement for inferring outflow rates (Equation 4.7). Observationally, outflow velocities are difficult to constrain (see discussions in e.g., [Veilleux et al., 2020](#); [Harrison et al., 2018](#); [Harrison & Ramos Almeida, 2024](#)) with issues including spectral resolution, projection effects and beam smearing (e.g., [Husemann et al., 2016](#); [Wang et al., 2024a](#)). One of the biggest challenges can be separating outflowing gas from non-outflowing gas. This is particularly problematic in spatially-unresolved observations, where a single emission line or absorption line profile is relied upon to determine an outflow velocity. There is a diversity of approaches in the literature for defining an outflow velocity from observations, and for deciding which gas is considered to be outflowing (e.g., [Carniani et al., 2015](#); [Cresci et al., 2015b](#); [Harrison et al., 2018](#); [Kang & Woo, 2018](#); [Veilleux et al., 2020](#)). These include taking a minimum velocity threshold of the emission/absorption line wings to separate outflowing and non-outflowing material (e.g., [Kakkad et al., 2020](#); [Lamperti et al., 2022](#)); taking the maximum velocity to represent the whole outflow to compensate for potential underestimation due to projection effects (e.g., [Fiore et al., 2017](#)); performing a multi-component fit, with the broad component tracing the outflowing gas (see [Lutz et al., 2020](#); [Hervella Seoane et al., 2023](#); [Gatto et al., 2024](#), for example explorations of different methods); and constructing a dynamical model of the galaxy and assuming that the residuals are formed by outflowing gas (e.g., [Rupke et al., 2017](#); [Ramos Almeida et al., 2022](#); [Girdhar et al., 2022](#)). However, all these methods, to greater or lesser extents, are still subject to systematic uncertainties and limitations on the minimum outflow velocity that can be determined above the systemic movement of

<sup>3</sup>We note that, sometimes, an additional term is added to the kinetic luminosity, to account for turbulent gas motions.

the galaxy. This consequently has implications for estimating a total mass of outflowing material as masses are typically calculated by integrating the flux from a given line over a velocity range, and then converting this to a mass measurement.

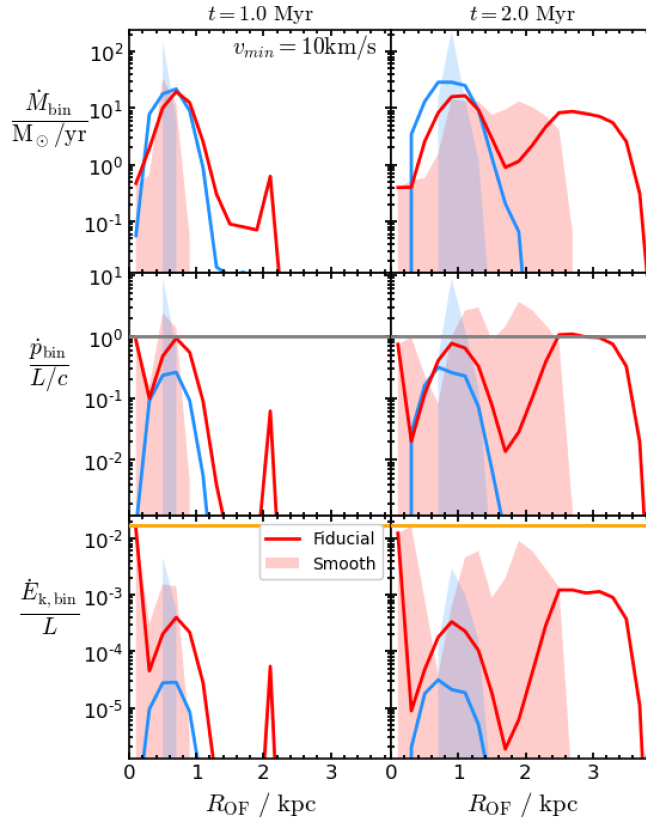
In our simulations we thus explore using two radial velocity cuts to define which gas is outflowing: a ‘theoretical’ value of  $v_{\min}=10 \text{ km s}^{-1}$  and a representative ‘observational limit’ value of  $v_{\min}=100 \text{ km s}^{-1}$  (see Section 4.2.6). In Figure 4.7, we show the mass outflow rate calculated both for  $v_{\min}=10 \text{ km s}^{-1}$  (solid circles) and  $v_{\min}=100 \text{ km s}^{-1}$  (hollow circles). We showed that if the outflow has a spherical shell-like morphology (smooth ISM case), the choice of  $v_{\min}$  has little impact on the resulting mass outflow rate. *However, if the initial medium is clumpy, there is no single characteristic outflow velocity, with both phases showing a broad range in radial velocity distribution* (see Figure 4.6) and thus no straightforward way to isolate AGN-driven outflowing motion. This results in the outflow mass being highly sensitive to  $v_{\min}$ , especially for the cold phase, as the bulk of this gas is travelling at lower velocities (see also [Costa et al., 2018b](#)). Increasing  $v_{\min}$  from  $10 \text{ km s}^{-1}$  to  $100 \text{ km s}^{-1}$  results in around a factor of eight lower outflow rate for the cold gas, and around 10% lower in the hot phase. This difference is much more significant than the variance in the measured outflow rate due to the initial condition parameters such as the clumpiness, disc height or AGN wind velocity (Figure 4.7). For this reason, throughout this Section, we will consider the impact of using two different velocity cuts on outflow properties.

### Outflow phases

Another factor in measuring the mass of an outflow is which gas phase is detected. Although outflows have been detected in multiple phases (e.g., [Liu et al., 2013](#); [Rupke & Veilleux, 2013](#); [Carniani et al., 2015](#); [González-Alfonso et al., 2017](#); [Girdhar et al., 2022, 2024](#)), observations are usually limited to tracing one phase, or only a few phases in a limited temperature range ([King & Pounds, 2015](#); [Veilleux et al., 2020](#); [Harrison & Ramos Almeida, 2024](#)). As shown in Section 4.3, when the initial medium is clumpy, the cold phase ( $T < 10^{4.5} \text{ K}$ ) carries the bulk of the mass (mass outflow rates a factor of 5 higher; Figure 4.7), but the hot phase is much more energetic, with an order of magnitude higher kinetic energy coupling efficiency (Figure 4.5). These results are in agreement with observational studies, which typically find the colder outflow phases are lower velocity, less spatially extended, and have higher mass outflow rates ([Vayner et al., 2021](#); [Girdhar et al., 2022](#); [Speranza et al., 2024](#)). However, studying the hottest outflow phases (i.e.,  $T \gtrsim 10^6 \text{ K}$ ) is challenging due to the low density of the X-ray emitting gas, but there has been some success (e.g., [Greene et al., 2014](#); [Veilleux et al., 2014](#); [Lansbury et al., 2018](#)). *The results presented in Section 4.3 clearly show that a multi-wavelength approach is crucial as otherwise a significant amount of mass/energy of the outflow will not be observed if only hot/cold phases are studied.* For the remainder of this section, we thus investigate the effect of separating our outflow into distinct phases through a simple temperature cut at  $T=10^{4.5} \text{ K}$  (see Figure 4.4).

### Outflow radius and resolved outflow rates

The radial extent of detected AGN outflows varies from 10s pc - 10s kpc across multiple gas phases (e.g., see [Veilleux et al., 2020](#); [Harrison & Ramos Almeida, 2024](#)). Measuring the exact value of  $R_{\text{OF}}$  can be challenging, especially in seeing-limited conditions, with a resolution of 1" corresponding to 2 kpc at  $z = 0.1$  or 9 kpc at  $z = 1$  (e.g., [Husemann et al., 2016](#); [Tadhunter et al., 2018](#)), although spectroastrometric techniques can help improve the effective resolution ([Carniani et al., 2015](#); [Lamperti et al., 2022](#)). If the outflow is unresolved, the beam size or fibre width is often taken as an upper limit on the radius. Some



**Figure 4.8:** Radial evolution of the outflow, showing the localised mass outflow rate (top), momentum boosting (middle) and energy coupling efficiency (bottom) for the fiducial simulation (lines) and the smooth case (shaded). On the left is shown  $t=1$  Myr and on the right  $t=2$  Myr. The grey horizontal line in the middle panels shows the momentum-conserving value of  $\dot{p}/(L/c) = 1$  and the orange horizontal line in the bottom panels shows our injected energy rate ( $\dot{E}_k/L = 1.7\%$ ). We can see that the cold phase (blue) dominates the mass at low radii, but the hot phase (red) dominates the energy, especially at large spatial scales.

studies with good spatial resolution (order of 100s pc) utilise long-slit spectroscopic or integral field unit observations to analyse the radial evolution of the outflow (e.g., [Crenshaw & Kraemer, 2000](#); [Revalski et al., 2018, 2021](#); [Meena et al., 2023](#); [Riffel et al., 2023](#)). In this Section, we emulate this approach by taking spherical bins of width  $\Delta r = 200$  pc and calculating the total mass outflow, momentum and energy rate within each spatial bin (following [González-Alfonso et al., 2017](#)) as

$$\dot{M}_{\text{bin}} = \frac{M_{\text{bin}} \cdot v_{\text{med}}}{\Delta r}, \quad (4.10)$$

where  $M_{\text{bin}}$  is the outflowing mass within each bin and  $v_{\text{med}}$  is the median velocity within the bin. Momentum and energy are calculated with further powers of  $v_{\text{med}}$ . We take an outflow velocity cut of  $v_{\text{min}}=10 \text{ km s}^{-1}$ , and note that because the calculation involves taking a median velocity, the results do not differ significantly from the higher  $v_{\text{min}}=100 \text{ km s}^{-1}$  cut. The reduction in mass due to a higher  $v_{\text{min}}$  cut is compensated by a higher median velocity. Our results are also insensitive to the bin width chosen. We stress that these *radially resolved* values differ from the *globally averaged* values used in the rest of the paper.

The radial profile of outflow properties is shown in Figure 4.8. We show the mass outflow rate (top row), momentum flux (middle row) and kinetic energy coupling (bottom row) at outflow times of  $t=1$  Myr (left column) and  $t=2$  Myr (right column), with the colours showing the phase split as before (blue: cold, red: hot gas). The solid lines show

the results from the fiducial (clumpy ISM) simulation and the shaded areas show the results for an initially smooth medium. In the clumpy case, we can see that the cold phase dominates the mass outflow rate, but only at small radii. The hot phase is moving faster and thus dominates the mass outflow rate at halo scales. The hot phase dominates the momentum flux and kinetic luminosity at all radii as its higher velocity more than compensates for its lower mass loading. This hot gas shows a two-humped structure, similar to that seen in the ionised phase of some observed outflows (Revalski et al., 2018, 2021), which could point to two outflow structures: an equatorial outflow travelling at a slower speed and the uninhibited polar outflow. The outflow through the initially smooth disc (shaded areas; Figure 4.8) is concentrated at a single radius, especially at  $t=1$  Myr, due to its shell-like morphology (see Section 4.3.1; Figure 4.3). Interestingly, the kinetic energy coupling reached by the cold outflow in this smooth disc is two orders of magnitude higher than the maximum seen in clumpy case. This is likely due to its much more efficient confinement of the energetic hot gas (Section 4.3.1).

A spatially-resolved approach shows that the two phases have very different radial extents, with the hot outflow reaching  $r = 3.5$  kpc by  $t=2$  Myr and the cold outflow remaining within  $r < 1.5$  kpc. As demonstrated in Figure 4.2 and quantified in Figure 4.8, *when the initial medium is clumpy, there is no single characteristic outflow radius*, making it challenging to calculate single mass, momentum flux, or kinetic luminosity outflow rates, as is performed for observations where outflows are not spatially resolved.

#### 4.4.2 Challenges for Interpreting Momentum Fluxes and Kinetic Luminosities

Alongside the challenges for measuring outflow properties, there are difficulties in using these quantities to interpret the resulting momentum boost and kinetic luminosities. Observational studies (e.g., Ciccone et al., 2014; Carniani et al., 2015; Fiore et al., 2017; Musiimenta et al., 2023) often use the momentum flux to assess whether the outflow is energy- ( $\dot{p}/(L/c) > 1$ ) or momentum-driven ( $\dot{p}/(L/c) < 1$ ). This is motivated from analytical models that predict that large-scale, energy-driven winds can have momentum fluxes as high as  $\dot{p}/(L/c) \approx 20$  (King, 2003; Zubovas & King, 2012; Faucher-Giguère & Quataert, 2012). Likewise, the kinetic energy coupling efficiency ( $\dot{E}_k/L$ ) is often calculated in observations to assess how efficiently the AGN wind couples to the ISM (e.g., Fiore et al., 2017). The observed kinetic luminosity is sometimes compared to the input feedback efficiencies set in cosmological simulations, which generally are  $\approx 5 - 20\%$  (Schaye et al., 2015; Weinberger et al., 2017). Such comparisons are misleading, as these values represent the subgrid feedback efficiency assumed in the simulations – which is often chosen for numerical, not physical reasons – and cannot be straightforwardly translated into galaxy-wide outflow kinetic luminosities (see further discussion in Harrison et al. 2018). Although it may be possible for a perfectly energy-driven wind to achieve efficiencies of  $\dot{E}_k/L=5\%$  (King, 2005), in practice, the large-scale outflow may have an order of magnitude lower kinetic energy due to gravitational work and radiative cooling losses (Veilleux et al., 2017; Richings & Faucher-Giguère, 2018b; Costa et al., 2020).

In this Section, we explore the predictions of momentum boosts and kinetic coupling efficiencies from our simulations, and discuss the resulting challenges this might pose for observers in using these quantities for interpretation.

#### Radial profiles of outflow properties

In Figure 4.8, we showed the spatially-resolved radial evolution of outflow momentum and kinetic energy fluxes. We mark the commonly assumed boundary between the energy- and momentum-driven regime ( $\dot{p}/(L/c) = 1$ ) as a grey line in the second row of panels.

In the clumpy case (solid lines), the cold outflow would always be seen as momentum-conserving, despite having a global value that places it in marginally energy-driven regime ( $\dot{p}/(L/c) = 2.5$ ; Figure 4.5). The hot outflow has consistently higher values (due to its higher velocity; Figure 4.6) but would also still be interpreted as momentum-driven. In the smooth case, however, the cold phase can be seen as energy-driven at  $R_{\text{OF}} = 600$  pc at  $t=1$  Myr and  $R_{\text{OF}} = 1$  kpc at  $t=2$  Myr. This is due to all the cold gas being constrained to a thin shell travelling at  $v_r \approx 400$  km s $^{-1}$ . At  $t=1$  Myr, the hot phase is still mostly constrained to this shell so we also see an energy-driven outflow at  $R_{\text{OF}} = 600$  pc, but by  $t=2$  Myr, enough has broken out of the disc to extend out to  $R_{\text{OF}} = 2.2$  kpc.

The bottom row of Figure 4.8 shows the radial evolution of the kinetic energy coupling efficiency. The orange line in the bottom row shows the energy flux of the small-scale wind  $\dot{E}_k/L=1.7\%$ . The hot outflow completely dominates over the cold at all radii, with the cold outflow only having a peak value of  $\dot{E}_k/L \approx 0.003\%$ . The hot outflow varies significantly with radius, but has a peak value of  $\dot{E}_k/L \approx 0.1\%$  in the outer shock front ( $R_{\text{OF}} \approx 3$  kpc). The smooth case again has higher peak values, with the cold phase peaking at  $\dot{E}_k/L \approx 0.4\%$  and the hot at almost 1% at  $R_{\text{OF}} \approx 2$  kpc at  $t=2$  Myr. Furthermore, there is a large variation in radius with these quantities, for example, the energy and momentum fluxes of the hot outflow drop by almost two orders of magnitude from  $R_{\text{OF}} = 1$  kpc to  $R_{\text{OF}} = 1.8$  kpc before rising by the same amount or more again at  $R_{\text{OF}} = 2.4$  kpc, once the outflow accelerates into the halo.


In conclusion, the location at which the outflow is measured can have a large effect on how the momentum fluxes and kinetic luminosities are interpreted. When studies spatially resolve the outflow, as we have emulated in Figure 4.8, even if the global outflow is energy-driven, the wide spread in radii and velocities can result in no one radial bin implying an energy-driven solution based on measured momentum fluxes. For studies that do not spatially resolve the outflow and instead take a global approach, it is important to note that the two phases dominate at different spatial scales. In addition, a significant amount of energy and momentum may be missed when observations are unable to measure the hottest gas phases at the largest scales.

### Time evolution of global quantities

We also investigated how the time at which the outflow is measured may affect how the global, time-averaged energetics are interpreted. Figure 4.9 shows the global time evolution of the mass outflow rate (top panel), momentum boost (middle panel) and kinetic energy coupling efficiency (bottom panel), with the cold/hot phases in blue/red. The solid lines show results for outflow velocity cut of  $v_{\text{min}}=10$  km s $^{-1}$  (our ‘theoretical value’) and the dashed lines show  $v_{\text{min}}=100$  km s $^{-1}$  (representative of an ‘observational’ limit, see Section 4.4.1). We can see that, using  $v_{\text{min}}=10$  km s $^{-1}$ , the cold outflow dominates the mass outflow rate, rising rapidly to a peak of  $\dot{M}_{\text{OF}} \approx 200 M_{\odot} \text{ yr}^{-1}$  by  $t=1.5$  Myr, before flattening. The hot outflow shows a peak of  $\dot{M}_{\text{OF}} \approx 65 M_{\odot} \text{ yr}^{-1}$  at  $t=0.4$  Myr, before settling to  $\dot{M}_{\text{OF}} \approx 40 M_{\odot} \text{ yr}^{-1}$  by  $t=1.5$  Myr. However, using the higher outflow velocity cut of  $v_{\text{min}}=100$  km s $^{-1}$  drastically reduces the inferred cold gas outflow rate, which now peaks at  $\dot{M}_{\text{OF}}=40 M_{\odot} \text{ yr}^{-1}$  and shows a steeper tail-off, dropping by a factor of 2 in the 4 Myr after the peak. As previously discussed (Section 4.4.1), a higher outflow velocity cut has less impact on the hot phase, resulting in only a  $\approx 15\%$  decrease across most of the simulation time.

A similar picture emerges when considering the momentum flux (middle panel of Figure 4.9). The momentum flux in the cold outflow again rapidly rises to a peak at  $\dot{p}/(L/c) \approx 2.5$  by  $t=1$  Myr for  $v_{\text{min}}=10$  km s $^{-1}$ . However, when using the cut of  $v_{\text{min}}=100$  km s $^{-1}$ , the peak value reduces to  $\dot{p}/(L/c) \approx 1$ , before dropping further by 60% by  $t=5$  Myr. We can see that both the outflow time and the minimum velocity cut used affect whether






**Figure 4.9:** Time evolution of outflow in our fiducial simulation showing the mass outflow rate (top), momentum boosting (middle) and energy coupling efficiency (bottom). Solid lines show  $v_{\min}=10 \text{ km s}^{-1}$  and dashed lines show  $v_{\min}=100 \text{ km s}^{-1}$ . The cold phase is in blue and the hot phase is in red. The grey horizontal line (middle panel) shows the momentum-conserving value of  $\dot{p}/(L/c) = 1$  and the orange horizontal line (bottom panel) shows the injected energy rate ( $\dot{E}_k/L = 1.7\%$ ). The cold phase dominates the mass outflow rate and the hot phase dominates the momentum flux and kinetic luminosity. However, the measured properties of cold phase is particularly sensitive to the minimum radial velocity cut ( $v_{\min}$ ).

a momentum boost of  $\dot{p}/(L/c) > 1$  or  $\dot{p}/(L/c) < 1$  is measured. If only the cold phase is considered, the outflow may be incorrectly categorised as ‘momentum-driven’ despite the overall outflow solution being energy-driven.

Looking at the third row of panels, the hot outflow dominates the energy of the outflow with a steady value of  $\dot{E}_k/L \approx 0.5\%$ . However, the kinetic luminosity in the cold phase is much lower, peaking at  $\dot{E}_k/L = 0.04\%$  and showing a strong time evolution, especially with the higher minimum velocity limit, dropping to  $\dot{E}_k/L = 0.01\%$  at  $t=5 \text{ Myr}$ . Compared to the kinetic luminosity of the small-scale wind ( $\dot{E}_k/L = 1.7\%$ ; Equation 4.2), we can see that the total kinetic luminosity of the large-scale outflow is lower by a factor of three. If only the cold phase were observed, this difference amounts to two orders of magnitude.

In this section we have shown that *where, when and in which phase* the outflow is being measured can have a large impact on the inferred mass outflow rates, momentum boosts and kinetic coupling efficiencies. These challenges are exacerbated in the case of an inhomogeneous medium where there is a wider range of outflow velocities than in the smooth case. If only a single phase is measured, this can lead to the outflow being mischaracterised as momentum-driven despite the overall energy-driven nature of the outflow.



**Figure 4.10:** Outflow momentum flux against velocity for our simulation suite at  $t=1$  Myr for  $v_{\min}=10$  km s $^{-1}$  (solid points) and  $v_{\min}=100$  km s $^{-1}$  (hollow points). Blue/red represent the cold/hot phases and the shape of the point shows different AGN luminosities. Also shown are analytic expectations for energy- and momentum conservation (grey lines) based on our input velocity of  $v_{\text{AGN}}=10^4$  km s $^{-1}$ , with the dotted line showing the energy-conserving case for  $v_{\text{AGN}}=3 \times 10^4$  km s $^{-1}$ . These define the energy-driven (pink) and the momentum- or energy driven (orange) regimes. The black circles and green squares show an observational compilation from [Fiore et al. \(2017\)](#) for comparison, representing molecular and ionised gas outflows, respectively (although we caution against a direct comparison with our hot/cold phase split). We can see that, despite the outflow being energy-driven overall, measurements in just the cold phase can result in momentum fluxes  $\dot{p}/(L/c) < 1$ .

### Energy & momentum conservation

Figure 4.10 shows momentum flux versus outflow velocity; a plane commonly used by observers to infer whether the measured outflow is momentum- or energy-conserving (e.g., [Bischetti et al., 2019](#); [Marasco et al., 2020](#); [Longinotti et al., 2023](#); [Bischetti et al., 2024](#)). The grey lines represent analytic expectations, with the horizontal line showing the momentum-conserving case of  $\dot{p}/(L/c) = 1$ . The diagonal lines show a maximally energy-conserving outflow ( $\dot{p}/(L/c) = v_{\text{OF}}/v_{\text{AGN}}$ ), representing the maximum momentum flux the large-scale outflow would have if it converted all the energy from the small-scale AGN wind. The solid diagonal shows the expectation for our fiducial AGN wind velocity of  $v_{\text{AGN}}=10^4$  km s $^{-1}$  and the dashed diagonal shows how this changes for a higher small-scale wind velocity of  $v_{\text{AGN}}=3 \times 10^4$  km s $^{-1}$ . We show the momentum flux of our small-scale wind as the red triangle at  $v_{\text{AGN}}=10^4$  km s $^{-1}$ . The other blue/red points show the momentum flux for the cold/hot gas at  $t=1$  Myr, calculated globally using Equation 4.5. The outflow velocity is computed as a mass-weighted mean of the radial velocity for gas above a cut of  $v_{\min}=10$  km s $^{-1}$  (solid points) and  $v_{\min}=100$  km s $^{-1}$  (hollow points).

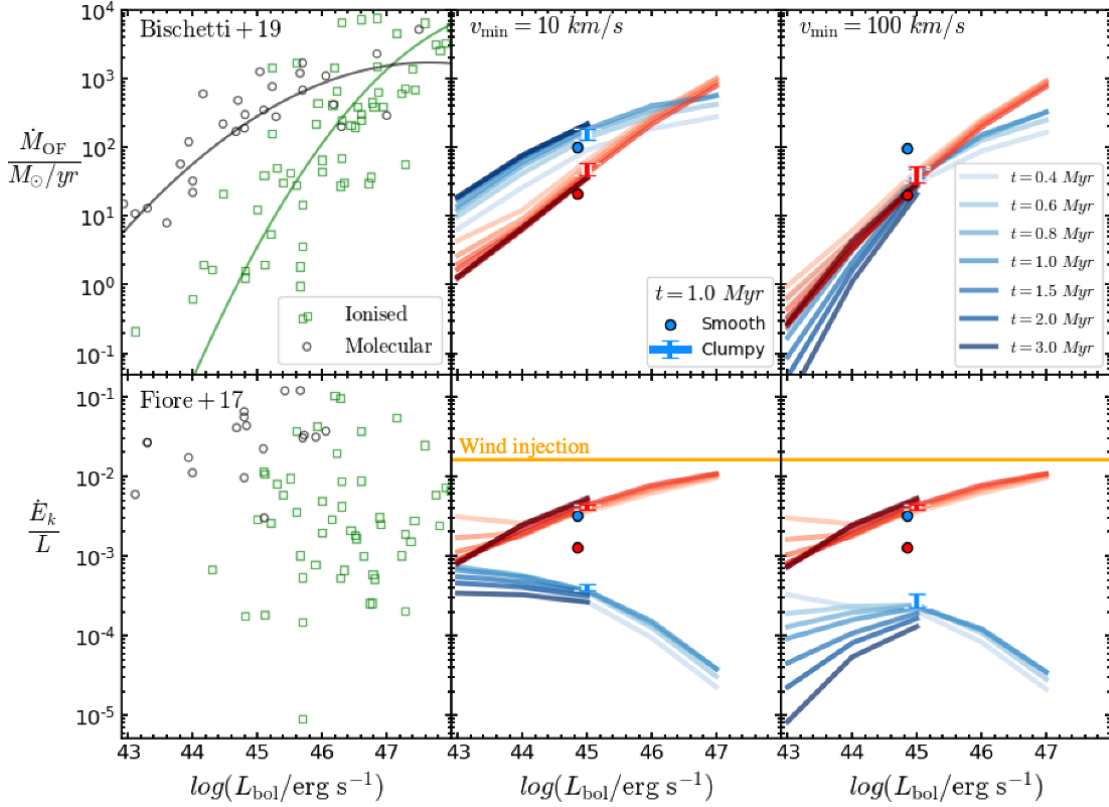
For the fiducial simulation, the momentum fluxes are in the range  $\dot{p}/(L/c) \approx 2 - 3$  for both the hot and cold phases, dropping to  $\dot{p}/(L/c) \approx 1$  in the cold case for the  $v_{\min}=100$  km s $^{-1}$  cut. As shown in Figure 4.6, the hot phase has a higher velocity than the cold phase, with

a mass-weighted mean velocity a factor of 4 – 5 higher. We show results for a range of AGN luminosities, all with the fiducial disc parameters (Table 4.1). The highest value for  $\dot{p}/(L/c)$  is measured in the lowest luminosity ( $L_{\text{bol}}=10^{43}$  erg s $^{-1}$ ) simulation, with  $\dot{p}/(L/c)\approx 9$  in the cold phase. However, the velocities at these faint luminosities are low, resulting in a much greater difference when we vary  $v_{\text{min}}$ . Even for the fiducial luminosity ( $L_{\text{bol}}=10^{45}$  erg s $^{-1}$ ), increasing the velocity cut to  $v_{\text{min}}=100$  km s $^{-1}$  results in the cold outflow having  $\dot{p}/(L/c)\lesssim 1$ , which would result in it being misinterpreted as momentum-driven. Additionally, the circular points show the results for the smooth case. Here, both phases have approximately the same velocity, and the cold phase dominates the momentum flux with  $\dot{p}/(L/c)\approx 6$ . We note that none of our simulated outflows reach the analytic expectation for maximal energy conservation (diagonal grey line). The closest point is the cold phase in the smooth outflow and this is still a factor of 5 times lower than the maximal value. Even if the global outflow produced by our model is energy-driven (King, 2003; Costa et al., 2020), the global momentum flux of the cold component can be low ( $\dot{p}/(L/c)\lesssim 1$ ), as it couples weakly to the clumpy ISM.

In Figure 4.10 we show an observational compilation from Fiore et al. (2017) for comparison, with the black circles and green squares representing molecular and ionised outflows, respectively. For clarity, we do not show error bars on the observational points, but these are generally large, spanning around an order of magnitude. We are not attempting a one-to-one correspondence with the observations, so caution against a direct comparison. In particular, we stress that our hot/cold temperature split does not map directly to observed molecular/ionised phases, but should be considered as just a generalised trend for gas at different temperatures. Additionally, in this sample, Fiore et al. (2017) use the maximum outflow velocity to compensate for orientation effects, which differs from our mass-weighted mean approach. However, we can observe some general trends with gas phase. For example, we can see that the ionised outflows span a large range of  $\dot{p}/(L/c)$ , but the molecular outflows all have higher momentum fluxes, at or near the energy-conserving relation. This is in contrast to our simulations where we find that the inhomogeneous structure of the ISM results in weak coupling to the cold gas phase, resulting in modest momentum fluxes and lower outflow velocities. It is surprising therefore that these observed molecular outflows have such high momentum fluxes. This could be due to the observational assumption that all the gas is moving at the same velocity in a thin shell, which is in tension with our results which find entrained cold clumps outflowing at a range of radii. We note that even in an inhomogeneous ISM, there may be local conditions where the ambient medium behaves homogeneously, causing the outflow to propagate in a shell; for example if the AGN is embedded within a large cloud (see discussion in Bieri et al. 2017), or if the cold phase is arranged in a ‘mist’ (Gronke & Oh, 2020b) with a high covering factor that efficiently traps the wind (equivalent to an even smaller  $\lambda_{\text{max}}$  than in our simulations). However, overall, our results suggest that the cold outflowing phase does not readily show a shell-like morphology for the parameter space explored in this study.

### 4.4.3 Outflow Scaling Relations with AGN Luminosity

One of the ways in which observational studies assess the role of AGN in launching outflows and the potential impact they have on galaxy evolution is to construct scaling relations between outflow, AGN and galaxy properties (Fiore et al., 2017; González-Alfonso et al., 2017; Leung et al., 2019; Bischetti et al., 2019; Lamperti et al., 2022; Musiimenta et al., 2023). Some simulations have used these observational constraints as an input for their AGN models (Rennehan et al., 2024). As we have discussed, there are many difficulties and uncertainties when measuring these quantities related to characterising outflow velocities and masses (Section 4.4.1), radii (Section 4.4.1), contributions from different phases (Section 4.4.1), and variations in time (Section 4.4.2) and location (Section



**Figure 4.11:** Scaling relations between the AGN luminosity, mass outflow rate and kinetic energy coupling efficiency. The blue/red lines show the cold/hot outflow in the fiducial simulation with darker lines showing later times. We note that the higher luminosity simulations ( $L_{\text{bol}} \geq 10^{46}$  erg s $^{-1}$ ) were only evolved to  $t=1$  Myr due to computational and boxsize constraints). The errorbars show the minimum and maximum values across all the clumpy simulations at  $t=1$  Myr to show the expected scatter in properties due to the intrinsic ISM structure of the galaxy. The circular points show the result from the smooth case. The orange line (bottom panels) shows the small-scale injected energy rate ( $\dot{E}_k/L = 1.7\%$ ). On the left, we show an compiled observational sample from [Fiore et al. \(2017\)](#); [Bischetti et al. \(2019\)](#) for comparison, although we note that our hot/cold temperature split should not be directly compared to observed ionised/molecular phases. Our simulations predict a positive scaling relation between the mass outflow rate and AGN luminosity.

4.4.2). The fact that AGN vary faster than the outflow properties ([Zubovas & Nardini, 2020](#)), and various target selection effects (see discussion in [Ramos Almeida et al., 2022](#); [Harrison & Ramos Almeida, 2024](#)) both pose additional challenges. There is therefore still debate about whether such scaling relations exist, with some studies finding them ([Fiore et al., 2017](#); [Musiimenta et al., 2023](#)) and others not finding tight correlations ([Davies et al., 2020](#); [Lamperti et al., 2022](#)).

### Simulations predict scaling relations

Using the results of our fiducial simulations performed over the AGN luminosity range  $L_{\text{bol}}=10^{43-47}$  erg s $^{-1}$ , we calculate global mass outflow rates and kinetic coupling efficiencies (following Section 4.2.6). We also explore the impact of variations in the initial disc clumpiness and explore results for a smooth medium at a single luminosity of  $L_{\text{bol}}=10^{45}$  erg s $^{-1}$ . We present these results in Figure 4.11, with the top row showing the mass outflow rate and the bottom showing the kinetic coupling efficiency as a function of AGN luminosity. The columns show our different outflow definitions: the middle shows the low velocity cut ( $v_{\text{min}}=10$  km s $^{-1}$ ) and the right shows the higher velocity

( $v_{\min}=100 \text{ km s}^{-1}$ ) outflow definitions. The solid lines show the results from our fiducial simulation, with the blue/red colour showing the cold/hot temperature cut and the gradient of the line showing the time evolution, with later times shown darker.

Looking at the middle column in the top row ( $v_{\min}=10 \text{ km s}^{-1}$ ), we can see that both gas phases have higher mass outflow rates with increasing  $L_{\text{bol}}$  (see also [Costa et al., 2020](#)). The hot phase shows the steepest correlation with the mass outflow rate increasing by 3 dex over the 4 dex luminosity range. The cold outflow rate also increases strongly at low  $L_{\text{bol}}$ , but then begins to taper off, reaching a more pronounced turnover point at  $L_{\text{bol}} > 10^{46} \text{ erg s}^{-1}$ , which allows the hot phase to exceed it in outflow rate at the highest luminosities. This turnover possibly demonstrates a critical luminosity above which the AGN is too powerful to allow the cold clumps to survive in the wind. This result echoes [Zubovas & Bourne \(2017\)](#), who identified a critical AGN luminosity of  $L_{\text{AGN}} \approx 5 \times 10^{46} \text{ erg s}^{-1}$  at which AGN-induced cloud fragmentation is maximally effective and above which the wind ejects gas too efficiently for the gas to cool and fragment. There is some subtle time evolution, with the mass outflow rate in the hot phase decreasing with time and the cold phase increasing over the period  $t \approx 0.5 - 3 \text{ Myr}$  (see Figure 4.9). However, this time evolution is subdominant compared to the overall positive trend with  $L_{\text{bol}}$ .

For the high velocity cut  $v_{\min}=100 \text{ km s}^{-1}$ , the same qualitative trends remain, with both phases showing a positive correlation with an even steeper gradient. Raising the velocity cut has an even greater effect on the low-luminosity ( $L_{\text{bol}}=10^{43-44} \text{ erg s}^{-1}$ ) systems, with the cold mass outflow rate dropping by two orders of magnitude. However, at the brighter end ( $L_{\text{bol}} > 10^{45} \text{ erg s}^{-1}$ ),  $v_{\min}$  plays less of a role and we still see the cold phase turnover at the same luminosity of  $L_{\text{bol}}=10^{46} \text{ erg s}^{-1}$ .

In Figure 4.11 we also show the results for an initially smooth medium as circular points. As shown in Figure 4.7, a smooth medium results in a factor  $\approx 2$  lower mass outflow rate. However, because the characteristic velocity of the resulting shell is higher than both choices of  $v_{\min}$ , the smooth case is unaffected by raising the minimum velocity cut (unlike the clumpy case). It thus results in a cold outflow rate that is a factor of 3 larger than the clumpy case for  $v_{\min}=100 \text{ km s}^{-1}$ . To evaluate the contribution of initial ISM clumpiness on the resulting outflow properties, we show the contributions to the scatter on these scaling relations as the blue and red errorbars on the plot. These represent the maximum/minimum mass outflow rate of any of the simulations with  $\lambda_{\text{max}}=40 - 330 \text{ pc}$  at  $t=1 \text{ Myr}$ . We can see that the change in mass outflow rate due to the initial clumpiness is smaller than the changes caused due to time evolution.

In the bottom row of Figure 4.11, the orange line shows the kinetic luminosity of the injected small-scale wind ( $\dot{E}_k/L = 1.7\%$ ). The hot phase (red) shows a slight positive correlation with  $L_{\text{bol}}$ , increasing by 1 dex across the full AGN luminosity range probed by our simulations. The cold phase (blue), however, shows a negative correlation, with the trend accelerating downwards at higher luminosities  $L_{\text{bol}} > 10^{45} \text{ erg s}^{-1}$ . When taking the higher radial velocity cut ( $v_{\min}=100 \text{ km s}^{-1}$ ), the kinetic coupling efficiency in lower  $L_{\text{bol}}$  systems shows a strong decline with time. This time variation is most prominent at fainter AGN luminosities, suggesting that the intrinsic scatter in observed scaling relations should be highest for  $L_{\text{bol}} \lesssim 10^{45} \text{ erg s}^{-1}$ . Finally, we note that the difference between the clumpy and smooth cases on the inferred outflow kinetic coupling efficiencies now becomes even more noticeable, with the cold phase kinetic coupling typically more than an order of magnitude higher in the homogeneous medium compared to the clumpy simulations.

### Observational comparison

In the left panel of Figure 4.11, we show observational compilations from [Fiore et al. \(2017\)](#) and [Bischetti et al. \(2019\)](#) split into measurements for molecular (black circles) and ionised phases (green squares). These have primarily been estimated based on CO and [O III] line emission, respectively. The molecular phase dominates the mass outflow rate at lower luminosities, but flattens off around  $L_{\text{bol}}=10^{46}$  erg s<sup>-1</sup>. The ionised phase increases consistently, becoming dominant in the highest- $L_{\text{bol}}$  systems. The observed kinetic luminosities do not show any obvious trend with  $L_{\text{bol}}$ . The molecular phase has high (> 1%) coupling efficiencies, up to  $\dot{E}_k/L \approx 10\%$ ; higher on average than the ionised phase, although there is limited overlap in this observational sample.

It is important to note that we cannot draw direct comparison between these molecular and ionised phases and our simulated cold/hot outflows, as our simulations do not account for low temperature cooling, molecular chemistry or radiative transfer. Instead we focus on qualitative trends for the two gas phases. Our simulations predict similar mass outflow rate trends, with our cold phase also dominating at lower  $L_{\text{bol}}$  before flattening off at  $L_{\text{bol}} > 10^{46}$  erg s<sup>-1</sup>. Our results for the kinetic coupling efficiency, however, differ from the observational trends. We find weak positive (hot phase) and negative (cold phase) trends with  $L_{\text{bol}}$ , which are not seen in the observations. Additionally, our values for  $\dot{E}_k/L\%$  are much lower than those observed. As we have shown, our wind couples weakly to the inhomogeneous ISM and we lose energy to cooling and mixing (Figure 4.5). The effect of this can be seen by noting that the cold phase in the smooth simulations has an  $\dot{E}_k/L\%$  an order of magnitude higher than the clumpy case. It is therefore surprising that such high ( $\dot{E}_k/L \approx 1 - 10\%$ ) kinetic coupling efficiencies are found in the observations where we may also expect inhomogeneous ISM conditions.

As we have discussed, there are many observational difficulties and uncertainties when measuring these quantities related to characterising outflow velocities and masses (Section 4.4.1); radii (Section 4.4.1); and variations in time (Section 4.4.2) and location (Section 4.4.2). These challenges could add scatter to the observational results, with additional scatter also being driven by physical differences across individual galaxies such as varying disk masses, ISM distribution and initial wind velocities (see Section 4.3.3). There are also additional uncertainties that affect outflow properties derived from observations, including constraining electron densities, and more generally, conversion factors between observed line fluxes and total gas masses (e.g., [Rose et al., 2018](#); [Lamperti et al., 2022](#); [Holden et al., 2023](#); [Holden & Tadhunter, 2023](#)), with more recent analysis suggesting that previous work overestimated the outflow rates in the ionised phase by a factor of a few or more, especially in the  $L_{\text{bol}} \lesssim 10^{45}$  erg s<sup>-1</sup> regime ([Davies et al., 2020](#)). Additionally, if the observed outflow is assumed to be a spherical shell, it could lead the bulk outflow velocity being overestimated, which could have a large impact on the inferred kinetic luminosity.

It is also possible that our simulations genuinely under-predict the outflow rate; for example, by neglecting additional driving mechanisms for the outflow, such as radiation pressure, cosmic rays or star-formation. Another possibility is that the small-scale wind should be even faster and more energetic than considered here. More massive and centrally-concentrated gas reservoirs or gas configurations with larger covering fractions than considered in our study may also result in more mass-loaded outflows. Missing physical ingredients, such as metal-line cooling, could increase the cold gas in the outflow.

In this section, we demonstrated that our simulations predict scaling relations between the mass outflow rate and AGN luminosity. We showed that the scatter on these relations can be affected by the time, disc clumpiness and the minimum radial velocity sensitivity. We compared our results to an observational sample, finding some similarities (such as

a turnover in the cold gas outflow rate at  $L_{\text{bol}} \gtrsim 10^{46}$  erg s $^{-1}$ ), but also that our simulations have significantly lower kinetic luminosities than those implied from the comparison observational sample.

## 4.5 Conclusions & Outlook

We performed controlled experiments simulating a physically-motivated AGN wind embedded in a clumpy ISM disc. By manually setting the initial ISM structure, we investigated the effect this has on the energetics and multiphase structure of the resulting outflow. We used the AREPO code (Springel, 2010b; Pakmor et al., 2016) and the AGN wind model BOLA (Costa et al., 2020) with AGN luminosities ( $L_{\text{AGN}} = 10^{43-47}$  erg s $^{-1}$ ). We divided our results into two main sections. In Section 4.3 we investigated the effect of an AGN wind on our clumpy setup and characterised the energetics of the resulting outflow. Our main findings are:

- **Multiscale structure:** the small-scale ultra-fast outflow (UFO) launches large outflow bubbles into the halo, reaching  $R = 6$  kpc by  $t = 5$  Myr (Figure 4.2). The initial disc inhomogeneities allow the hot gas to vent through high-velocity ‘chimneys’. The cold outflow is formed from small ( $10 - 20$  pc), dense ( $n \approx 10^3$  cm $^{-3}$ ) clouds. The venting of the hot gas creates a strong velocity differential between the two phases, with the hot gas streaming at up to  $v_r \approx 1000$  km s $^{-1}$ , but the bulk of the cold gas moving at  $v_r \approx 100$  km s $^{-1}$  (Figure 4.6). Despite the strong ram pressure this creates on the cold clouds, they are able to survive over the  $t = 5$  Myr timescale of the simulations, possibly due to efficient cooling at the phase boundary (Figure 4.2).
- **Multiphase gas energetics:** the outflow is clearly separated into two phases (Figure 4.4), which we define as ‘hot’ ( $T > 10^{4.5}$  K) and ‘cold’ ( $T < 10^{4.5}$  K). The cold outflow carries most of the mass with our fiducial case having  $\dot{M}_{\text{OF}} = 185$  M $_{\odot}$ /yr at  $t = 1$  Myr compared to the hot phase at  $\dot{M}_{\text{OF}} = 40$  M $_{\odot}$ /yr (both for  $v_{\text{min}} = 10$  km s $^{-1}$ ). However, because the hot phase has a much higher velocity, the momentum rates are split roughly evenly between the two phases. In clumpy media, the energy budget is dominated by the hot gas, with around an order of magnitude greater kinetic luminosity (Figure 4.5).
- **Comparison to smooth disc:** we also investigated our AGN wind in a homogeneous disc (Figure 4.3). We found that the resulting outflow differs significantly from the clumpy case, showing a much narrower spread in radial velocity, with a characteristic speed of  $v_r \approx 400$  km s $^{-1}$  in both phases (Figure 4.6). This results in the cold phase of the gas being much more energetic than in the clumpy setup, containing  $\sim 70\%$  of the kinetic luminosity of the system at  $t = 1$  Myr and a higher momentum flux of  $\dot{p}/(L/c) \approx 7$  (Figure 4.5).
- **Sensitivity to setup:** we tested the impact on varying the initial galaxy, for example, by reducing the initial density of the disc or altering the sizes of the initial clumps (Section 4.3.3). We found that the initial clump size made a modest impact, varying the mass outflow rates by up to around 30%. Changing the density and height of the disc reduced the mass outflow rates by around a factor of two.

In Section 4.4 we discussed the implications of our findings for observational studies of AGN outflows. In particular, we found:

- **Outflow measurements:** we discussed the difficulty of observationally measuring outflow properties such as the radius, mass and velocity (Section 4.4.1). Many observational studies assume a spherical shell-like outflow, but, as we have seen, an outflow originating from a more realistic, clumpy environment has very different morphology and energetics to a spherical shell. This could lead to incorrect assumptions about the outflow, for example, that all of the gas is moving at a characteristic velocity (Figure 4.6). This could result in the derived outflow rates being significantly overestimated, especially in colder gas phases.
- **Sensitivity to minimum outflow velocity:** we found that a major source of uncertainty in calculating outflow properties for the cold phase is the minimum radial velocity cut used to define gas as outflowing  $v_{\min}$ , which can lead to a factor of 8 difference in mass outflow rate between  $v_{\min}=10 \text{ km s}^{-1}$  and  $v_{\min}=100 \text{ km s}^{-1}$  (Figure 4.7). Many observational methods for calculating the outflow rates make an explicit or implicit assumption for  $v_{\min}$ . This is particularly challenging when trying to de-couple outflow from non-outflow kinematics (e.g., galaxy rotation) in colder gas phases. A greater value of  $v_{\min}$  will lead to a large proportion of the outflowing mass to be missed, with this effect proportionally much worse for the colder phases.
- **Inferring driving mechanisms and kinetic efficiencies:** we found that the, despite our outflow being energy-conserving overall, values of  $\dot{p}/(L/c) < 1$  (momentum-driven) could still be inferred if only a single phase was measured, the full radial extent of the outflow was not captured (Figure 4.8) or the outflow was observed past the momentum peak (Figure 4.9). This makes it difficult for observations to accurately determine the driving forces behind any outflow seen (Figure 4.10). Additionally, the derived kinetic energy coupling efficiencies were seen to be highly dependent on the phase, time and location of the outflow (Figures 4.8 and 4.9). This makes inferences about the efficiency of the wind from large-scale measurements of the kinetic coupling efficiency ( $\dot{E}_k/L$ ) challenging.
- **Scaling relations:** we found that our simulations predict a positive correlation between the mass outflow rate and AGN luminosity in both the hot and cold phase (Figure 4.11). The cold phase dominates at lower  $L_{\text{bol}}$ , but flattens off at  $L_{\text{bol}} > 10^{46} \text{ erg s}^{-1}$ . This turnover point is also seen in the observational compilation of [Bischetti et al. \(2019\)](#). However, we find lower kinetic coupling efficiencies than observed, especially in the cold phase. Future work warrants a more comprehensive comparison to observations, including accounting for systematic uncertainties and other potential sources of scatter (see Section 4.4.3).

The observation of cold clouds entrained in galactic outflows (e.g., [Di Teodoro et al., 2019](#); [Veilleux et al., 2020](#)) is puzzling as the typical cloud crushing timescale is smaller than the outflow timescale ([Klein et al., 1994](#); [Zhang et al., 2017](#); [Schneider & Robertson, 2017](#)). In our simulations, we do find cold gas clouds surviving on Myr timescales. Cold, dense gas can be fast, travelling at velocities up to  $800 \text{ km s}^{-1}$  (see Figure 4.4), but does not appear to reach the extreme velocities  $> 1000 \text{ km s}^{-1}$  suggested by some observations (e.g. [Lutz et al., 2020](#)). A full study of the formation and evolution of the cold clouds in our simulations is beyond the scope of this paper, but constitutes an important future work direction.

In this study we only considered primordial cooling down to  $T \approx 10^4 \text{ K}$ . Metal-line cooling boosts the cooling rate at  $T \approx 10^{5-7} \text{ K}$  and allows cooling to  $T < 10^4 \text{ K}$ . Our values for mass outflow rates in the cold component can therefore be considered conservative lower estimates. In particular, how much of our ‘cold phase’ eventually turns molecular remains to be understood. Our estimated outflows rates  $\approx 10^{2-3} M_{\odot}$  still fall short of



observational reports of outflow rates  $> 10^3 M_\odot$  (Fiore et al., 2017) even for our brightest simulated AGN. Performing new simulations with metal-line and low temperature cooling is will thus be important in future studies. Furthermore, our initial cold clumps have densities up to  $n \lesssim 10^3 \text{ cm}^{-3}$  which may not capture the highest densities seen in the cores of molecular clouds ( $n \approx 10^{2-6} \text{ cm}^{-3}$ ; Ferriere 2001). We showed in Figure 4.7 that reducing the mean initial density resulted in less cold outflowing gas, but slightly more hot gas. Raising the density could therefore increase the cold phase outflow rate, which could result in a closer match to the observations (Figure 4.11), however, increasing the density of the disc is unlikely to result in a higher outflow velocity for the cold gas, which is the main cause of the differences we see with the observations.


There are additional physical processes neglected here that will have to be examined in future studies. These include magnetic fields, AGN radiation and cosmic rays. Previous work on single-cloud simulations including magneto-hydrodynamics, for instance, has found that magnetic fields can have a range of effects, for example, enhancing thermal instabilities (Ji et al., 2018), or perhaps suppressing them (Gronke & Oh, 2020a). In a recent study, Hidalgo-Pineda et al. (2024) found that the interplay between magnetic fields and radiative cooling can reduce the size of entrained cold clouds and lead to more rapid entrainment. In the case of cosmic rays, the pressured built around cold clouds via the ‘bottleneck effect’ may play a major role in producing cold gas outflows (e.g. Brüggén & Scannapieco, 2020).

Since we find radiative cooling operates both in shocks and in mixing layers (Figure 4.2), it will be important to further understand the observational imprints of the associated cooling emission. Since mixed gas has a temperature of  $T \approx 10^6$  K in our simulations, it is possible this produces extended X-ray emission. Interestingly, there are X-ray ‘chimneys’ in our own Milky Way linking the galactic nucleus to the Fermi bubbles (Ponti et al., 2019, 2021) resembling the low-density gaps through which hot gas vents through in our simulations. Generating multi-wavelength predictions based on our simulations will be important to test AGN feedback through winds.

## 4.6 Chapter Appendix: Numerical Convergence

To test the numerical convergence of our results, we performed simulations at target mass resolutions of  $M_{\text{target}}=1000 M_\odot$  (‘low resolution’),  $M_{\text{target}}=100 M_\odot$  (‘fiducial resolution’) and  $M_{\text{target}}=10 M_\odot$  (‘high resolution’). These were performed with our fiducial parameters, i.e., a disc clumpiness of  $\lambda_{\text{max}}=170 \text{ pc}$  ( $k_{\text{min}}=12$ ) and AGN luminosity of  $L_{\text{AGN}}=10^{45} \text{ erg s}^{-1}$  (see Table 4.1 for a full description of our fiducial model). Due to the significant computational cost of the high resolution simulation, we only perform it to  $t=1 \text{ Myr}$ . Figure 4.12 shows the distribution of the cell diameters for the three resolution test simulations, assuming a spherical geometry for the Voronoi cells. Improving the mass resolution by a factor of 10 should improve the spatial resolution by a factor of  $\sqrt[3]{10} \approx 2.2$  which is broadly in line with what we see here. Of particular interest is the minimum cell diameter probed as this will be in the densest gas where we probe the structure of the entrained cold gas clouds. The low/fiducial/high resolution simulations have minimum spatial resolutions of  $d_{\text{cell,min}} \approx 1.5, 0.5, 0.2 \text{ pc}$  respectively. Our fiducial simulation has higher resolution than other similar simulations of AGN interacting in idealised clumpy discs (e.g., Mukherjee et al. 2018 and Tanner & Weaver 2022 use grid-based systems with resolutions of 6 pc and 10 pc respectively) and our high-resolution simulation offers unprecedented resolution for galaxy-scale simulations.

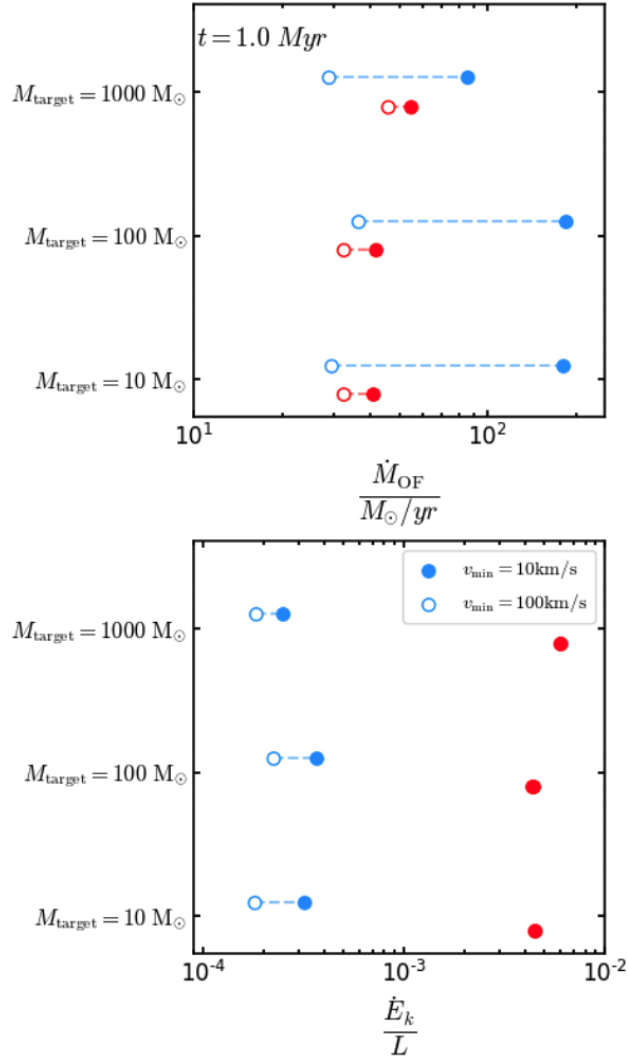
To investigate the numerical convergence of our global outflow properties, we computed the mass outflow rate and kinetic energy coupling efficiency for our three resolution tests, shown in Figure 4.13. The symbols are the same as Figure 4.7, with blue/red showing the cold/hot phase and the filled/hollow markers showing the effect of varying the minimum



**Figure 4.12:** The distribution of cell diameters in our three resolution test simulations at  $t=1$  Myr. The minimum cell diameter for each resolution is roughly  $d_{\text{cell,min}} \approx 1.5, 0.5, 0.2$  pc, for our target mass resolutions of  $M_{\text{target}}=1000, 100, 10 M_{\odot}$  respectively.

outflow velocity ( $v_{\text{min}}$ ). We can see that the low resolution simulations (top) has lower mass outflow rates and kinetic energies in the cold phase than the fiducial simulation (middle), but slightly higher values in the hot phase. This could be because the larger cold gas clumps seen in the low resolution simulation (Figure 4.14) are harder to accelerate and have a lower surface area resulting in less phase mixing and cooling than in the fiducial simulation. However, the differences between the fiducial simulation and the high resolution simulation (bottom) are much smaller: the hot phase is essentially the same, and the cold phase only differs by a slight decrease in mass outflow rate for  $v_{\text{min}}=100 \text{ km s}^{-1}$ . This demonstrates that the global outflow properties are well-converged at our fiducial resolution.

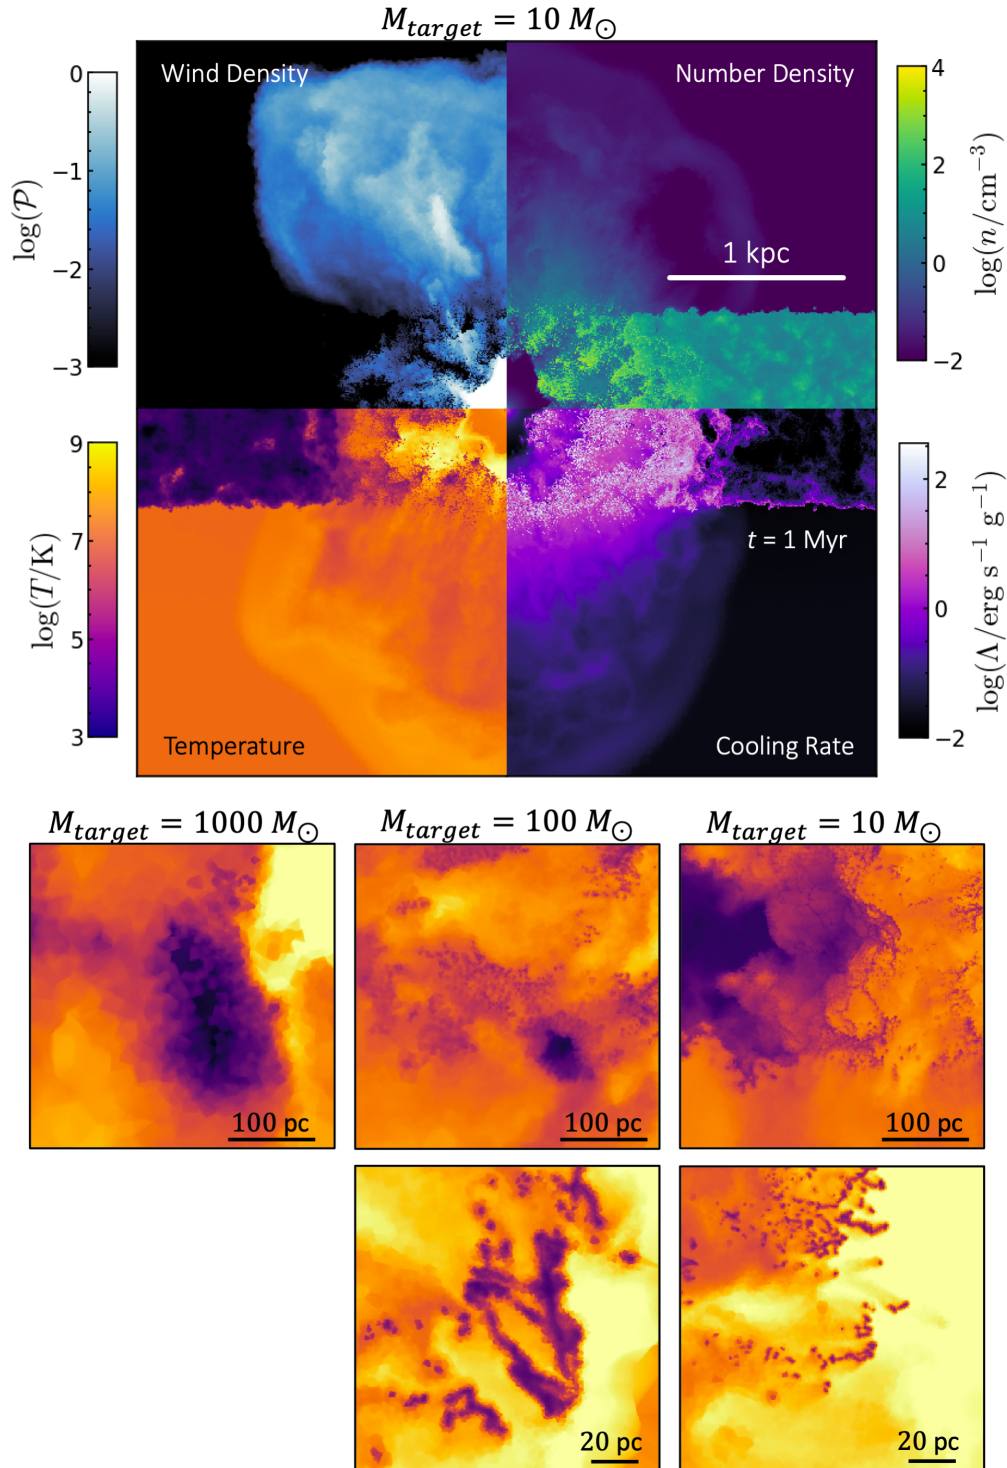
In Figure 4.14 we show a similar plot to Figure 4.2 for our high-resolution simulation at  $t=1$  Myr. The top section shows an edge-on view of our disc showing (clockwise from top left) the wind density, number density, cooling rate and temperature. At this kpc-scale, the high-resolution outflow looks broadly similar to the fiducial resolution seen in Figure 4.2. This corroborates our finding in Figure 4.13 that the global outflow is well-converged at our fiducial resolution. However, the bottom section of this panel shows the morphology of the entrained cold clouds the three different resolutions (columns) at different spatial scales (rows). We can see that increasing the resolution decreases the size of the smallest cloudlets seen: at our fiducial resolution ( $M_{\text{target}}=100 M_{\odot}$ ; middle column), the smallest clouds are on the scale of  $\approx 10$  pc whereas in the high resolution simulation ( $M_{\text{target}}=10 M_{\odot}$ ; right column), they can be as small as  $\approx 1$  pc. The shapes of the cold gas clouds are similar, showing initially dense cores surrounded by smaller fragments and filaments, but the scale clearly varies with resolution. However, even high-resolution simulations of single clouds interacting with a wind do not always find convergence, with Yirak et al. (2010) finding that a ratio of cloud-to-resolution element ratio of  $r_{\text{cl}}/d_{\text{cell}} > 100$  is not always sufficient for self-convergence. Conversely, Gronke & Oh (2020a) and Banda-Barragán et al. (2021) argue that their global quantities (e.g., mass entrainment rate) converge at  $r_{\text{cl}}/d_{\text{cell}} \geq 8$ , even if the exact morphology of the cold fragments requires much higher resolution to converge. This feature of global properties converging while the structure of the gas continues to granulate may be a common and unavoidable feature of the cold gas phase (Hummels et al., 2019; van de Voort et al., 2019; Nelson et al., 2020). Nevertheless, as we have shown, the exact spatial structure of the cold clouds does not affect the global outflow properties, which is the focus of this work.



**Figure 4.13:** Convergence of global properties, adapted from Figure 4.7. We show the mass outflow rate (top) and kinetic luminosity (bottom) for the three resolutions ( $M_{\text{target}}=10, 100, 1000 M_{\odot}$ ). The blue/red points represent our cold/hot phases, and the solid/hollow points show minimum radial velocity cuts of  $v_{\text{min}}=10 \text{ km s}^{-1}$  and  $v_{\text{min}}=100 \text{ km s}^{-1}$ , respectively. We can see that there is little change in the global outflow properties between the fiducial and high resolution simulations, demonstrating these properties have converged at  $M_{\text{target}}=100 M_{\odot}$ .

## Acknowledgements

The author acknowledges computing facilities from the Computational Center for Particle and Astrophysics (C2PAP), part of the ORIGINS Excellence Cluster. ORIGINS is funded by the Deutsche Forschungsgemeinschaft (DFG; German Research Foundation) under Germany’s Excellence Strategy: EXC-2094-390783311. This work also used the DiRAC Memory Intensive service (Cosma8) at Durham University, managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility ([www.dirac.ac.uk](http://www.dirac.ac.uk)). The DiRAC service at Durham was funded by BEIS, UKRI and STFC capital funding, Durham University and STFC operations grants. DiRAC is part of the UKRI Digital Research Infrastructure.



**Figure 4.14:** The results of our high-resolution simulation ( $M_{\text{target}}=10 M_{\odot}$ ). The top panel shows the galaxy edge-on at  $t=1 \text{ Myr}$ . We show (clockwise from top left) the wind tracer density, gas number density, cooling rate and temperature. The outflow at this spatial scale is broadly similar to that of our fiducial resolution, shown in Figure 4.2. At the bottom of the Figure, we show a comparison of the structure of the small-scale cold gas clouds for the three resolution simulations. The columns show the resolution ( $M_{\text{target}}=1000, 100, 10 M_{\odot}$ ) and the rows show two different spatial scales, shown by the scale bars. The fiducial resolution simulation shows cold clouds of sizes down to  $\approx 10 \text{ pc}$ , but in the high resolution simulation has even smaller clouds, down to  $\approx 1 \text{ pc}$ . This emphasises the difficulty in resolving the cold gas phase in galaxy-scale simulations.

# 5

## **X-ray Emission from Wind-ISM Interactions**

**Mixing between AGN winds and ISM clouds produces  
luminous Bremsstrahlung emission**

**Samuel Ruthven Ward; T.Costa; C.M.Harrison; V.Mainieri**  
(in prep.)

---

Active galactic nuclei (AGN) drive powerful, multiphase outflows into their host galaxies which are expected to play a key role in galaxy evolution. Despite containing a large fraction of the energy of these outflows, the hot ionised phase ( $T \gtrsim 10^6$  K) is difficult to observe due to its low density and long cooling times. However, in a previous work featuring an AGN wind embedded in an idealised galaxy disc with a manually-distributed clumpy ISM, we found that interactions between the AGN wind and the ISM clumps resulted in enhanced radiative cooling as the wind mixed with the clumps. This motivated us to investigate the X-ray emission associated with this process. In this paper, we find that this mixed material makes a dominant contribution to the Bremsstrahlung emission of the outflow, in contrast to analytic expectations from a homogeneous medium which predict that the shocked ambient medium is the dominant origin of the X-rays. This mixing-induced free-free emission is strongest in the part of the outflow propagating equatorially through the disc. The size of the emitting region reaches a diameter of  $\approx 3$  kpc by  $t=3$  Myr. This extended region can be resolved up to distances of  $d_A \approx 400$  Mpc ( $z \approx 0.11$ ; assuming an instrumental resolution of  $0.5''$ ), which would allow it to be distinguished from the point-source nuclear emission. We also show that, for  $L_{\text{AGN}} \gtrsim 10^{45}$  erg s $^{-1}$ , the total free-free emission exceeds the X-ray contribution due to star formation for reasonable estimates of the SFR of local quasars. We predict a strong scaling between the free-free emission and AGN luminosity ( $L_X \propto L_{\text{AGN}}$ ). Finally, we create synthetic *Chandra* images to demonstrate that this emission is observable in nearby galaxies.

---

## 5.1 Introduction

Feedback from active galactic nuclei (AGN) has become an essential component of modern cosmological models and simulations of galaxy evolution (e.g., [Springel et al., 2005a](#); [Somerville et al., 2008](#); [Schaye et al., 2015](#); [Dubois et al., 2016](#); [Weinberger et al., 2018](#); [Davé et al., 2019](#)), with the energy released from accreting supermassive black holes (SMBHs) thought to play a key role in quenching star formation in galaxies. However, a key uncertainty in models of feedback is how efficiently the energy from the AGN can couple to the galaxy’s multiphase interstellar medium (ISM). Constraining this coupling efficiency is therefore key to understanding the impact of AGN on their host galaxies, but extrapolating observed quantities to measure the coupling is highly challenging (see discussions in [Harrison et al., 2018](#); [Harrison & Ramos Almeida, 2024](#); [Ward et al., 2024](#)).

To assess the impact of AGN, studies target kiloparsec-scale outflows, which are launched when energy released by the AGN in the form of accretion disc winds, jets and/or radiation pressure sweeps up the surrounding ISM and accelerates it to high velocities ( $v_r \gtrsim 100 \text{ km s}^{-1}$ ). These outflows are multiphase, containing entrained clouds of molecular or neutral gas (e.g., [González-Alfonso et al., 2017](#); [Veilleux et al., 2020](#); [Lamperti et al., 2022](#)), warm ionised gas (e.g., [Harrison et al., 2014](#); [Rose et al., 2018](#); [Molyneux et al., 2019](#)), and hot X-ray emitting gas (e.g., [Tombesi et al., 2013](#); [Lansbury et al., 2018](#)). Therefore, to fully characterise the outflow, and thus infer its effect on the galaxy, a multiwavelength approach is needed to capture all the mass and energy ([Cicone et al., 2018](#); [Girdhar et al., 2022](#); [Harrison & Ramos Almeida, 2024](#)). In particular, the hot ionised phase ( $T \gtrsim 10^6 \text{ K}$ ) is expected to dominate the energetics of the outflow, although the long cooling time and low density ( $n \approx 10^{-3} \text{ cm}^{-3}$ ) of this phase means it is not expected to radiate efficiently. However, in our previous work ([Ward et al., 2024](#)), we found that mixing between the ISM and an AGN wind enhanced the radiative cooling of the hot outflow, which could result in observable X-ray emission. Investigating this mixing-driven X-ray emission is the focus of this paper.

There have been several observational successes in detecting kiloparsec-scale hot gas in AGN outflows using X-ray imaging, mostly from the *Chandra* space telescope (e.g., [Croston et al., 2008](#); [Greene et al., 2014](#); [Di Gesu et al., 2017](#); [Lansbury et al., 2018](#)). However it is not clear whether the emission seen in these systems is caused by cooling from a shock-heated outflow bubble or photoionisation caused by an AGN or star formation (e.g., [Wang et al., 2010](#); [Somalwar et al., 2020](#); [Wang et al., 2024a](#)). The Milky Way also contains kiloparsec-scale X-ray bubbles ([Predehl et al., 2020](#)) as well as narrow ‘chimneys’ linking them to the galactic centre ([Ponti et al., 2019, 2021](#)), making even our own galaxy a potential test-bed for signatures of AGN feedback. However, there are very limited direct predictions from different models or simulations of what the expected observational signatures are of the X-ray emitting hot phase and its relative contribution to total X-ray emission compared to other relevant processes. There have been some numerical works looking at predicted X-ray emission in the CGM of galaxies ([Pillepich et al., 2021](#)) including at high ( $z > 6$ ) redshift ([Costa et al., 2014a](#); [Bennett et al., 2024](#)) but these are based on cosmological simulations with heuristic models for AGN feedback.

A strong candidate for creating the kiloparsec-scale outflows seen in observations (including in X-rays) are accretion disc winds. In this picture, the accretion disc launches a small-scale wind which interacts with the surrounding medium, shocking it to high temperatures, and creating a large-scale outflow. An analytic model for this process was introduced in [King \(2003, 2005\)](#) and developed further in [Zubovas & King \(2012\)](#); [Faucher-Giguère & Quataert \(2012\)](#); [Costa et al. \(2014b\)](#). A numerical implementation in the hydrodynamic code AREPO was presented in [Costa et al. \(2020\)](#). An important contribution was made by [Nims et al. \(2015\)](#) who provided analytic arguments for observational signatures

of this model, discussing both thermal (inverse-Compton and Bremsstrahlung) and non-thermal (synchrotron, non-thermal inverse-Compton and pion decay) mechanisms. They concluded that the shocked ambient medium should produce observable X-ray emission via Bremsstrahlung radiation, on scales of a few kiloparsecs. However, these analytic results assume the wind propagates in an isotropic, homogeneous medium. As shown in [Ward et al. \(2024\)](#), outflow properties can differ significantly from a homogeneous case when the parsec-scale structure of the ISM is taken into account. Thus it is unclear what the effect of a multiphase ISM structure would be on the resulting X-ray emission.

In this study, we investigate the Bremsstrahlung emission from interactions between an AGN wind and a clumpy ISM, building on our previous work ([Ward et al. 2024](#); Paper I) where we introduced a series of numerical experiments featuring an AGN wind model embedded in both a smooth and clumpy ISM. This paper is structured as follows: in Section 5.2 we summarise the simulations first presented in Paper I and describe how we estimate the X-ray emission and construct mock X-ray images; in Section 5.3 we show the main results of our work; in Section 5.4 we present synthetic *Chandra* observations, and discuss the observational implications of our findings; and finally in Section 5.5 we summarise our results and conclude with ideas for further investigation. We assume a flat,  $\Lambda$ CDM cosmology throughout, using values from [Planck Collaboration \(2016\)](#) of  $H_0 = 67.7 \text{ kms}^{-1} \text{ Mpc}^{-1}$ ,  $\Omega_m = 0.3$  and  $\Omega_\Lambda = 0.7$ .

## 5.2 Methods

In this Section, we briefly describe the setup of the simulation suite and then discuss how the X-ray emission from the outflow can be estimated. For full details of the simulations, the reader is referred to [Ward et al. \(2024\)](#).

### 5.2.1 The ACDC Simulations

In this study, we use the ACDC (AGN in Clumpy DisCs) simulation suite presented in Paper I. These feature an isotropic AGN wind situated within a galaxy disc with a clumpy ISM sub-structure. This idealised setup allowed us to perform a series of controlled experiments investigating how disc sub-structure affects the propagation of multiphase outflows. The simulations are performed using the moving-mesh hydrodynamic code AREPO ([Springel, 2010b](#)) which features an unstructured Voronoi mesh that moves with the fluid and refines/de-refines in regions of high/low density. An exact Riemann solver is used at the cell interfaces to calculate hydrodynamic fluxes between each cell ([Pakmor et al., 2016](#)). This yields both accurate shock-capturing and high spatial resolution in dense regions.

#### Clumpy disc setup

It is computationally challenging to produce a realistic ISM structure in galaxy simulations. Therefore we follow studies such as [Sutherland & Bicknell \(2007\)](#); [Cooper et al. \(2008\)](#); [Wagner & Bicknell \(2011\)](#); [Mukherjee et al. \(2016\)](#); [Bieri et al. \(2017\)](#); [Banda-Barragán et al. \(2020\)](#); [Tanner & Weaver \(2022\)](#) in manually setting the spatial distribution and phase structure of the ISM. This is achieved using the PyFC<sup>1</sup> package ([Wagner et al., 2012](#)) to create a log-normal density distribution with a fractal spatial distribution ([Lewis & Austin, 2002](#); [Sutherland & Bicknell, 2007](#)). This creates a random distribution of cold, dense clouds, with the clump size parameterised by the average largest cloud size,  $\lambda_{\text{max}}$ . We investigate three such cloud sizes:  $\lambda_{\text{max}}=40 \text{ pc}$  (‘small’ clouds);  $\lambda_{\text{max}}=170 \text{ pc}$  (‘medium’ clouds); and  $\lambda_{\text{max}}=330 \text{ pc}$  (‘large’ clouds). We also show the results for a ‘smooth’ disc which has a homogeneous density distribution for comparison.

<sup>1</sup><https://pypi.org/project/pyFC/>



We use a disc with diameter 4 kpc and height 1 kpc. The initial mean density of the disc is set to  $\langle n_0 \rangle = 5 \text{ cm}^{-3}$  and the mean temperature to  $\langle T_{0,\text{disc}} \rangle = 10^4 \text{ K}$ , leading to initial clumps with densities  $n_0 = 10^{1-3} \text{ cm}^{-3}$ . Cells with temperatures above  $T_{\text{crit}} = 3 \times 10^4 \text{ K}$  are replaced by hot background gas, to generate porosity in the disc. The background is a static halo in pressure equilibrium with the disc and an initial constant temperature and density of  $T_{0,\text{bkg}} = 10^7 \text{ K}$  and  $n_{0,\text{bkg}} = 10^{-2} \text{ cm}^{-3}$ . Primordial cooling is included and gravity is neglected (see [Ward et al. 2024](#) for a justification of this idealised setup).

Our fiducial simulations have a mass resolution of  $M_{\text{target}}=100 M_{\odot}$  which gives us a spatial resolution down to  $d_{\text{cell}} \approx 1 \text{ pc}$  in the highest density regions. In Paper I, we performed a numerical convergence test and found that global outflow properties are well-converged at our fiducial resolution.

### AGN wind model

We use BOLA ([Costa et al., 2020](#)) to generate a fast, small-scale wind. Two spherical shells of cells are fixed in place at the centre of the disc and mass, momentum and energy are then injected across the boundary of these cells. We model a spherical, ultra-fast outflow, namely a wind velocity of  $v_{\text{AGN}} = 10^4 \text{ km s}^{-1}$ , a momentum boost factor of  $\tau = \dot{p}/(L_{\text{AGN}}/c) = 1$ , and a fiducial AGN luminosity of  $L_{\text{AGN}}=10^{45} \text{ erg s}^{-1}$  (we investigate the effect of AGN luminosity on the resulting X-ray emission in Sections 5.3.5 & 5.3.6). A passive scalar is also injected along with the wind ( $\mathcal{P}$ ) in order to trace this component throughout time. This parameter represents the mass fraction of a cell that originated from the wind mass injection. An additional cell refinement scheme is imposed to refine the cells further in regions of high wind tracer density, to increase the resolution in the outflow by a factor of 10.

### 5.2.2 Bremsstrahlung Emission

The hot ionised gas phase ( $T \approx 10^{5-8} \text{ K}$ ) is expected to emit X-rays via thermal processes. The most important of these is Bremsstrahlung (free-free) radiation, emitted by free electrons interacting with charged ions. To estimate the bolometric X-ray luminosity from the hot gas, we use the Bremsstrahlung approximation (see also [Sijacki & Springel, 2006](#); [Bennett & Sijacki, 2022](#))

$$L_X = 1.4 \times 10^{27} T^{\frac{1}{2}} n_e n_i Z^2 \bar{g} \quad [\text{erg s}^{-1}] \quad (5.1)$$

where  $T$  is the gas temperature,  $n_e$  and  $n_i$  are the electron and ion number density, respectively,  $Z$  is the mean charge of the ions, and  $\bar{g} \approx 1$  is the Gaunt factor. This approximation assumes that Bremsstrahlung emission is the dominant source of radiation from the gas, neglecting thermal inverse-Compton scattering and metal-lines. [Nims et al. \(2015\)](#) compared the expected emission various mechanisms for an analytic AGN wind and found that free-free emission from the outflow dominated inverse-Compton emission at the spatial scales ( $\gtrsim 10 \text{ pc}$ ) and densities ( $n \approx 10 \text{ cm}^{-3}$ ) relevant for the outflows in this work. We note that this calculation yields the broad-band X-ray emission, so care must be taken when comparing to observed data which is often split into soft (0.1 – 2 keV) and hard (2 – 10 keV) bands.<sup>2</sup> We discuss this further in Section 5.4.

<sup>2</sup>These bands are not strictly defined, and may vary based on the instrument and science case being discussed.

### 5.2.3 Synthetic Observations

In addition to the Bremsstrahlung approximation, we also generate synthetic observations of our simulations as if they were observed with the *Chandra* X-ray telescope. We use the PyXSIM<sup>3</sup> Python package (ZuHone & Hallman 2016, based on the earlier code PHOX; Biffi et al. 2012, 2013) to generate synthetic X-ray emission from our simulation output which is then fed to the SOXS<sup>4</sup> instrument simulator (ZuHone et al., 2023) to model the *Chandra* observations.

The workflow for generating these synthetic observations is as follows:

- A source model for the X-ray emission is defined using PyXSIM. We use a thermal model, assumed to be in collisional ionisation equilibrium (CIE). As discussed above, this is likely to be dominated by free-free processes.
- Using this emission model, each AREPO cell in our simulation is assigned photons which are then cosmologically redshifted according to the source distance, creating a *photon list*. At this stage, a large number of photons are generated to be sampled from in the following steps.
- The photons are then projected along the chosen line-of-sight, taking into account any Doppler shifts from gas motion within the simulation, and optionally applying Galactic foreground absorption. This creates an *event list* of photons that have been detected.
- This event list is then passed to SOXS which performs mock observations using the parameters of a given telescope, such as the point spread function (PSF), detector and background noise, bandwidth, and effective area. Functionality is included for a range of instruments for current and future X-ray missions, such as *Lynx*, *Chandra* and *Athena*.

In Section 5.4 we show the resulting mock *Chandra* images of our simulations.

## 5.3 Results


### 5.3.1 Wind-Cloud Interactions Produce X-ray Emission

Figure 5.1 shows the resulting X-ray emission from the fiducial simulations presented in Paper I at  $t=1$  Myr. The top row shows the full galaxy disc edge-on (left panel) and top-down (right panel). These panels show a projection of the integrated X-ray luminosity from Bremsstrahlung emission and a density-weighted mean of the temperature, electron density and wind tracer density, clockwise from top left.<sup>5</sup> We can see that by  $t=1$  Myr the outflow has travelled about 1 kpc within the disc (equatorial outflow) and around 2 kpc outside the disc into the halo (polar outflow). The equatorial outflow has the effect of heating (to  $T \approx 10^{6-7}$  K) and compressing the gas within the disc (top-right subpanel), which increases the free electron density (bottom-right subpanel). This results in bright X-ray emission within the disc. As the gas distribution in the initial disc was clumpy, the X-ray emission is also inhomogeneous; we discuss how the initial clumps lead to bright X-rays in the next section. The polar outflow results in roughly symmetric X-ray bubbles rising in the halo. However, these are less bright than the X-ray emission from the

<sup>3</sup><https://hea-www.cfa.harvard.edu/~jzuhone/pyxsim>

<sup>4</sup><https://hea-www.cfa.harvard.edu/soxs>


<sup>5</sup>Note that in Paper I, we showed midplane slices of the galaxy to more clearly show the effect of the wind interacting with the gas clumps. In this plot, we show projections through the whole disc to more easily compare to the integrated X-ray emission.



**Figure 5.1:** An overview of our simulations, showing the resulting X-ray luminosity from Bremsstrahlung emission at  $t=1$  Myr. The top row shows a projection of the galaxy disc side-on (left) and top-down (right). Across all the panels, the X-ray luminosity has been integrated along the line of sight and the other quantities are density-weighted averages (temperature, electron density and wind tracer density). The bottom panels show a zoom-in of an initially dense gas cloud and the tail formed from stripped and mixed gas behind it. We find that the strongest X-ray emission due to the quasar-driven outflow comes from gas mixing with the wind in the tail behind the dense clump, forming a luminous chimney, on a scale of a few hundred parsecs.

equatorial outflow. We can see that the electron density of these bubbles is only slightly higher than the background, resulting in weak X-ray emission. However, this may be due to our static, isobaric background which is rather a simplistic model of the environment around galaxies. Therefore, in a more realistic galaxy halo, the outflow may be able to sweep up more gas in these polar bubbles, resulting in higher electron density and/or higher temperatures and thus higher X-ray luminosity. Exploration of this is left to future work.

Linking the disc to these halos, we see some bright ‘chimneys’ of X-ray emitting gas. In the lower panels of Figure 5.1, we show a closer view of one of these structures which represent some of the most luminous gas in the X-ray. We show slices of a  $d \approx 100$  pc cloud and its tail, showing from left to right the integrated X-ray luminosity, and density-weighted averages of the temperature, electron density and wind tracer density. The AGN wind is approaching from the top. This cold cloud ( $T \approx 10^4$  K) has survived for  $t=1$  Myr despite strong ram pressure from the wind. Behind the cloud, we can see a long tail of gas which is constrained by two channels either side where the wind is venting through at high velocity. However, despite these vents containing the hottest gas, their low density



**Figure 5.2:** Histogram of X-ray luminosity as a function of wind tracer density ( $\mathcal{P}$ ). Each bin shows the sum of the X-ray emission within it. We split the wind into pure, mixed and shocked ISM. The pink and yellow lines show the results from an initially clumpy setup, with  $\lambda_{\max}=170$  pc and  $\lambda_{\max}=40$  pc respectively. The blue line shows a smooth setup. We can see that the smooth case has a much larger contribution from low- $\mathcal{P}$  values (shocked ISM).

( $n_e \approx 10^{-2} \text{ cm}^{-3}$ ) results in low X-ray emission. We find that the X-ray emission is strongest in region behind the cold clump (the ‘tail’), which has both a moderate density ( $n_e \approx 1 \text{ cm}^{-3}$ ) and temperature ( $T \approx 10^{6-7} \text{ K}$ ).

In the rightmost panel, we show the wind tracer density. This allows us to see the contribution of the injected wind to the X-ray luminosity. We split the wind tracer into three regimes: ‘pure wind’ ( $\mathcal{P} > 0.5$ ) which is dominated by the injected AGN wind; ‘mixed wind’ ( $10^{-3} < \mathcal{P} < 0.5$ ) where the injected wind has thoroughly mixed with the initial medium; and ‘shocked ISM’ ( $\mathcal{P} < 10^{-3}$ ) which are regions which have only very weakly mixed with the wind fluid. We can see that the initial clump has low wind tracer values as the wind has not penetrated the dense gas and the region above the clump is dominated by the freely-expanding wind. We can also see the ‘vents’ either side of the gas tail where the wind is escaping along low-density channels. This region has low X-ray luminosity, suggesting that strong X-ray emission anti-correlates with the locations where the AGN wind is venting through low-density regions. The region behind the cloud, where the brightest X-ray emission is produced, is in the mixed wind phase showing that this is where free-free emission is most efficient.

In the next section we quantify the contribution of each of these wind phases to the total X-ray luminosity of the galaxy.

### 5.3.2 Wind Tracer Phases

In Figure 5.2 we show a histogram of the calculated X-ray as a function of the density of the passively advected wind tracer. Each bin shows the sum of the X-ray emission within that bin. The pink line shows the fiducial simulation with medium-sized initial clumps (fiducial;  $\lambda_{\max}=170$  pc), the orange line shows the results from the setup with small initial clumps ( $\lambda_{\max}=40$  pc), and the blue line shows an initially homogeneous disc. We split the wind tracer density into three regimes, as described in the previous section.

Considering first a wind propagating in a smooth medium, Nims et al. (2015) used an analytic model, based on the work of Faucher-Giguère & Quataert (2012), to estimate the free-free emission, assuming an isotropic wind, and a spherically-symmetric, homogeneous ambient medium. They argued that the shocked ambient medium would contribute more

to the free-free emission than the shocked wind phase analysed in [Faucher-Giguère & Quataert \(2012\)](#). The blue line in Figure 5.2 shows our results when using our smooth disc. Our results concur with [Nims et al. \(2015\)](#) that the emission from the shocked wind (‘pure wind’;  $\mathcal{P} > 0.5$ ) is subdominant compared to the shocked ambient medium (‘shocked ISM’;  $10^{-3} < \mathcal{P} < 0.5$ ). Although the shocked wind is hot ( $T \approx 10^9$  K), its low density results in weak X-ray emission. We find that most of the X-ray emission from the smooth disc is from gas with  $\mathcal{P} \lesssim 10^{-3}$ , representing gas that has barely mixed with the AGN wind.

However, our results differ from the findings of [Nims et al. \(2015\)](#) and our smooth model when we consider a clumpy medium, shown in pink (medium clumps) and orange (small clumps) in Figure 5.2. Firstly, we can see that the emission from the pure wind is even lower than in the smooth case, possibly because the free expansion phase of the wind (see, e.g., [Faucher-Giguère & Quataert, 2012](#); [Costa et al., 2020](#)) has been disrupted by mixing with dense clouds. The emission peaks for intermediate values of  $\mathcal{P}$  – this is ISM gas which has been ablated from the cold clumps and mixed with the injected wind (see bottom panels in Figure 5.1). The emission declines again for low wind tracer values, unlike in the smooth case which shows strong emission in the shocked ISM phase. This shows that when the ISM is arranged in a clumpy structure, it is much more difficult for shocks to cause it to radiate. This confirms what was postulated by [Nims et al. \(2015\)](#), who predicted that an AGN wind would struggle to shock-heat dense clumps to high enough temperatures for luminous free-free emission. However, we find this is partly compensated for by enhanced mixing between the ISM and wind, resulting in more emission for the mixed wind phase in the clumpy case.


Additionally, we find that the size of the initial clumps plays a role in the resulting emission. We find that smaller initial clumps ( $\lambda_{\max}=40$  pc; orange line) have more emission across all values of  $\mathcal{P}$  than the medium clumps ( $\lambda_{\max}=170$  pc; pink line. We didn’t find significant difference between medium and large ( $\lambda_{\max}=330$  pc) clumps so we do not show those here). The enhancement in emission in the mixed wind phase could be due to a greater surface area for the wind to interact with smaller ISM clumps, resulting in more mixing between the hot wind and the cold ISM. Additionally, the slightly higher emission in the shocked ISM regime could be because the more tightly spaced clumps could be more efficient at trapping the expanding outflow, resulting in less wind venting and more shock-heating. This is supported by the result found in Paper I that small clumps resulted in a higher mass outflow rate in the hot gas phase.

The exact boundary value between mixed wind and shocked ISM phase is somewhat of an arbitrary choice. The value we have chosen ( $\mathcal{P}_{\text{mix}} \approx 10^{-3}$ ) represents a mass fraction of  $< 0.1\%$  of the cell being composed of mass initially injected by the wind for ‘shocked’ or unmixed ISM. From Figure 5.2, we can see that the clumpy case has far less X-ray emission at low  $\mathcal{P}$  than the smooth case, showing the emission from clumpy discs is dominated by mixed gas. This broad conclusion is insensitive to the exact choice for  $\mathcal{P}_{\text{mix}}$ .

### 5.3.3 Total X-ray Contribution by Wind Phase

In Figure 5.3 we show the total X-ray luminosity, summed across all cells in the box at  $t=1$  Myr. We show the cumulative emission as a function of radial velocity, split by wind phase (see Section 5.3.2). The left panel shows our fiducial run, with medium clump sizes ( $\lambda_{\max}=170$  pc), the middle panel shows the simulation with small clump sizes ( $\lambda_{\max}=40$  pc), and the right panel shows an initially inhomogeneous setup.

For the medium clumps (left panel), we can see that the highest velocity gas is all contained in the pure wind phase (blue), starting at the injection velocity of  $v_r=10^4$  km s $^{-1}$ . However, this makes a small overall contribution to the overall X-ray luminosity ( $L_X \approx 3 \times 10^{37}$  erg s $^{-1}$ ) and there is no gas moving slower than  $v_r \lesssim 1000$  km s $^{-1}$  in this phase emitting any X-rays. At a radial velocity of  $v_r \lesssim 3000$  km s $^{-1}$ , the mixed wind phase starts




**Figure 5.3:** Cumulative total X-ray luminosity as a function of outflow velocity, split into the wind phases as presented in Figure 5.2 with the mixed-ISM boundary set at  $\mathcal{P}_{\text{mix}} = 10^{-3}$ . The horizontal dashed line shows the emission over the whole velocity space (i.e., including the static background). We show (from left to right) the results from our medium clumps (fiducial), small clumps and smooth simulations. In the two clumpy cases, the pure wind dominates at the fastest velocities  $v_r \gtrsim 3000 \text{ km s}^{-1}$  (i.e., in the centre near injection and high-velocity vents) but this is negligible overall. Gas that has significantly mixed with the wind dominates from  $v_r \approx 100 - 3000 \text{ km s}^{-1}$ . Overall, the shocked ISM and the mixed wind have roughly equal contributions to the total X-ray luminosity for the clumpy case. However, in the smooth case the emission from unmixed gas ( $\mathcal{P}_{\text{mix}} \lesssim 10^{-3}$ ) dominates the other sources.

dominating over the pure wind before flattening off at  $v_r \approx 100 \text{ km s}^{-1}$ . As mentioned, this phase represents outflowing gas that has mixed significantly with the wind. At lower velocities, the shocked ISM phase has similar X-ray emission to the mixed wind. The shocked ISM phase is made up of gas that has been driven in front of the wind shock without mixing with it and cold clumps that have been heated but not destroyed by the wind. This phase is mostly moving at  $v_r \approx 100 - 500 \text{ km s}^{-1}$ . The inset shows that the total X-ray emission from the shocked ISM and mixed wind phases is similar, although the mixed phase has a slightly higher total of  $L_X \approx 9 \times 10^{40} \text{ erg s}^{-1}$  compared to  $L_X \approx 8 \times 10^{40} \text{ erg s}^{-1}$  in the shocked phase. This gives a total emission in the fiducial simulation of  $L_X = 1.8 \times 10^{41} \text{ erg s}^{-1}$ .

In the case with small initial clumps (middle panel of Figure 5.3), we can see a similar distribution to the fiducial run. However, the total luminosity is higher ( $L_X = 4.0 \times 10^{41} \text{ erg s}^{-1}$ ) which is driven by increased emission in both the shocked and mixed phases. The increase in the mixed phase shows that the wind is able to interact with a greater amount of cold gas, possibly due to the higher surface area of the smaller clumps. In Paper I, we showed that the velocity distribution of the tightly clumped case started to trend towards that of the smooth case, suggesting the outflow was being more efficiently trapped by the smaller clumps. This could also be the reason for the elevated contribution from the shocked phase.

In the rightmost panel we plot the results for the smooth disc. We can see that the pure wind contribution is higher than in the clumpy case, however, it is still subdominant to the mixed and shocked phases, as also predicted by earlier works (e.g., [Faucher-Giguère & Quataert, 2012](#); [Nims et al., 2015](#)). The emission from the shocked phase rises more steeply than in the clumpy case – i.e., there is a narrower velocity range where most of the emission is occurring. The total emission starts to flatten out at  $v_r \approx 300 \text{ km s}^{-1}$  compared to  $v_r \approx 100 \text{ km s}^{-1}$  in the clumpy case. This is similar to our finding in Paper I, where, in the smooth case, the bulk of the outflowing material had a characteristic velocity of  $v_r \approx 400 \text{ km s}^{-1}$ , whereas in the clumpy case, there was a wider range of velocities ( $v_r \approx 10 - 300 \text{ km s}^{-1}$ ). The total luminosity in the smooth run is  $L_X = 3.8 \times 10^{41} \text{ erg s}^{-1}$ ;



**Figure 5.4:** The radial evolution of the X-ray producing outflow. We show the X-ray luminosity in increasing radial shells, normalised by the volume of the shell. The peak of the outflow decreases and broadens with time, but still shows strong emission at radii  $R < 2$  kpc. There is some larger-scale emission in the halo ( $R > 2$  kpc), but this is much fainter than the outflow within the disc.


around double the value of the medium clumps, and just less than the small clumps. This demonstrates that the enhanced X-ray emission from wind-ISM mixing somewhat compensates the lack of emission from the shocked ISM in the clumpy case (see Figure 5.2).

The grey dashed line in all three panels of Figure 5.3 shows the X-ray emission summed over the whole simulation and the grey solid histogram shows the sum of the three outflow phases cumulative emission by radial velocity. In the fiducial case, the sum of the three outflow phases (solid grey line) is around  $2 \times 10^{40}$  erg s $^{-1}$  lower than the overall total (dashed grey line). This difference comes from gas with  $v_r \lesssim 5$  km s $^{-1}$  associated to the static, hot background. To confirm this, we calculated the X-ray emission from our no-AGN control run which summed to  $L_{X,\text{bkg}} = 2.5 \times 10^{40}$  erg s $^{-1}$ . For the fiducial run of  $L_{\text{AGN}} = 10^{45}$  erg s $^{-1}$ , this is a negligible contribution, but it becomes more dominant for lower-luminosity runs. For this reason, in Sections 5.3.5 & 5.3.6, we subtract this background value from the total emission.

### 5.3.4 Radial Evolution

Figure 5.4 shows the radial evolution of the X-ray producing gas for a range of times up to  $t=3$  Myr. The galaxy is binned in spherical shells of radius  $R$  and the luminosity is divided by the volume of each shell. Darker lines show later times. The X-ray emission starts strongly peaked, before widening and reducing in peak flux. The peak of the emission is moving at  $v \approx 250$  km s $^{-1}$ , demonstrating it is dominated by the equatorial outflow. By  $t=3$  Myr, the emission has broadened significantly across a width of  $R \approx 1.5$  kpc. As shown in Figure 5.1, in our setup, most emission is formed in the disc, rather than in the halo bubble. This is because the overall emission is dominated by ISM gas that has strongly mixed with the wind, rather than the wind material itself (Figures 5.2 & 5.3). Once the outflow breaks out of the disc in the polar direction, there is less material to mix with, due to the halo’s low density. We can see in Figure 5.4 that at large radii ( $R \gtrsim 2$  kpc) the emission is significantly weaker than within the disc, although this may be driven by our simplistic halo density model, as discussed in Section 5.3.1.

Our prediction for X-ray emission on extended ( $D \approx 3$  kpc) scales within the disc suggests that it could be distinguished from the point-source emission from the AGN itself. We discuss the detectability of this emission further in Section 5.4.



**Figure 5.5:** The time evolution of the total integrated X-ray luminosity (background subtracted). The pink lines show the medium clumped disc at a range of AGN luminosities (as shown by variable line-style). The dark blue lines show an initially smooth disc at  $L_{\text{AGN}}=10^{45}$  erg s $^{-1}$ , and the light blue/orange lines show initially large/small clump sizes.

### 5.3.5 Time Evolution

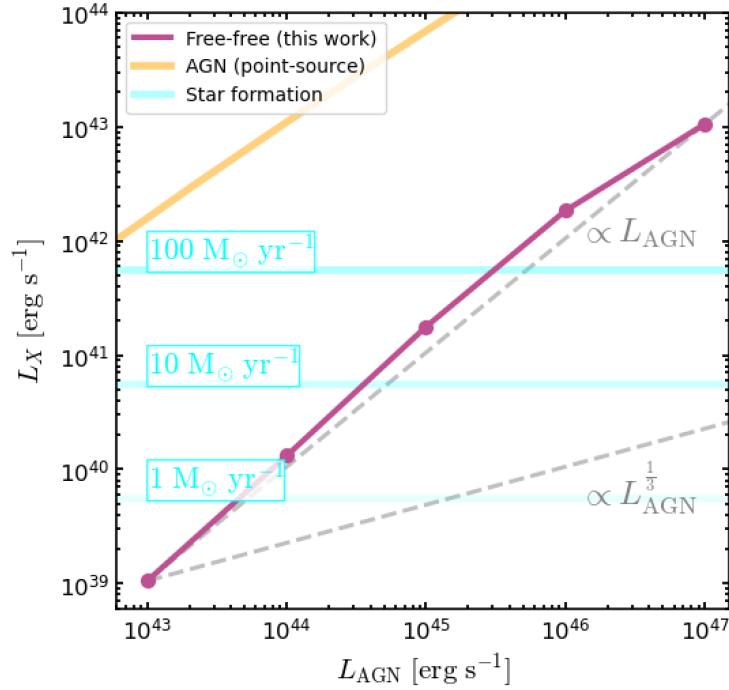
Figure 5.5 shows the time evolution for the total X-ray luminosity from our simulations. We show a range of AGN luminosities ( $L_{\text{AGN}}=10^{43-47}$  erg s $^{-1}$ ; dashed/dotted lines) in discs with our fiducial clumpiness ( $\lambda_{\text{max}}=170$  pc; ‘medium’ clumps). For an AGN luminosity of  $L_{\text{AGN}}=10^{45}$  erg s $^{-1}$ , we also show results for variations in clumpiness, including large clumps ( $\lambda_{\text{max}}=330$  pc; light blue), small clumps ( $\lambda_{\text{max}}=40$  pc; orange), and the smooth setup (dark blue). We subtract the background emission of  $L_{\text{X,bkg}} = 2.5 \times 10^{40}$  erg s $^{-1}$  from these values (see Section 5.3.3).

In our fiducial run (solid pink line), the luminosity rises to a peak of  $L_{\text{X}} \approx 2.5 \times 10^{41}$  erg s $^{-1}$  at  $t \approx 0.3$  Myr before declining slowly to  $L_{\text{X}} \approx 6 \times 10^{40}$  erg s $^{-1}$  by  $t \approx 4$  Myr. The large clump case (light blue) is very similar to the fiducial at  $t \lesssim 1$  Myr but then declines slightly faster to  $L_{\text{X}} \approx 5 \times 10^{40}$  erg s $^{-1}$ . The small clumps (orange) show a higher and slightly later peak in emission of  $L_{\text{X}} \approx 5 \times 10^{41}$  erg s $^{-1}$  and declines less rapidly than the medium and large clumps. This could be due to increased interactions between the wind and the clouds due to the higher overall surface area of the small clumps. The smooth setup shows a rapid increase in luminosity to a peak of  $L_{\text{X}} \approx 3 \times 10^{42}$  erg s $^{-1}$  at  $t \approx 0.8$  Myr before rapidly declining. This corresponds to the ‘post-shock-cooling’ time discussed in Paper I. At this time, the amount of cold gas in the outflow rapidly increases due to efficient cooling, but Figure 5.5 shows that this results in an equally sudden drop in X-ray emission due to much less hot gas in the outflow. By  $t \approx 1.5$  Myr, the smooth case has the lowest emission of any of the  $L_{\text{AGN}}=10^{45}$  erg s $^{-1}$  simulations.

In Figure 5.5 we also show the time evolution for a range of AGN luminosities.<sup>6</sup> The highest luminosity AGN show the brightest free-free emission, with the X-ray luminosities for  $L_{\text{AGN}} \geq 10^{44}$  erg s $^{-1}$  all showing a similar shape; rising to a peak at  $t \approx 0.3$  Myr before declining slowly, although the peak in the  $L_{\text{AGN}}=10^{44}$  erg s $^{-1}$  track is quite small. However,

<sup>6</sup>The highest AGN luminosity runs ( $L_{\text{AGN}}=10^{46-47}$  erg s $^{-1}$ ) are evolved to shorter times due to computational and boxsize constraints.





**Figure 5.6:** Correlations between AGN luminosity and the resulting Bremsstrahlung emission. In pink, we show our results for the fiducial clumpy simulation at  $t=1$  Myr. We compare this to two theoretical predictions for the  $L_X - L_{\text{AGN}}$  correlation (Nims et al., 2015). The results are compared to two other sources of X-rays: emission from the point-source AGN (shown in orange; Marconi et al. 2004), and from star formation (shown in blue; Pereira-Santaella et al. 2011). Both of these components are integrated across the range 0.5 – 10 keV. Although the point-source AGN dominates our predicted free-free emission, it may be detectable against the diffuse star formation component, if sufficiently resolved.

the lowest AGN luminosity ( $L_{\text{AGN}}=10^{43}$  erg s $^{-1}$ ) shows a different trend. The emission peaks at the start before rapidly declining to  $L_X \approx 10^{39}$  erg s $^{-1}$  by  $t \approx 0.5$  Myr. This could suggest the AGN is not powerful enough for the outflow to impact and mix with a significant amount of ISM gas.

We can see that the simulations with different  $L_{\text{AGN}}$  have roughly equally-space increases in X-ray emission. In the next section, we will investigate the correlation between the AGN luminosity and resulting free-free emission.

### 5.3.6 Scaling Relations between AGN and X-ray Luminosities

In Figure 5.6 we show the predicted scaling relations between the AGN luminosity and various sources of X-rays within a galaxy. The results for this work, where we estimated the free-free emission from hot gas and found it was primarily driven by wind-ISM mixing, are shown in pink. We show our results at a time of  $t=1$  Myr which is after the peak in emission shown in Figure 5.5. We use our fiducial setup, with medium initial clumps and subtract the background emission (see Section 5.3.3). We find a positive correlation between the AGN and free-free luminosities across the whole parameter range of  $L_{\text{AGN}}=10^{43-47}$  erg s $^{-1}$ .

For comparison, in grey we plot theoretical predictions for scalings of  $L_X \propto L_{\text{AGN}}$  and  $L_X \propto L_{\text{AGN}}^{1/3}$ . The normalisation of these tracks has been arbitrarily set to the emission from the lowest luminosity simulation. Nims et al. (2015) predict a scaling between the AGN luminosity and free-free emission with a power of 1/3 in the case of inefficient cooling of the shocked ISM and a power of 1 if the ISM can cool rapidly. Our results match the rapid cooling case, with a gradient of just over unity from  $L_{\text{AGN}}=10^{43-46}$  erg s $^{-1}$ , and perhaps a slight tail off in the brightest AGN. This tail-off hints at the result in Ward et al.

(2024) where we found that extremely bright AGN start to destroy the clouds before they are entrained in the outflow. The tail-off in this case is more subtle than that found for the cold mass outflow rate, but if the bright AGN is destroying the clouds too efficiently then the cloud crushing time may drop below the cooling time, reducing the free-free emission.

The scaling of  $\approx 1$  we find can also be justified as follows. In the AGN wind model used in this work (Costa et al., 2020), the density of the wind is proportional to the AGN luminosity:

$$n_w \propto L_{\text{AGN}} \quad (5.2)$$

As the wind mixes with the ISM, the density of this mixed material is:

$$n_{\text{mix}} \approx \sqrt{n_w n_{\text{ISM}}} \quad (5.3)$$

From the Bremsstrahlung approximation given in Equation 5.1, the free-free X-ray emission is thus given by:


$$L_X \propto n_{\text{mix}}^2 \propto n_w n_{\text{ISM}} \propto L_{\text{AGN}} n_{\text{ISM}} \quad (5.4)$$

which yields the linear relationship found in Figure 5.6. It also suggests that the X-ray luminosity should be linearly dependent on the initial ISM density, a prediction that will be tested in future work.

Now we have explored the free-free emission caused by the AGN wind, the next question is: is this observable? The two other major sources of X-rays in galaxies are star formation and the AGN itself, through X-ray emission that has been up-scattered from the accretion disc. In Figure 5.6 we plot tracks showing the rough contribution of each of these to the X-ray emission.

Star formation produces X-rays mostly through the creation of X-ray binaries (XRBs). We plot the expected 0.5 – 10 keV X-ray contribution of these sources to the overall galaxy X-ray budget in blue by summing the soft and hard contributions found by Pereira-Santaella et al. (2011) for a sample of local galaxies, showing increasing star formation rates (SFR) in darker lines. At our fiducial AGN luminosity of  $L_{\text{AGN}}=10^{45}$  erg s<sup>-1</sup>, the X-ray luminosity associated with the outflow of  $L_X = 2 \times 10^{41}$  erg s<sup>-1</sup> is equivalent to a star formation rate of around  $\text{SFR} \approx 35M_{\odot} \text{ yr}^{-1}$  which is a high star formation rate for local galaxies. Additionally, the free-free emission is likely to be centrally concentrated than the star formation emission, making it a dominant contributor to the total X-ray emission for radii of a few kiloparsecs. For an AGN luminosity of  $L_{\text{AGN}}=10^{46}$  erg s<sup>-1</sup>, the comparable star formation is  $\text{SFR} \approx 350M_{\odot} \text{ yr}^{-1}$  which exceeds all but the most luminous quasar-hosting galaxies (e.g., Stanley et al., 2017; Zhuang et al., 2021).

The orange line in Figure 5.6 shows the expected X-ray emission from the AGN itself, based on summing the soft and hard X-ray-bolometric corrections of Marconi et al. (2004). AGN are luminous in the X-ray, so unsurprisingly this emission dominates over both the star formation and free-free emission, being around two orders of magnitude higher than the free-free track in pink. However, these X-rays are produced by the AGN accretion disc which is unresolvable, resulting in point-source emission. Therefore, if the spatial resolution of the observation is sufficient to resolve the free-free emission, it should be able to be disentangled from the point-source accretion disc. In Section 5.4, we discuss the maximum distance of a galaxy for the free-free region to be spatially resolved.



**Figure 5.7:** Mock observations of our predicted free-free emission using *Chandra* for our galaxy at  $t=3$  Myr with an exposure time of  $t_{\text{exp}} = 100$  ks. The top panel is face-on to the disc, and the bottom panel is side-on. The white outlines show the dimensions of the initial disc. At the angular diameter distance shown ( $d_A = 50$  Mpc;  $z \approx 0.11$ ), 1 kpc is approximately  $4''$  (scale bar in the lower right). We show the PSF of the *Chandra-ACIS* in the lower left (FWHM =  $0.5''$ ).

## 5.4 Discussion: Observations with Chandra

In this section we will discuss whether the prediction emission from wind-ISM mixing could be observable, using the *Chandra* X-ray observatory as a baseline instrument. *Chandra* is one of the foremost X-ray telescopes currently in operation, boasting a high spatial resolution of  $0.5''$ , which has allowed detailed mapping of AGN outflows in several nearby galaxies (e.g., Croston et al., 2008; Greene et al., 2014; Di Gesu et al., 2017; Lansbury et al., 2018). Using the method outlined in Section 5.2.3, based on the PyXSIM and SOXS packages, we created synthetic *Chandra* observations of our simulations, which are shown in Figure 5.7.

We place our galaxy at a simulation time of  $t=3$  Myr at an angular diameter distance of  $d_A = 50$  Mpc (representing a fairly local galaxy) and show the results for a wide X-ray band of  $0.2 - 10$  keV with an exposure time of  $t_{\text{exp}} = 100$  ks. At this distance, 1 kpc spans  $4''$  (shown by the scale bar in the lower right). We use the in-built SOXS model for the ACIS (Advanced CCD Imaging Spectrometer) instrument on *Chandra*, and show the PSF of this instrument in the lower left (FWHM =  $0.5''$ ).

We stress that this forward modelling is highly idealised – we only consider the free-free emission from our gas (assuming collisional ionisation equilibrium), and neglect all other X-ray producing processes in the galaxy. For example, point-source nuclear emission will dominate in the centre and may also be scattered more widely; star formation will be diffuse through the galaxy, probably peaking in the centre; and photoionisation from the AGN’s radiation field will also cause X-ray emission. We have also neglected any obscuration by intervening gas, including galactic absorption. Nevertheless, we can still use this result to discuss the potential for the free-free emission predicted by our model to be observed.

Firstly, we can see the X-ray emission is concentrated within the disc of the galaxy (shown in white outlines;  $D = 4$  kpc). As discussed in Section 5.3.3, in a clumpy medium, the Bremsstrahlung emission is dominated by mixed material being stripped from cold clouds. This makes the equatorial outflow X-ray bright, and shows us that the total free-

free emission is a good tracer for the amount of ISM the wind has interacted with. The side-on view (lower panel of Figure 5.7) shows some chimneys extending out the top and bottom of the disc (Ponti et al., 2019, 2021). However, we don't see large X-ray bubbles in the halo of the galaxy, which are seen in observations (Greene et al., 2014; Lansbury et al., 2018), some cosmological simulations (Pillepich et al., 2021), and in the eROSITA bubbles of the Milky Way (Predehl et al., 2020). We can see bubbles in Figure 5.1, but as we mentioned in Section 5.3.1, the low density and temperature contrast between these bubbles and the background results in weak free-free emission which we could not detect in our forward modelling, despite our relatively deep exposure time. Our inability to detect these features is probably due to our simplistic halo modelling – we use a static, isobaric background with a constant density ( $n_{\text{bkg}} = 10^{-2} \text{ cm}^{-3}$ ) which does not reflect the complex density profile and multiphase structure seen in the CGM (Tumlinson et al., 2017). This results in the outflow interacting with less mass than is physical once it leaves the disc, and thus producing only weak X-ray emission. Alternatively, some previous studies that have observed large X-ray bubbles in the CGM (e.g., Lansbury et al., 2018) have also found spatially coincident radio emission consistent with a jet. This could suggest that additional energy injection is required to create a bright X-ray bubble. However, another system with a bright bubble (Greene et al., 2014) found that the outflow could be described with a wind alone. Future simulations with a more realistic treatment of the CGM are vital to explore whether energy injection from winds are sufficient to inflate luminous X-ray bubbles. However, we argue that the emission seen within the disc due to the interaction between the clouds and the wind is a genuine prediction of our model. We will now discuss whether this emission could be observed above other processes in the galaxy, using *Chandra-ACIS* as a baseline instrument.

As shown in Figure 5.6, the dominant source of X-rays is expected to be the nuclear emission, caused by photons being up-scattered from the corona and reflected by the accretion disc. However, this emission is spatially unresolved and thus acts as a point-source. It could therefore be separated from the free-free emission if this component is well-resolved. At  $t=3$  Myr, the diameter of the free-free region is  $D \approx 3$  kpc (Figure 5.4). We assume the structure to be sufficiently resolved at an angular size of three resolution elements, which for *Chandra's* FWHM = 0.5" PSF is  $\theta = 1.5''$ . Therefore, for the  $D = 3$  kpc free-free emission, this suggests that we could resolve this region at angular diameter distances of  $d_A \lesssim 410$  Mpc, corresponding to a redshift of  $z \lesssim 0.11$ . However, we note that scattering could blur the nuclear emission beyond just a point-source, so this represents an optimistic scenario.

Diffuse emission from star formation (i.e., X-ray binaries) will also contribute to the total X-rays observed (e.g., Lehmer et al., 2010). However, as we have shown in Figure 5.6, the free-free emission should dominate over the XRB contribution in luminous quasars ( $L_{\text{AGN}} \gtrsim 10^{45} \text{ erg s}^{-1}$ ), except for in the most extreme starbursts. Additionally, supernova-driven outflows could drive shocks into the ISM, resulting in similar free-free emission to our results for an AGN wind. Wang et al. (2014) find diffuse hard X-rays in the central 5 kpc of a nearby galaxy merger. They attribute only 5% of this emission to XRBs, and find the rest is produced by shock-heating from SNe.

A more challenging component to distinguish this emission from photoionisation caused by the radiation field of the AGN (e.g., Wang et al., 2024b). In Di Gesu et al. (2017), extra-nuclear emission was found on kiloparsec scales. The spectra was fit both with a collisional excitation model (caused by an outflowing wind) and a photoionisation model. They found that either scenario was in good agreement with their data, but were unable to differentiate between the two. A similar result was found in Wang et al. (2010), with neither scenario being strongly preferred.

To break this degeneracy, more detailed theoretical expectations are required. This would require radiative transfer modelling to track the spectra of the photoionised plasma, alongside the collisionally ionised gas from the outflow. Using a similar idealised setup to ours, Meenakshi et al. (2022a,b) post-processed their jet-ISM simulations. They found that the fraction of dense ( $n > 100 \text{ cm}^{-3}$ ) gas affected by photoionisation was much lower (only  $\approx 5\%$ ) than the gas that was collisionally ionised by their jet-induced shock. This suggests that the extended X-ray emission seen in the aforementioned observational works is likely to be mostly caused by shocks rather than radiation. However, more theoretical work is needed to explore the parameter space and verify if this conclusion holds for wind-driven outflows.

## 5.5 Conclusions & Outlook

In this study, we used the simulations of an AGN wind in a clumpy ISM from Ward et al. (2024) and investigated the resulting X-ray emission, using the Bremsstrahlung approximation and synthetic imaging techniques. Our main results are as follows:

- Wind-ISM interactions produce luminous free-free emission due to mixing between the AGN wind and ISM gas. This X-ray emission is strongest in the ‘tails’ behind dense clumps, where the stripped gas from the clouds is efficiently mixed with the wind. This is in contrast to analytic results for a wind in a homogeneous medium which find that the emission is dominated by shocked ISM gas.
- We predict a strong scaling between the AGN luminosity and the resulting X-ray emission, with  $L_X \propto L_{\text{AGN}}$ . This demonstrates our outflow is in the rapid cooling limit where the shocked ambient medium cools on the outflow timescale (Nims et al., 2015).
- This emission should be detectable in nearby galaxies as spatially extended X-rays from the central 1-4 kpc. The free-free emission can be detected over the diffuse contribution from star formation, but is dominated by the central point source X-ray emission from the AGN itself. This emission could be resolved up to a distance of  $d_A \approx 400 \text{ Mpc}$  ( $z \approx 0.1$ ) for an instrument with a resolution of  $0.5''$ . However, it may be challenging to separate this emission from that caused by photoionisation.

This study represents a first look at the X-ray emission expected from wind-cloud interactions, and a proof-of-concept for being able to observe this emission with *Chandra*. However, various improvements could be made to the model to broaden the scope of our results.

The main limitation is the simplified halo model used which limits our ability to predict the emission from X-ray bubbles rising into the CGM. In contrast to several observational studies (e.g., Croston et al., 2008; Greene et al., 2014; Di Gesu et al., 2017) we struggled to find any detectable emission on  $\approx 10 \text{ kpc}$  scales. An improved CGM model could result in higher densities in the outflowing bubble, allowing us to predict emission on these scales. It is also possible that such large X-ray bubbles require an additional driving mechanism to reach the densities required for such emission; for example, the observed X-ray bubble in Lansbury et al. (2018) was found to be aligned along a kiloparsec-scale jet (Harrison et al., 2015). Furthermore, our simulations only include primordial cooling. Including metal-line cooling may boost the X-ray luminosity.

An additional method to separate the free-free emission from the star formation and the AGN accretion disc could be to investigate the expected hardness of the emitted X-rays. Star formation is expected to dominate at softer ( $\lesssim 2 \text{ keV}$ ) X-rays, and nuclear

emission at harder ( $\gtrsim 2$  keV) X-rays. The Bremsstrahlung approximation mostly used in this study is bolometric, however, investigating the thermal collisional ionisation model in PyXSIM further could allow us to predict the spectra of our X-rays. This could provide an additional key test to differentiate the emission from both the star formation and accretion disc, alongside their spatial extents. Additionally, radiative transfer modelling (Meenakshi et al., 2022a,b) could be performed to compare the contribution of collisionally ionised gas caused by the outflow and photoionised gas caused by the AGN radiation field.

To investigate the predictions presented in this study, we could use an observational sample such as the Quasar Feedback Survey (QFeedS; Jarvis et al. 2021). QFeedS targets the multiphase outflows of a sample of  $z \approx 0.05 - 0.2$  luminous quasars ( $L_{\text{AGN}} \gtrsim 10^{45}$  erg s $^{-1}$ ) with typical star formation of  $\text{SFR} \approx 8 - 80 M_{\odot} \text{ yr}^{-1}$  (Jarvis et al., 2020). Our findings in this study suggest that these systems would be an ideal test-bed for studying the emission predicted from our model.

Finally, we have demonstrated that *Chandra* is a critical tool to be able to detect this emission thanks to its exceptional spatial resolution and sensitivity. However, the future of this telescope is currently uncertain, with potential de-funding of mission operations from October 2024<sup>7</sup>. The next generation of X-ray observatories such as *Athena*<sup>8</sup> from ESA and *Lynx*<sup>9</sup> from NASA will not be launched until the mid-2030s. Therefore, it is critical to keep emphasising the unique science that can be done with *Chandra* to justify its continued funding.

---

<sup>7</sup>[www.savechandra.org](http://www.savechandra.org)

<sup>8</sup>[www.the-athena-x-ray-observatory.eu/en](http://www.the-athena-x-ray-observatory.eu/en)

<sup>9</sup>[www.lynxobservatory.com](http://www.lynxobservatory.com)

# 6

## Summary and Outlook

*But even after everything we've seen,  
we've barely caught a glimpse of what it means.  
In the architecture of the soul  
the universe began with our eyes closed.*

*—Overture from *Darkness, Sleeping at Last**



## 6.1 Conclusions

The last two decades have seen a growing understanding of the crucial role played by supermassive black holes in the evolution of galaxies across cosmic time. As introduced in Chapter 1, negative feedback from active galactic nuclei has been invoked to solve several outstanding problems, such as the correlation between black hole mass and bulge velocity dispersion, the ‘cooling flow’ problem in galaxy clusters, and the quenching of massive galaxies.

However there are still many mysteries yet to be uncovered. For example, why don’t we see evidence for AGN feedback on population level? How do AGN winds couple to the ISM differently in inhomogeneous environments and what are the limitations of assuming a smooth ambient medium? And what are the observational signatures of the interaction between winds and ISM clouds? This thesis has attempted to answer some of these questions by seeking to reconcile results from numerical studies of feedback with techniques used by observers to try and find observational evidence of feedback. In particular, our main conclusions are:

### 6.1.1 AGN preferentially reside in gas-rich galaxies

Various observational studies have looked at global properties across galaxy samples to find population-level evidence that AGN negatively affect their host galaxies. However, such studies have not found clear, ‘smoking-gun’ evidence of AGN suppressing the star formation or molecular gas content of their host galaxies, instead finding that these properties show flat or positive correlations with AGN luminosity, and that AGN-selected and non-AGN samples have similar SFR and gas fractions.

To investigate this apparent contradiction, in **Chapter 3** we analysed three contemporary cosmological simulations: IllustrisTNG (Marinacci et al., 2018; Naiman et al., 2018; Nelson et al., 2018; Pillepich et al., 2018; Springel et al., 2018), EAGLE (Crain et al., 2015; Schaye et al., 2015), and SIMBA (Davé et al., 2019), using post-processed models to estimate the molecular gas masses. All three of these simulations rely on AGN feedback to reproduce key observations (see Chapter 2), although the subgrid implementation of feedback in each differs significantly. We applied similar tests to those used by observers to assess the gas fractions and star formation rates of AGN, comparing to observational datasets both in the local universe ( $z = 0$ ) and at Cosmic Noon ( $z = 2$ ).

We first investigated correlations between the AGN luminosity and the specific star formation rate (sSFR) or molecular gas fraction ( $f_{\text{H}_2}$ ) for star-forming galaxies. We found no significant negative correlations with AGN luminosity, instead finding flat or positive trends. We showed that, in the simulations, AGN are preferentially located in gas-rich and star-forming galaxies. This conclusion holds for both a luminosity- and Eddington ratio-based AGN selection. Finally, we constructed a mass-matched comparison sample and showed that AGN always have a lower fraction of gas-depletion than the control sample. These results all concur, at least qualitatively, with the observational samples used for comparison, demonstrating that finding AGN in gas-rich and star-forming galaxies is not in tension with galaxy evolution models that strongly rely on AGN feedback to quench galaxies.

Despite all three cosmological simulations agreeing with these conclusions, we did find variations between them, driven by their different approaches to modelling AGN activity. For example, in TNG, there is a critical black hole mass ( $M_{\text{BH}} \approx 10^{8.2} M_{\odot}$ ) above which the galaxy quenches rapidly (see also Weinberger et al., 2018; Terrazas et al., 2020). This is due to a sudden switch to the more efficient kinetic mode of feedback. In SIMBA, we see a small population of galaxies with high luminosities but also in the ‘jet’ mode of feedback which is efficiently quenching the galaxy. It is unclear which of these differences are due to

physical reasons that may be observed in real galaxies, or simply to the choice of subgrid modelling or calibration. More dialogue is needed between simulators and observers to establish testable predictions from the simulations that be investigated observationally.

### 6.1.2 ISM structure is key to understanding multiphase outflows

We showed in Chapter 3 that it is challenging to find obvious fingerprints of AGN feedback on population-level studies of global galaxy properties. However, another method of assessing the impact of feedback is to investigate how AGN are able to couple their energy to the interstellar medium by studying galaxies containing outflowing gas. There are a range of observational studies across a variety of wavelengths that find such outflows, demonstrating their multiphase nature. However, there is not a clear consensus on how these outflows are affecting the gas within the galaxy, with observational studies finding a range of momentum boosts and kinetic energy coupling efficiencies. Part of the problem is the lack of theoretical expectations for how an AGN outflow interacts with a realistic multiphase ISM. This is because achieving the spatial resolution required to model the ISM structure is computationally challenging.

In **Chapter 4**, we presented a suite of simulations to investigate how an AGN-driven wind couples to a clumpy ISM structure. Using the AREPO code, we initialised an idealised galaxy disc with a manually-set ISM structure made of fractally-distributed clumps of cold gas, surrounded by a tenuous hot phase. At the centre we placed a physically-motivated model for an AGN wind (BOLA; [Costa et al. 2020](#)) which is able to resolve the shock structure of a wind-driven outflow (e.g., [King, 2003](#); [Faucher-Giguère & Quataert, 2012](#)).

We investigated the multiscale structure of the resulting outflow, which contained kiloparsec-scale bubbles expanding into the halo; high-velocity ‘vents’ where the hot wind could break out of underdense regions; and entrained cold ( $T \lesssim 10^4$  K) clouds that can survive on  $\gtrsim 5$  Myr timescales thanks to efficient cooling on their surfaces. We split the outflow into ‘hot’ ( $T > 10^{4.5}$  K) and ‘cold’ ( $T < 10^{4.5}$  K) phases, finding that the cold outflow carries most of the mass, but that the hot outflow is more energetic, due to its higher velocity. We compared our results to an outflow propagating in an homogeneous medium. This resulted in an outflowing shell of cooling gas, with a clear characteristic velocity and radius. This is in contrast to the range of velocities and radii found in a clumpy medium where the wind is able to vent out of low density regions.

A key aim of this study was to be in dialogue with observational works who analyse the properties of outflows observed mostly in CO (molecular) and [O III] (warm ionised) gas. We discussed how the assumption that the outflow is a thin spherical shell is invalid in the case of a clumpy ISM and demonstrated that the minimum outflow velocity the observational technique is sensitive to is a dominant source of uncertainty, especially in the cold phase. We found that, despite our outflow being energy-conserving overall, values of  $\dot{p}/(L/c) < 1$  (i.e., normally interpreted as momentum-driven) could still be measured if only a single gas phase was measured or the full radial extent of the outflow was not captured. Finally, we investigated scaling relations between the outflow properties and AGN luminosity, as have been investigated in e.g., [Fiore et al. \(2017\)](#); [Musiimenta et al. \(2023\)](#); [Ramos Almeida et al. \(2022\)](#). We found a positive correlation between the mass outflow rate and AGN luminosity, with the cold phase showing a turnover in the highest  $L_{\text{bol}}$  sources (also seen observationally; e.g., [Bischetti et al. 2019](#)). This could suggest that there is a critical AGN luminosity above which the wind destroys the cold clouds before they can become entrained, although further work is needed looking specifically at the cold cloud formation and survival in these simulations.

### 6.1.3 Wind-cloud mixing produces luminous X-ray emission

We showed in Chapter 4 that the hot phase is expected to dominate the energetics of AGN outflows. This hot ionised gas is difficult to observe, due to its long radiative cooling time and low density. However, our simulations showed that mixing between ISM clouds and the AGN wind can enhance cooling at the interface between these two phases, potentially creating observable X-ray emission. In **Chapter 5**, we used our suite of simulations presented in Chapter 4 and estimated the Bremsstrahlung emission from the hot gas, finding bright ‘chimneys’ in regions behind initial gas clumps where the ablated gas could strongly mix with the wind material. This created a luminous equatorial outflow that propagated through the galaxy disc. By using the density of the passively advected wind tracer as a measure for how strongly mixed the X-ray emitting gas was with the AGN wind, we found that X-ray luminosity was dominated by strongly mixed gas. This is in contrast to an initially smooth ambient medium where the free-free emission mostly occurs in shocked ISM driven in front of the outflow. We investigated the scaling relation between AGN luminosity and predicted free-free emission, finding a correlation of  $L_{\text{AGN}} \propto L_X$  which matches analytic expectations for an efficiently cooling outflow (Nims et al., 2015; Costa et al., 2020). We discussed whether this emission could be observable against the bright, point-source AGN itself and the diffuse emission star formation. Although the X-rays from the AGN itself are much brighter than those from the mixing hot gas, they are concentrated as a central point source, unlike the extended ( $D \approx 3$  kpc) free-free emission region we predicted. Therefore, a telescope such as *Chandra* with 0.5” resolution could resolve this emission up to distances of  $d_A \approx 400$  Mpc ( $z \approx 0.11$ ). This emission could be the source of the extended X-rays seen on kiloparsec scales in the centres of some nearby AGN (e.g., Wang et al., 2010; Di Gesu et al., 2017; Lansbury et al., 2018), although more work is needed to model the photoionised gas which could also contribute on these spatial scales. Finally, we produced a mock observation using the *Chandra* space telescope to demonstrate the observability of our results.

## 6.2 Outlook

### 6.2.1 Linking Simulations & Observations to Constrain Feedback

In Chapter 3, we saw that looking for AGN residing in gas depleted and quiescent galaxies is not a good test for feedback. An AGN can slowly quench a galaxy by restricting the supply of cold gas which would result in a long decline in star formation and AGN activity due to them both sharing a common gas reservoir. Alternatively, the AGN could rapidly blow out all the gas, quickly removing future fuel for itself and star formation. Both of these scenarios result in the same outcome: we only observe AGN in gas-rich and star forming galaxies. Observations and models have suggested there may be ‘rapid’ ( $t \approx 100$  Myr) and a ‘slow’ ( $t \approx 1$  Gyr) quenching pathways (e.g., Wu et al., 2018; Belli et al., 2019; Akins et al., 2022) but it remains to be seen which one is the most common and what role, if any, AGN have to play in it. Studying this quenching timescale with cosmological simulations is limited due to their phenomenological models for AGN feedback – for example, TNG shows a rapid quenching of galaxies at a black hole mass of  $M_{\text{BH}} \approx 10^{8.2} M_{\odot}$ , but this is because it’s when the subgrid model switches to the more efficient kinetic mode (Weinberger et al., 2018; Terrazas et al., 2020). Therefore, to understand the timescale on which AGN affect their galaxies it is therefore crucial to use more physically-motivated feedback models.

In the last few years there has been significant progress in developing models of AGN feedback more closely linked to physical processes, for example, spin-driven jets (Talbot et al., 2021, 2022, 2024), quasar winds with resolved shock phases (Costa et al., 2020), and direct radiation pressure (Bieri et al., 2017; Costa et al., 2018a,b; Ishibashi et al., 2018).

Just as important is understanding how gas flows back into the galaxy to build a central gas reservoir (Verwilghen et al., 2024), as this will provide future fuel for AGN activity and central star formation. However, these models are computationally intensive and require higher spatial resolution than is currently feasible for a big-box cosmological simulation. Furthermore, a direct comparison of how these models impact galaxies is challenging, due to the differing codes used and galaxy setups investigated (although there have been some studies comparing different driving mechanisms directly; e.g., Cielo et al. 2018; Huško et al. 2024).


To distinguish between the models and understand when these different driving mechanisms are dominant, we need to use these experiments to make robust and testable observational predictions. However, making direct observational predictions for, e.g., a certain line profile, remains challenging. There have been some recent successes in predicting line emission by post-processing simulations with radiative transfer codes, for example, Meenakshi et al. (2022b) studied the [O III] emission expected from jet-ISM interactions and Costa et al. (2022) investigated predicted Lyman- $\alpha$  halos around high-redshift quasars. Predicting the emission from cold gas tracers, such as CO, is significantly more challenging, as it requires robust models for dust and molecular formation and destruction and non-equilibrium chemical networks. This adds significant computational complexity and have thus been limited to very idealised setups (e.g., Richings & Faucher-Giguère, 2018b; Richings et al., 2021).

To overcome these difficulties in producing observational predictions, there are ways in which simulators and observers can ‘meet in the middle’, for example, using gas or stellar kinematics, or properties derived from observations such as star formation histories. The rise of integral field units (IFUs) has led to the ability to characterise the gas across a whole galaxy at sub-kiloparsec scales. Figure 6.1 shows two examples of high-resolution observations using IFUs. Girdhar et al. (2022) studied a radio-quasar that hosted an inclined jet impacting the ISM. They found enhanced velocity dispersion in regions perpendicular to the jet direction; possibly an example of preventative feedback (see also Venturi et al. 2021). Bessiere & Ramos Almeida (2022) found an enhanced fraction of younger stellar populations on the edge of an AGN outflow, suggesting localised positive feedback. These properties could be compared to cosmological zoom-in scale simulations (e.g., IllustrisTNG50: Pillepich et al. 2019; Nelson et al. 2019a) that have been re-simulated with a range of AGN feedback models, allowing us to better constrain the immediate impact of different types of AGN feedback on their host galaxies.

Ever improving telescopes and instrumentation will further improve the information we can learn from these systems, and push our ability to resolve these features to higher redshifts. For example, *JWST* is able to detect PAH features in galaxies at Cosmic Noon ( $z \approx 2$ ). PAHs are expected to be quickly destroyed by radiation from young stars, making them a tracer of very recent ( $< 10$  Myr) star formation. This will provide further constraining power on the star formation histories of these galaxies. Additionally, the HARMONI instrument on the Extremely Large Telescope (ELT; first light scheduled 2028) will be able to resolve gas kinematics of galaxies at Cosmic Noon (García-Bernete et al., 2021) that is currently only feasible in the local Universe. Together, *JWST* and the ELT will help push our understanding of the impact of AGN feedback to even higher redshifts.

### 6.2.2 Understanding Cold Gas in Outflows


In Chapter 4, we found that an AGN wind interacting with a spatially-resolved ISM was able to launch multiphase outflows, containing entrained clouds of cold gas. However, this is not a trivial result – the existence of high-velocity cold gas in observed outflows (see review in Veilleux et al., 2020) is a longstanding puzzle, as the cloud crushing time



**Figure 6.1:** Detailed IFU observations can unveil the gas kinematics and star formation histories on sub-kpc scales, and link these to AGN effects such as outflows and jets. *Left panel:* an inclined jet increases the velocity dispersion of gas perpendicular to its orientation (Girdhar et al., 2022). *Right panel:* a higher proportion of young stellar populations are observed at the edge of an outflow, hinting at positive feedback Bessiere & Ramos Almeida (2022).

of such clouds is expected to be much shorter than the outflow time (e.g., Klein et al., 1994; Zhang et al., 2017; Schneider & Robertson, 2017). Recent studies using a single cloud in a ‘wind-tunnel’ setup have found that cold gas can survive in certain conditions, emphasising that radiative cooling is key to enduring its survival (although the density structure of the cloud also likely plays a role; see Banda-Barragán et al. 2019). The wind-cloud interaction at the interface between the phases (a ‘radiative mixing layer’; Fielding et al. 2020) creates a mixed phase of gas which can then condensate in the tail of the cloud, maintaining or even growing the cold gas mass (Gronke & Oh, 2018, 2020a). In Chapter 5, we showed that this mixed material is likely to be luminous in X-rays, opening up a way to observationally investigate this effect. However, there are still many open questions about the structure of the cold gas phase, such as whether it exists in a mist or small clumps (Gronke & Oh, 2020b) and what the scale of this structure is (McCourt et al., 2018). In a galaxy environment, there are also complex effects such as the effect of cloud shielding, which means that most cold gas clumps won’t experience a streamlined wind as is the case. There has been some work on how shocks would interact with multicloud systems (Banda-Barragán et al., 2020, 2021) but more work needs to be done to place this into a galactic environment.

In this context, our simulations presented in Chapter 4 can be considered a multiple cloud-crushing problem on a galactic scale. In Almeida et al. (prep) we dig deeper into the effect of the AGN wind on the cold gas in these simulations. By linking together connected cold gas cells, the distribution of clouds can be analysed. It was found that the initial clumpiness of the disc only has a small impact on the resulting size distribution of the clouds. This is because, as the wind hits the clump, it tends to fragment into smaller clumps that then become entrained. However, it was found that AGN luminosity plays the largest role in setting the outflowing cloud size, with brighter AGN shattering the clumps into smaller cloudlets. The number of cold gas clouds increases with luminosity until  $L_{\text{bol}} \approx 10^{46}$  erg s $^{-1}$ , beyond which it starts declining. This corroborates our findings in Chapter 4, where the mass outflow and energy rates have a turnover point at the same luminosity. This suggests that the brightest AGN are destroying the cold clouds before they can re-form through cooling



**Figure 6.2:** A targeted refinement method for cold, outflowing gas being developed by Almeida et al. (prep). The left column shows a low mass resolution of  $M_{\text{target}} = 10^6 M_{\odot}$  and the right column shows a high resolution of  $M_{\text{target}} = 2 \times 10^4 M_{\odot}$ . The middle column shows the targeted refinement scheme – the global mass resolution is the same as the low resolution case, but cells are refined further if they have a high Mach number (bottom row) or if they exceed their cooling mass (criteria from McCourt et al. 2018). We can see that the target refinement criteria capture the physics of the problem much more faithfully than the low resolution case, but at a much reduced computational cost compared to the high resolution example.

However, as we discussed in Chapter 4, it is incredibly challenging to spatially resolve the structure of the cold gas phase, with increasing resolution finding smaller and smaller clouds (e.g., Nelson et al., 2020). Uniformly increasing the resolution across the simulation domain is prohibitively expensive, but one solution is to use targeted refinement to only increase resolution in areas of interest (in this case, the cold gas). Other studies have used this targeted refinement approach on larger scales to look at cold gas in the CGM (Hummels et al., 2019; van de Voort et al., 2019; Ramesh & Nelson, 2024; Ramesh et al., 2024b,a). The simulations presented in this thesis already use targeted refinement in regions of high wind tracer density (Costa et al., 2020) and an additional refinement scheme for the cold clouds is being developed by Almeida et al. (prep). Gas cells are refined further if they meet two criteria: a Mach number of  $\mathcal{M} > 1$  (around  $v \approx 10 \text{ km s}^{-1}$  for the cold phase) and if their mass is greater than the cooling mass, above which the cloud should shatter into smaller cloudlets<sup>1</sup> Figure 6.2 shows how this targeted refinement scheme can more faithfully reproduce the high-resolution results, while also saving massively on computational cost.

To fully understand the evolution of the cold phase of the outflow, in addition to just better resolution, more physical effects also need to be included. Magnetic fields may enhance (Ji et al., 2018) or suppress (Gronke & Oh, 2020a; Jung et al., 2023) thermal instabilities. Alignment plays an important role, with aligned fields possibly increasing the phase mixing and transverse magnetic fields leading to ‘draping’ where the cloud is shielded by the field, but also squeezed which can lead to mass loss (Cottle et al., 2020).

<sup>1</sup>The cooling mass is defined as the mass associated to the cooling length, see McCourt et al. (2018).

Additionally, self-gravity of the clouds may play a role (Mandal et al., 2024), increasing their survival time. Low temperature cooling should also be considered, which requires careful modelling of dust and molecule formation, as well as radiative transfer.

To tie these models to observations, we could investigate the cold HI clouds seen entrained in the Milky Way’s galactic outflow (Di Teodoro et al., 2019, 2020). For example, Banda-Barragán et al. (2021) took their shock-multicloud simulation data and performed mock observations, degrading the simulated resolution to match various telescopes. A similar approach could be taken with our work: modifying our initial disc to better match the Milky Way would allow us to investigate the similarities or differences between our results and these observed clouds, for example, by looking at their morphological and velocity distributions. This could help reveal whether these clouds were launched by an AGN- or starburst-driven wind.

### 6.2.3 X-ray and Radio Emission from Wind-ISM Interactions

In Chapter 5, we estimated the Bremsstrahlung emission from our idealised simulations. We also showed a proof-of-concept synthetic image of this emission from the *Chandra* space telescope.

A key improvement to this work would be to estimate the spectrum or radio hardness of this emission. Star formation is expected to dominate in softer bands ( $\approx 0.1 - 2$  keV) and emission from the AGN itself at harder band ( $\approx 2 - 10$  keV). Establishing the expected hardness of the predicted free-free emission from wind-ISM mixing would be a testable prediction to compare against the observations; for example, observational work by Di Gesu et al. (2017) suggests that collisionally excited X-ray emission should peak in the soft band. The modelling involved in the synthetic *Chandra* image in Chapter 5 using the PyXSIM package did yield a simulated spectra for the X-ray emission. However, more work needs to be done to ensure this is robust, for example, establishing that the thermal collisional ionisation model used is appropriate for the case of wind-ISM mixing, and investigating the effect of including metal-line cooling on the resulting spectra.

Radiative transfer modelling would also play a useful role in further contextualising our results. We predicted that wind-cloud mixing should produce observable X-rays on kiloparsec scales within the centre of galaxies, beyond the point-source of the accretion disc. Some studies (e.g. Lansbury et al., 2018) have found extended X-rays on these scales, but it is unclear whether they are caused by mixing-induced Bremsstrahlung emission, or other effects such as photoionisation or scattering (Wang et al., 2010; Di Gesu et al., 2017). Meenakshi et al. (2022a) post-processed their idealised jet-ISM simulation with a radiative transfer code to investigate the interplay between collisional and radiative excitation, finding that the collisional excitation from the jet dominated. A similar approach could be used with our simulations to then also predict the X-ray emission from this ionised gas, allowing us to compare more directly with observed quasar targets and determine the relative contribution of wind-ISM mixing.

When considering how to combine simulations and observations, it is intriguing to note that our own Milky Way is a potential test-bed for models of AGN feedback, especially in the X-ray. The *Fermi*/*eROSITA* bubbles (Su et al., 2010; Predehl et al., 2020) provide evidence of a previous outburst from the galactic centre, although it is disputed whether this is caused by AGN (e.g., Zubovas et al., 2011; Zhang & Guo, 2020; Yang et al., 2021) or starburst (e.g., Lacki, 2014) activity. There have also been detections of narrow ‘chimneys’ of X-ray emitting gas or around 160 pc in length linked to the larger bubbles (Ponti et al., 2019, 2021). Our simulations also showed these features – in particular, we found that the clumpiness of the ISM played a major role in shaping the chimneys, with the region behind an initially high-density cloud showing high X-ray luminosity due to the tail being formed of strongly mixed material and the high-velocity vents either side shaping

this tail. A similar setup could be used to create a more Milky-Way like galaxy to see if similar features can be replicated and explore how sensitive these chimneys are to the ISM distribution in regions around the galactic centre.

We demonstrated that resolving the spatially-extended free-free emission from wind-ISM interactions is possible with high-resolution X-ray telescopes, such as *Chandra*. However, the future of *Chandra* is currently uncertain due to financial pressures, with mission operations facing significant budget cuts<sup>2</sup>. This would represent a significant loss to the astronomical community.

More optimistically, looking ahead to the next generation of X-ray instruments, the *Lynx* Space Telescope<sup>3</sup> is a mission concept aiming to be launched in the mid-2030s. Although *Chandra* has excellent 0.5'' spatial resolution, this is only in the centre, and blooms off-axis. *Lynx* will also achieve the same 0.5'' resolution, but will maintain this sub-arcsec resolution out to 10', representing a sixteen-fold improvement in area over *Chandra*. Combined with order-of-magnitude improvements in sensitivity and spectral resolution, *Lynx* will represent a significant improvement in our ability to study the X-rays sky, especially for science cases such as the one described here.

### Radio signatures of outflows

Radio observations are a powerful tool to study AGN feedback, however, the origin of radio emission, especially in radio-quiet (RQ) sources without large jets, is unclear. A significant contribution to this emission in RQ galaxies could be caused by synchrotron emission from wind-driven shocks in the ISM (Panessa et al., 2019). Several observational works have suggested that wind shocks could be responsible for the higher radio detection fraction in galaxies with observed outflows (e.g., Zakamska & Greene, 2014; Petley et al., 2022; Fawcett et al., 2023) but so far this subject has received limited theoretical attention.

As well as discussing the expected X-ray emission from AGN wind, Nims et al. (2015) also used their analytic model to calculate the contribution from non-thermal process, such as synchrotron emission. However, this was again performed with an analytic model assuming a homogeneous ambient medium. As we showed in Chapter 5, the resulting emission from a realistic ISM can differ markedly from an isotropic case. Therefore, it would be interesting to use the simulations presented in Chapter 4 to investigate the effect of ISM clumpiness on the predicted radio emission. Shock fronts around the cold clumps could accelerate electrons to relativistic speeds, resulting in synchrotron emission. To investigate this, our wind-ISM simulations could be post-processed with a shock finder (e.g., Schaal & Springel, 2015) to estimate the synchrotron emission. As with the X-rays in Chapter 5, this shock-driven radio emission would need to be disentangled from the point-source corona, and the diffuse star formation.

Beyond our idealised setup, modelling the non-thermal component self-consistently is extremely challenging due to the need to follow the production and evolution of cosmic rays (Ruszkowski & Pfrommer, 2023). However, there have been recent successes at incorporating cosmic ray modelling into hydrodynamic codes, including AREPO (Pfrommer et al., 2017; Winner et al., 2019). This has allowed the non-thermal component to be studied in galaxy-scale simulations (Werhahn et al., 2021a,b; Pfrommer et al., 2022). Combining these simulations with physically-motivated AGN models would be a fascinating way to create a truly multiwavelength view of AGN feedback at work.

It is an exciting time for radio astronomy, with upgrades to LOFAR and the Square Kilometre Array (SKA) both coming online over the next few years, thereby making robust theoretical predictions, especially at low frequencies vital to exploit the full potential of these instruments.

---

<sup>2</sup>[www.savechandra.org](http://www.savechandra.org)

<sup>3</sup>[www.lynxobservatory.com](http://www.lynxobservatory.com)



### 6.2.4 The Next Decade in Galaxy Simulations

As we saw in Chapter 2, the computational power available to astronomers has increased exponentially over the last few decades. However, we can't simply rely on faster processors solving all our problems in the future. Even at the current rate of progress, in a decade's time, the spatial resolution of cosmological-scale simulations will still only be  $\sim 100$  pc for big-boxes, and  $\sim 10$  pc for zoom-ins<sup>4</sup>, which is insufficient to study many of the processes discussed in this thesis, such as ISM substructure and cold gas entrainment. Therefore, new ways of exploring these processes in simulations need to be developed to continue making progress.

#### Subgrid cold gas

The interaction between the different phases of the gas is critical in determining the survival and entrainment of cold gas in galactic outflows (e.g., [Tan & Fielding, 2024](#); [Ward et al., 2024](#)). However, it requires extremely high spatial resolution. One approach to overcome this is to rethink how we model cold gas in simulations and include the interactions between different gas phases as subgrid models.

For example, in the ARKENSTONE model ([Smith et al., 2024](#)), supernovae inject hot and cold components separately, including 'cloud particles' which can exchange mass and momentum with the ambient medium before re-coupling with the gas. An even more general approach is to evolve the hot and cold gas as two separate fluids across the whole simulation domain by solving two sets of hydrodynamic equations which interact via source terms ([Weinberger & Hernquist, 2023](#); [Butsky et al., 2024](#)). In both cases, the key challenge is to constrain the mixing factors that model the subgrid interactions between the cold clouds and the hot background which are parameterised from high-resolution simulations (e.g., [Fielding & Bryan, 2022](#)).

The eventual aim of these projects is to be able to model the cold gas phase more accurately in cosmological-scale simulations. However, constraining the mixing terms between cold clouds and hot winds across a range of environments, such as in a multi-cloud problem where shielding is in play, is likely to pose a key challenge in the coming years. This is where my ACDC simulations could play a role in quantifying this interaction on galaxy scales.

#### Machine learning approaches

The last few years have seen machine learning applications break into the public sphere, such as ChatGPT for text-based interactions, and DALLÉ for image generation. These AI tools have the potential to radically transform the way astronomers interact with their data, especially as we enter the era of 'big-data' with forthcoming facilities such as Vera Rubin Observatory and the Square Kilometre Array (SKA) predicted to generate terabytes of data per second. There are endless use for machine learning in astrophysics; here I will briefly outline a few applications as a taste for what's to come in the next decade:

Machine learning can be used as an analysis tool for existing datasets. For example, [Piotrowska et al. \(2022\)](#) performed a random forest analysis on Illustris, IllustrisTNG and EAGLE to determine which galaxy property was the best predictor for the galaxy being quiescent. They found that in all three simulations, the black hole mass was the most important classifier which they interpreted as evidence for the integrated effect of AGN feedback.

---

<sup>4</sup>Assuming computing power doubles every 18 months, and recalling that doubling the spatial resolution requires  $16\times$  the processing power (see [Feldmann et al. 2023](#) for some recent trends in boxesizes/resolution).

As we have discussed, one of the biggest challenges when studying galaxy-wide AGN outflows is the high computational cost required to accurately capture the phases of the outflow. In [Stachenfeld et al. \(2021\)](#), a convolutional neural network was trained on high-resolution simulations of turbulence. It was then found to be able to replicate the turbulence at a coarser resolution more accurately than traditional solvers. This could offer a new way of designing subgrid models that reduces the computational cost required to study the effect of turbulence on galaxy scales, such as in and around AGN outflows.


Machine learning can also be used to speed-up the construction of cosmological datasets by reducing the computational cost required. For example, work by [Kamdar et al. \(2016\)](#); [Agarwal et al. \(2018\)](#); [Das et al. \(2024\)](#) trained machine-learning algorithms on hydrodynamic cosmological simulations, such as Illustris and SIMBA, to be able to predict baryonic properties of galaxies from only their dark matter content. Similar in aim to semi-analytic models, these ‘paint-on’ galaxy properties to large N-body simulations, allowing large datasets to be rapidly generated and studied.

Finally, on the largest scales, machine learning can be used to infer cosmological parameters. The CAMELS suite of simulations ([Villaescusa-Navarro et al., 2021](#); [Ni et al., 2023](#)) features over 4000 cosmological simulations with a range of astrophysical (AGN and SF feedback) and cosmological ( $\Omega_m$ ,  $\Omega_b$ ,  $\sigma_8$ ) parameters for the purpose of training machine-learning algorithms. This can allow for rapid forward-modelling to compare to observations from current and forthcoming cosmological surveys such as DESI or *Euclid*. One of the main free parameters of these models is AGN feedback, demonstrating how the work done in this Thesis can be used, even in a small way, to uncover the most profound questions about the nature of our Universe.

## 6.3 Final Remarks

It is an exciting time to be studying AGN: during the course of this thesis the Nobel Prize for Physics in 2020 was awarded for the confirmation that our own galaxy hosts a supermassive black hole, and in 2022, telescopes across the world combined to image it directly (Event Horizon Telescope Collaboration et al., 2022). The future for this field is bright, with the stunning data from *JWST* already revolutionising our understanding of black hole formation and growth at high redshift. Many more exciting telescopes are under construction or in the planning phase, such as the ELT, SKA, *Lynx*, *Athena* and *LISA*; all of which have the potential to open new frontiers in our understanding. Increasing computing power, the development of more efficient numerical techniques, and the integration of machine learning into the scientific process, will further push astronomers to be creative and devise new theories to try and understand this dizzying array of data.

I finish by showing the EHT images of the black hole in the centre of M87 and SgrA\* (Figure 6.3). These two images required a global network of telescope and a team of hundreds of scientists and engineers from across the world to produce. They stand as a testament to human curiosity, collaboration and community, and serve as a reminder that even our own galaxy has been shaped by one of the Universe’s strangest and most mysterious phenomena – supermassive black holes.



**Figure 6.3:** The first image of an AGN accretion disc in M87 (*left*) and the black hole in the centre of our own galaxy (*right*) observed using the Event Horizon Telescope. Image credit: EHT/ESO



# Bibliography

- Aarseth S. J., 1963, *MNRAS*, **126**, 223
- Aarseth S. J., 1966, *MNRAS*, **132**, 35
- Abazajian K. N., et al., 2009, *ApJS*, **182**, 543
- Abramowicz M. A., Fragile P. C., 2013, *Living Reviews in Relativity*, **16**, 1
- Agarwal S., Davé R., Bassett B. A., 2018, *MNRAS*, **478**, 3410
- Akins H. B., Narayanan D., Whitaker K. E., Davé R., Lower S., Bezanson R., Feldmann R., Kriek M., 2022, *ApJ*, **929**, 94
- Al Yazeedi A., Katkov I. Y., Gelfand J. D., Wylezalek D., Zakamska N. L., Liu W., 2021, *ApJ*, **916**, 102
- Almeida I., Nemmen R., 2020, *MNRAS*, **492**, 2553
- Almeida I., Nemmen R., Riffel R. A., 2023, *MNRAS*, **526**, 217
- Almeida I., Costa T., Harrison C. M., Ward S. R., in prep.
- Anglés-Alcázar D., Özel F., Davé R., Katz N., Kollmeier J. A., Oppenheimer B. D., 2015, *ApJ*, **800**, 127
- Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017, *MNRAS*, **464**, 2840
- Anglés-Alcázar D., et al., 2021, *ApJ*, **917**, 53
- Appleby S., Davé R., Kraljic K., Anglés-Alcázar D., Narayanan D., 2020, *MNRAS*, **494**, 6053
- Appleby S., Davé R., Sorini D., Storey-Fisher K., Smith B., 2021, *MNRAS*, **507**, 2383
- Arrabal Haro P., et al., 2023, *Nature*, **622**, 707
- Arrigoni Battaia F., Hennawi J. F., Prochaska J. X., Oñorbe J., Farina E. P., Cantalupo S., Lusso E., 2019, *MNRAS*, **482**, 3162
- Azadi M., et al., 2015, *ApJ*, **806**, 187
- Azadi M., et al., 2017, *ApJ*, **835**, 27
- Baldry I. K., Glazebrook K., Brinkmann J., Ivezić Ž., Lupton R. H., Nichol R. C., Szalay A. S., 2004, *ApJ*, **600**, 681
- Banda-Barragán W. E., Zertuche F. J., Federrath C., García Del Valle J., Brüggén M., Wagner A. Y., 2019, *MNRAS*, **486**, 4526
- Banda-Barragán W. E., Brüggén M., Federrath C., Wagner A. Y., Scannapieco E., Cottle J., 2020, *MNRAS*, **499**, 2173
- Banda-Barragán W. E., Brüggén M., Heesen V., Scannapieco E., Cottle J., Federrath C., Wagner A. Y., 2021, *MNRAS*, **506**, 5658
- Baron D., Netzer H., 2019, *MNRAS*, **486**, 4290
- Baron D., et al., 2018, *MNRAS*, **480**, 3993

- Baron D., Netzer H., Lutz D., Prochaska J. X., Davies R. I., 2022, *MNRAS*, **509**, 4457
- Baumgartner W. H., Tueller J., Markwardt C. B., Skinner G. K., Barthelmy S., Mushotzky R. F., Evans P. A., Gehrels N., 2013, *ApJS*, **207**, 19
- Beckmann R. S., et al., 2017, *MNRAS*, **472**, 949
- Begelman M. C., Blandford R. D., Rees M. J., 1984, *Reviews of Modern Physics*, **56**, 255
- Belli S., Newman A. B., Ellis R. S., 2019, *ApJ*, **874**, 17
- Bennett J. S., Sijacki D., 2022, *MNRAS*, **514**, 313
- Bennett J. S., Sijacki D., Costa T., Laporte N., Witten C., 2024, *MNRAS*, **527**, 1033
- Benson A. J., 2010, *Phys. Rep.*, **495**, 33
- Bernhard E., Mullaney J. R., Daddi E., Ciesla L., Schreiber C., 2016, *MNRAS*, **460**, 902
- Bertola E., et al., 2024, accepted A&A
- Bessiere P. S., Ramos Almeida C., 2022, *MNRAS*, **512**, L54
- Best P. N., Heckman T. M., 2012, *MNRAS*, **421**, 1569
- Bieri R., Dubois Y., Rosdahl J., Wagner A., Silk J., Mamon G. A., 2017, *MNRAS*, **464**, 1854
- Biffi V., Dolag K., Böhringer H., Lemson G., 2012, *MNRAS*, **420**, 3545
- Biffi V., Dolag K., Böhringer H., 2013, *MNRAS*, **428**, 1395
- Bischetti M., et al., 2017, *A&A*, **598**, A122
- Bischetti M., et al., 2019, *A&A*, **628**, A118
- Bischetti M., et al., 2021, *A&A*, **645**, A33
- Bischetti M., et al., 2024, *ApJ*, **970**, 9
- Blandford R. D., Znajek R. L., 1977, *MNRAS*, **179**, 433
- Blandford R., Meier D., Readhead A., 2019, *ARA&A*, **57**, 467
- Bluck A. F. L., et al., 2020, *MNRAS*, **499**, 230
- Blumenthal G. R., Faber S. M., Primack J. R., Rees M. J., 1984, *Nature*, **311**, 517
- Boehringer H., Voges W., Fabian A. C., Edge A. C., Neumann D. M., 1993, *MNRAS*, **264**, L25
- Bondi H., 1952, *MNRAS*, **112**, 195
- Booth C. M., Schaye J., 2009, *MNRAS*, **398**, 53
- Booth C. M., Schaye J., 2010, *MNRAS*, **405**, L1
- Borrow J., Schaller M., Bower R. G., Schaye J., 2022, *MNRAS*, **511**, 2367
- Bothwell M. S., et al., 2013, *MNRAS*, **429**, 3047
- Bourne M. A., Sijacki D., 2021, *MNRAS*, **506**, 488
- Bourne M. A., Yang H.-Y. K., 2023, *Galaxies*, **11**, 73
- Bourne M. A., Zubovas K., Nayakshin S., 2015, *MNRAS*, **453**, 1829
- Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole S., Lacey C. G., 2006, *MNRAS*, **370**, 645
- Brinchmann J., Charlot S., White S. D., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, **351**, 1151
- Brüggen M., Scannapieco E., 2020, *ApJ*, **905**, 19
- Bundy K., Ellis R. S., Conselice C. J., 2005, *ApJ*, **625**, 621

- Butsky I. S., Hummels C. B., Hopkins P. F., Quinn T. R., Werk J. K., 2024, [arXiv e-prints](#), p. [arXiv:2402.03419](#)
- Carniani S., et al., 2015, *A&A*, **580**, A102
- Cattaneo A., et al., 2009, *Nature*, **460**, 213
- Chartas G., et al., 2021, *ApJ*, **920**, 24
- Chatterjee K., Liska M., Tchekhovskoy A., Markoff S. B., 2019, *MNRAS*, **490**, 2200
- Chen Z., He Z., Ho L. C., Gu Q., Wang T., Zhuang M., Liu G., Wang Z., 2022, *Nature Astronomy*,  
Cheung E., et al., 2016, *Nature*, **533**, 504
- Choi E., Somerville R. S., Ostriker J. P., Naab T., Hirschmann M., 2018, *ApJ*, **866**, 91
- Cicone C., et al., 2014, *A&A*, **562**, A21
- Cicone C., Brusa M., Ramos Almeida C., Cresci G., Husemann B., Mainieri V., 2018, *Nature Astronomy*, **2**, 176
- Cielo S., Bieri R., Volonteri M., Wagner A. Y., Dubois Y., 2018, *MNRAS*, **477**, 1336
- Ciotti L., Ostriker J. P., 2007, *ApJ*, **665**, 1038
- Ciotti L., Ostriker J. P., 2012, in Kim D.-W., Pellegrini S., eds, *Astrophysics and Space Science Library Vol. 378*, *Astrophysics and Space Science Library*. p. 83 ([arXiv:1104.2238](#)),  
[doi:10.1007/978-1-4614-0580-1\\_4](#)
- Circosta C., et al., 2021, *A&A*, **646**, A96
- Cole S., Lacey C. G., Baugh C. M., Frenk C. S., 2000, *MNRAS*, **319**, 168
- Concas A., Popesso P., Brusa M., Mainieri V., Thomas D., 2019, *A&A*, **622**, A188
- Condon J. J., Yin Q. F., 1990, *ApJ*, **357**, 97
- Cooper J. L., Bicknell G. V., Sutherland R. S., Bland-Hawthorn J., 2008, *ApJ*, **674**, 157
- Costa T., Sijacki D., Trenti M., Haehnelt M. G., 2014a, *MNRAS*, **439**, 2146
- Costa T., Sijacki D., Haehnelt M. G., 2014b, *MNRAS*, **444**, 2355
- Costa T., Sijacki D., Haehnelt M. G., 2015, *MNRAS*, **448**, L30
- Costa T., Rosdahl J., Sijacki D., Haehnelt M. G., 2018a, *MNRAS*, **473**, 4197
- Costa T., Rosdahl J., Sijacki D., Haehnelt M. G., 2018b, *MNRAS*, **479**, 2079
- Costa T., Pakmor R., Springel V., 2020, *MNRAS*, **497**, 5229
- Costa T., Arrigoni Battaia F., Farina E. P., Keating L. C., Rosdahl J., Kimm T., 2022, *MNRAS*,  
**517**, 1767
- Cottle J., Scannapieco E., Brüggem M., Banda-Barragán W., Federrath C., 2020, *ApJ*, **892**, 59
- Cox D. P., 2005, *ARA&A*, **43**, 337
- Crain R. A., van de Voort F., 2023, *ARA&A*, **61**, 473
- Crain R. A., et al., 2015, *MNRAS*, **450**, 1937
- Creasey P., Theuns T., Bower R. G., Lacey C. G., 2011, *MNRAS*, **415**, 3706
- Crenshaw D. M., Kraemer S. B., 2000, *ApJ*, **532**, L101
- Cresci G., et al., 2015a, *A&A*, **582**, A63
- Cresci G., et al., 2015b, *ApJ*, **799**, 82
- Croston J. H., Hardcastle M. J., Kharb P., Kraft R. P., Hota A., 2008, *ApJ*, **688**, 190
- Croton D. J., et al., 2006, *MNRAS*, **365**, 11
- Cui W., Davé R., Peacock J. A., Anglés-Alcázar D., Yang X., 2021, *Nature Astronomy*, **5**, 1069

- Curtis M., Sijacki D., 2015, *MNRAS*, 454, 3445
- Curtis-Lake E., Bluck A., d'Eugenio F., Maiolino R., Sijacki D., 2023a, *Nature Astronomy*, 7, 247
- Curtis-Lake E., et al., 2023b, *Nature Astronomy*, 7, 622
- Dalla Vecchia C., Schaye J., 2008, *MNRAS*, 387, 1431
- Das P. K., Davé R., Cui W., 2024, *arXiv e-prints*, p. arXiv:2406.16103
- Dasyra K. M., Combes F., 2011, *A&A*, 533, L10
- Davé R., Thompson R., Hopkins P. F., 2016, *MNRAS*, 462, 3265
- David L. P., Nulsen P. E. J., McNamara B. R., Forman W., Jones C., Ponman T., Robertson B., Wise M., 2001, *ApJ*, 557, 546
- Davies J. J., Crain R. A., Oppenheimer B. D., Schaye J., 2020, *MNRAS*, 491, 4462
- Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
- Dekel A., Birnboim Y., 2006, *MNRAS*, 368, 2
- Dekel A., Silk J., 1986, *ApJ*, 303, 39
- Di Gesu L., Costantini E., Piconcelli E., Kaastra J. S., Mehdipour M., Paltani S., 2017, *A&A*, 608, A115
- Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
- Di Teodoro E. M., et al., 2019, *ApJ*, 885, L32
- Di Teodoro E. M., McClure-Griffiths N. M., Lockman F. J., Armillotta L., 2020, *Nature*, 584, 364
- Diemer B., et al., 2018, *ApJS*, 238, 33
- Diemer B., et al., 2019, *MNRAS*, 487, 1529
- Donnari M., et al., 2019, *MNRAS*, 485, 4817
- Dubois Y., Gavazzi R., Peirani S., Silk J., 2013, *MNRAS*, 433, 3297
- Dubois Y., et al., 2014, *MNRAS*, 444, 1453
- Dubois Y., Peirani S., Pichon C., Devriendt J., Gavazzi R., Welker C., Volonteri M., 2016, *MNRAS*, 463, 3948
- Dubois Y., et al., 2021, *A&A*, 651, A109
- Dunn R. J. H., Fabian A. C., 2006, *MNRAS*, 373, 959
- Ellison S. L., Patton D. R., Hickox R. C., 2015, *MNRAS*, 451, L35
- Ellison S. L., et al., 2021, *MNRAS*, 505, L46
- Emsellem E., Renaud F., Bournaud F., Elmegreen B., Combes F., Gabor J. M., 2015, *MNRAS*, 446, 2468
- Event Horizon Telescope Collaboration et al., 2022, *ApJ*, 930, L12
- Fabian A. C., 1994, *ARA&A*, 32, 277
- Fabian A. C., 1999, *MNRAS*, 308, L39
- Fabian A. C., 2012, *ARA&A*, 50, 455
- Fanaroff B. L., Riley J. M., 1974, *MNRAS*, 167, 31P
- Faucher-Giguère C.-A., Quataert E., 2012, *MNRAS*, 425, 605
- Fawcett V. A., et al., 2023, *MNRAS*, 525, 5575
- Feldmann R., Mayer L., 2015, *MNRAS*, 446, 1939
- Feldmann R., et al., 2023, *MNRAS*, 522, 3831



- Ferriere K. M., 2001, *Reviews of Modern Physics*, **73**, 1031
- Fielding D. B., Bryan G. L., 2022, *ApJ*, **924**, 82
- Fielding D. B., Ostriker E. C., Bryan G. L., Jermyn A. S., 2020, *ApJ*, **894**, L24
- Fiore F., et al., 2017, *A&A*, **601**, 143
- Fischera J., Dopita M. A., Sutherland R. S., 2003, *ApJ*, **599**, L21
- Florez J., et al., 2020, *MNRAS*, **497**, 3273
- Fluetsch A., et al., 2019, *MNRAS*, **483**, 4586
- Frias Castillo M., et al., 2024, *A&A*, **683**, A211
- Furlong M., et al., 2015, *MNRAS*, **450**, 4486
- García-Bernete I., et al., 2021, *A&A*, **645**, A21
- Gatto L., Storchi-Bergmann T., Riffel R. A., Riffel R., Rembold S. B., Schimoia J. S., Mallmann N. D., Ilha G. S., 2024, *MNRAS*, **530**, 3059
- Gent F. A., Shukurov A., Fletcher A., Sarson G. R., Mantere M. J., 2013, *MNRAS*, **432**, 1396
- Genzel R., Eisenhauer F., Gillessen S., 2010, *Reviews of Modern Physics*, **82**, 3121
- Ghez A. M., et al., 2008, *ApJ*, **689**, 1044
- Girdhar A., et al., 2022, *MNRAS*, **512**, 1608
- Girdhar A., et al., 2024, *MNRAS*, **527**, 9322
- Girichidis P., et al., 2020, *Space Sci. Rev.*, **216**, 68
- Gnedin N. Y., Kravtsov A. V., 2011, *ApJ*, **728**, 88
- Godunov S. K., 1959, *Mat. Sbornik*, **47**(89), 271–306
- Gofford J., Reeves J. N., McLaughlin D. E., Braitto V., Turner T. J., Tombesi F., Cappi M., 2015, *MNRAS*, **451**, 4169
- González-Alfonso E., et al., 2017, *ApJ*, **836**, 11
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, **622**, 759
- Greene J. E., Ho L. C., 2005, *ApJ*, **630**, 122
- Greene J. E., Pooley D., Zakamska N. L., Comerford J. M., Sun A.-L., 2014, *ApJ*, **788**, 54
- Greene J. E., Strader J., Ho L. C., 2020, *ARA&A*, **58**, 257
- Gronke M., Oh S. P., 2018, *MNRAS*, **480**, L111
- Gronke M., Oh S. P., 2020a, *MNRAS*, **492**, 1970
- Gronke M., Oh S. P., 2020b, *MNRAS*, **494**, L27
- Gültekin K., et al., 2009, *ApJ*, **698**, 198
- Gürkan G., et al., 2015, *MNRAS*, **452**, 3776
- Habouzit M., et al., 2021, *MNRAS*, **503**, 1940
- Habouzit M., et al., 2022a, *MNRAS*, **509**, 3015
- Habouzit M., et al., 2022b, *MNRAS*, **511**, 3751
- Harrison C. M., 2017, *Nature Astronomy*, **1**, 0165
- Harrison C. M., Ramos Almeida C., 2024, *Galaxies*, **12**, 17
- Harrison C. M., et al., 2012, *ApJ*, **760**, L15
- Harrison C. M., Alexander D. M., Mullaney J. R., Swinbank A. M., 2014, *MNRAS*, **441**, 3306

- Harrison C. M., Thomson A. P., Alexander D. M., Bauer F. E., Edge A. C., Hogan M. T., Mullaney J. R., Swinbank A. M., 2015, *ApJ*, **800**, 45
- Harrison C. M., Costa T., Tadhunter C. N., Flütsch A., Kakkad D., Perna M., Vietri G., 2018, *Nature Astronomy*, **2**, 198
- Hayward C. C., Hopkins P. F., 2017, *MNRAS*, **465**, 1682
- Heckman T. M., Armus L., Miley G. K., 1990, *ApJS*, **74**, 833
- Henriques B. M. B., White S. D. M., Thomas P. A., Angulo R., Guo Q., Lemson G., Springel V., Overzier R., 2015, *MNRAS*, **451**, 2663
- Henriques B. M. B., White S. D. M., Lilly S. J., Bell E. F., Bluck A. F. L., Terrazas B. A., 2019, *MNRAS*, **485**, 3446
- Hervella Seoane K., Ramos Almeida C., Acosta-Pulido J. A., Speranza G., Tadhunter C. N., Bessiere P. S., 2023, *A&A*, **680**, A71
- Hickox R. C., Mullaney J. R., Alexander D. M., Chen C.-T. J., Civano F. M., Goulding A. D., Hainline K. N., 2014, *ApJ*, **782**, 9
- Hidalgo-Pineda F., Farber R. J., Gronke M., 2024, *MNRAS*, **527**, 135
- Hirashita H., Murga M. S., 2020, *MNRAS*, **492**, 3779
- Hirschmann M., Dolag K., Saro A., Bachmann L., Borgani S., Burkert A., 2014, *MNRAS*, **442**, 2304
- von Hoerner S., 1960, *ZAp*, **50**, 184
- Holden L. R., Tadhunter C. N., 2023, *MNRAS*, **524**, 886
- Holden L. R., Tadhunter C. N., Morganti R., Oosterloo T., 2023, *MNRAS*, **520**, 1848
- Holmberg E., 1941, *ApJ*, **94**, 385
- Hopkins P. F., 2015, *MNRAS*, **450**, 53
- Hopkins P. F., Quataert E., 2011, *MNRAS*, **415**, 1027
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, **445**, 581
- Hummels C. B., et al., 2019, *ApJ*, **882**, 156
- Husemann B., Scharwächter J., Bennert V. N., Mainieri V., Woo J. H., Kakkad D., 2016, *A&A*, **594**, A44
- Husemann B., Davis T. A., Jahnke K., Dannerbauer H., Urrutia T., Hodge J., 2017, *MNRAS*, **470**, 1570
- Huško F., Lacey C. G., Schaye J., Nobels F. S. J., Schaller M., 2024, *MNRAS*, **527**, 5988
- Ichikawa K., et al., 2019, *ApJ*, **870**, 31
- Inayoshi K., Visbal E., Haiman Z., 2020, *ARA&A*, **58**, 27
- Irodoutou D., et al., 2022, *MNRAS*, **513**, 3768
- Ishibashi W., Fabian A. C., Maiolino R., 2018, *MNRAS*, **476**, 512
- Jackson T. M., Rosario D. J., Alexander D. M., Scholtz J., McAlpine S., Bower R. G., 2020, *MNRAS*, **498**, 2323
- Jarvis M. E., et al., 2019, *MNRAS*, **485**, 2710
- Jarvis M. E., et al., 2020, *MNRAS*, **498**, 1560
- Jarvis M. E., et al., 2021, *MNRAS*, **503**, 1780
- Ji S., Oh S. P., McCourt M., 2018, *MNRAS*, **476**, 852

- Ji Z., Giavalisco M., Kirkpatrick A., Kocevski D., Daddi E., Delvecchio I., Hatcher C., 2022, *ApJ*, **925**, 74
- Jung S. L., Grønnow A., McClure-Griffiths N. M., 2023, *MNRAS*, **522**, 4161
- Kakkad D., et al., 2017, *MNRAS*, **468**, 4205
- Kakkad D., et al., 2020, *A&A*, **642**, A147
- Kamdar H. M., Turk M. J., Brunner R. J., 2016, *MNRAS*, **457**, 1162
- Kang D., Woo J.-H., 2018, *ApJ*, **864**, 124
- Kashlinsky A., Rees M. J., 1983, *MNRAS*, **205**, 955
- Kauffmann G., Haehnelt M., 2000, *MNRAS*, **311**, 576
- Kellermann K. I., Sramek R., Schmidt M., Shaffer D. B., Green R., 1989, *ApJ*, **98**, 1195
- Kennicutt R. C., 1998, *ARA&A*, **36**, 189
- Khandai N., Di Matteo T., Croft R., Wilkins S., Feng Y., Tucker E., DeGraf C., Liu M.-S., 2015, *MNRAS*, **450**, 1349
- Khochfar S., Ostriker J. P., 2008, *ApJ*, **680**, 54
- Khoperskov S., Haywood M., Di Matteo P., Lehnert M. D., Combes F., 2018, *A&A*, **609**, A60
- Kim C., et al., 2022, *ApJ*, **928**, 73
- King A., 2003, *ApJ*, **596**, L27
- King A., 2005, *ApJ*, **635**, L121
- King A., Pounds K., 2015, *ARA&A*, **53**, 115
- Kirchschlager F., Mattsson L., Gent F. A., 2022, *MNRAS*, **509**, 3218
- Kirkpatrick A., Sharon C., Keller E., Pope A., 2019, *ApJ*, **879**, 41
- Klein R. I., McKee C. F., Colella P., 1994, *ApJ*, **420**, 213
- Kormendy J., Ho L. C., 2013, *ARA&A*, **51**, 511
- Koss M., et al., 2017, *ApJ*, **850**, 74
- Koss M. J., et al., 2021, *ApJS*, **252**, 29
- Koudmani S., Henden N. A., Sijacki D., 2021, *MNRAS*, **503**, 3568
- Koudmani S., Sijacki D., Smith M. C., 2022, *MNRAS*, **516**, 2112
- Lacki B. C., 2014, *MNRAS*, **444**, L39
- Lagos C. D. P., et al., 2015, *MNRAS*, **452**, 3815
- Lamperti I., et al., 2022, *A&A*, **668**, A45
- Langan I., et al., 2023, *MNRAS*, **521**, 546
- Langan I., Popping G., Ginolfi M., Gentile F., Valentino F., Kaasinen M., 2024, *arXiv e-prints*, p. [arXiv:2406.01673](https://arxiv.org/abs/2406.01673)
- Lansbury G. B., Jarvis M. E., Harrison C. M., Alexander D. M., Del Moro A., Edge A. C., Mullaney J. R., Thomson A. P., 2018, *ApJ*, **856**, L1
- Larson R. L., et al., 2023, *ApJ*, **953**, L29
- van Leer B., 1979, *Journal of Computational Physics*, **32**, 101
- Lehmer B. D., Alexander D. M., Bauer F. E., Brandt W. N., Goulding A. D., Jenkins L. P., Ptak A., Roberts T. P., 2010, *ApJ*, **724**, 559
- Leung G. C. K., et al., 2019, *ApJ*, **886**, 11

- Lewis G. M., Austin P. H., 2002, 11th Conference on Atmospheric Radiation, ed. G. H. Smith & J. P. Brodie, American Meteorological Society Conference Series, 123
- Liska M., Tchekhovskoy A., Ingram A., van der Klis M., 2019, *MNRAS*, **487**, 550
- Liska M., Tchekhovskoy A., Quataert E., 2020, *MNRAS*, **494**, 3656
- Liska M. T. P., et al., 2022, *ApJS*, **263**, 26
- Liu G., Zakamska N. L., Greene J. E., Nesvadba N. P. H., Liu X., 2013, *MNRAS*, **436**, 2576
- Liu H.-Y., Liu W.-J., Dong X.-B., Zhou H., Wang T., Lu H., Yuan W., 2019, *ApJS*, **243**, 21
- Longinotti A. L., et al., 2013, *ApJ*, **766**, 104
- Longinotti A. L., et al., 2023, *MNRAS*, **521**, 2134
- Luo R., et al., 2021, *ApJ*, **908**, 221
- Lupi A., Quadri G., Volonteri M., Colpi M., Regan J. A., 2024, *A&A*, **686**, A256
- Lutz D., et al., 2010, *ApJ*, **712**, 1287
- Lutz D., et al., 2020, *A&A*, **633**, A134
- Madau P., Dickinson M., 2014, *ARA&A*, **52**, 415
- Madau P., Haardt F., Dotti M., 2014, *ApJ*, **784**, L38
- Mainieri V., et al., 2011, *A&A*, **535**, A80
- Maiolino R., et al., 2017, *Nature*, **544**, 202
- Man A., Belli S., 2018, *Nature Astronomy*, **2**, 695
- Mandal A., Mukherjee D., Federrath C., Nesvadba N. P. H., Bicknell G. V., Wagner A. Y., Meenakshi M., 2021, *MNRAS*, **508**, 4738
- Mandal A., Mukherjee D., Federrath C., Bicknell G. V., Nesvadba N. P. H., Mignone A., 2024, *MNRAS*, **531**, 2079
- Mannerkoski M., Rawlings A., Johansson P. H., Naab T., Rantala A., Springel V., Irodotou D., Liao S., 2023, *MNRAS*, **524**, 4062
- Marasco A., et al., 2020, *A&A*, **644**, A15
- Marconcini C., et al., 2023, *A&A*, **677**, A58
- Marconi A., Risaliti G., Gilli R., Hunt L. K., Maiolino R., Salvati M., 2004, *MNRAS*, **351**, 169
- Marinacci F., et al., 2018, *MNRAS*, **480**, 5113
- Martig M., Bournaud F., Teyssier R., Dekel A., 2009, *ApJ*, **707**, 250
- Martín-Navarro I., Pillepich A., Nelson D., Rodriguez-Gomez V., Donnari M., Hernquist L., Springel V., 2021, *Nature*, **594**, 187
- Matzeu G. A., et al., 2023, *A&A*, **670**, A182
- McAlpine S., et al., 2016, *Astronomy and Computing*, **15**, 72
- McAlpine S., Bower R. G., Harrison C. M., Crain R. A., Schaller M., Schaye J., Theuns T., 2017, *MNRAS*, **468**, 3395
- McCarthy I. G., Schaye J., Bower R. G., Ponman T. J., Booth C. M., Vecchia C. D., Springel V., 2011, *MNRAS*, **412**, 1965
- McCourt M., Oh S. P., O’Leary R., Madigan A.-M., 2018, *MNRAS*, **473**, 5407
- McKinney J. C., 2006, *MNRAS*, **368**, 1561
- McLure R. J., Dunlop J. S., 2004, *MNRAS*, **352**, 1390
- Meena B., et al., 2023, *ApJ*, **943**, 98

- Meenakshi M., Mukherjee D., Wagner A. Y., Nesvadba N. P. H., Morganti R., Janssen R. M. J., Bicknell G. V., 2022a, *MNRAS*, 511, 1622
- Meenakshi M., et al., 2022b, *MNRAS*, 516, 766
- Mercedes-Feliz J., et al., 2023, *MNRAS*, 524, 3446
- Mihos J. C., Hernquist L., 1996, *ApJ*, 464, 641
- Mizumoto M., Nomura M., Done C., Ohsuga K., Odaka H., 2021, *MNRAS*, 503, 1442
- Molyneux S. J., Harrison C. M., Jarvis M. E., 2019, *A&A*, 631, A132
- Morganti R., Tadhunter C. N., Oosterloo T. A., 2005, *A&A*, 444, L9
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Mukherjee D., Bicknell G. V., Sutherland R., Wagner A., 2016, *MNRAS*, 461, 967
- Mukherjee D., Bicknell G. V., Wagner A. Y., Sutherland R. S., Silk J., 2018, *MNRAS*, 479, 5544
- Mullaney J. R., et al., 2012, *ApJ*, 753, L30
- Murray N., Chiang J., Grossman S. A., Voit G. M., 1995, *ApJ*, 451, 498
- Musiimenta B., et al., 2023, *A&A*, 679, A84
- Naiman J. P., et al., 2018, *MNRAS*, 477, 1206
- Narayan R., Yi I., 1994, *ApJ*, 428, L13
- Nelson D., et al., 2018, *MNRAS*, 475, 624
- Nelson D., et al., 2019a, *Computational Astrophysics and Cosmology*, 6, 2
- Nelson D., et al., 2019b, *MNRAS*, 490, 3234
- Nelson D., et al., 2020, *MNRAS*, 498, 2391
- Ni Y., et al., 2023, *ApJ*, 959, 136
- Nims J., Quataert E., Faucher-Giguère C.-A., 2015, *MNRAS*, 447, 3612
- Noeske K. G., et al., 2007, *ApJ*, 660, L43
- Novak G. S., Ostriker J. P., Ciotti L., 2011, *ApJ*, 737, 26
- Novak G. S., Ostriker J. P., Ciotti L., 2012, *MNRAS*, 427, 2734
- O’Dea C. P., Payne H. E., Kocevski D., 1998, *AJ*, 116, 623
- Padovani P., Bonzini M., Kellermann K. I., Miller N., Mainieri V., Tozzi P., 2015, *MNRAS*, 452, 1263
- Padovani P., et al., 2017, *A&A Rev.*, 25, 2
- Page M. J., et al., 2012, *Nature*, 485, 213
- Pakmor R., Springel V., 2013, *MNRAS*, 432, 176
- Pakmor R., Springel V., Bauer A., Mocz P., Munoz D. J., Ohlmann S. T., Schaal K., Zhu C., 2016, *MNRAS*, 455, 1134
- Panessa F., Baldi R. D., Laor A., Padovani P., Behar E., McHardy I., 2019, *Nature Astronomy*, 3, 387
- Peebles P. J. E., 1972, *ApJ*, 178, 371
- Pereira-Santaella M., et al., 2011, *A&A*, 535, A93
- Perna M., et al., 2018, *A&A*, 619, 90
- Peterson J. R., et al., 2001, *A&A*, 365, L104
- Peterson J. R., Kahn S. M., Paerels F. B. S., Kaastra J. S., Tamura T., Bleeker J. A. M., Ferrigno C., Jernigan J. G., 2003, *ApJ*, 590, 207

- Petley J. W., et al., 2022, *MNRAS*, 515, 5159
- Pfrommer C., Pakmor R., Schaal K., Simpson C. M., Springel V., 2017, *MNRAS*, 465, 4500
- Pfrommer C., Werhahn M., Pakmor R., Girichidis P., Simpson C. M., 2022, *MNRAS*, 515, 4229
- Pillepich A., et al., 2017, *MNRAS*, 473, 4077
- Pillepich A., et al., 2018, *MNRAS*, 475, 648
- Pillepich A., et al., 2019, *MNRAS*, 490, 3196
- Pillepich A., Nelson D., Truong N., Weinberger R., Martin-Navarro I., Springel V., Faber S. M., Hernquist L., 2021, *MNRAS*, 508, 4667
- Piotrowska J. M., Bluck A. F. L., Maiolino R., Peng Y., 2022, *MNRAS*, 512, 1052
- Planck Collaboration 2016, *A&A*, 594, A13
- Planck Collaboration 2020, *A&A*, 641, A6
- Ponti G., et al., 2019, *Nature*, 567, 347
- Ponti G., Morris M. R., Churazov E., Heywood I., Fender R. P., 2021, *A&A*, 646, A66
- Popping G., et al., 2019, *ApJ*, 882, 137
- Pounds K. A., Reeves J. N., King A. R., Page K. L., O'Brien P. T., Turner M. J. L., 2003a, *MNRAS*, 345, 705
- Pounds K. A., King A. R., Page K. L., O'Brien P. T., 2003b, *MNRAS*, 346, 1025
- Powell M. C., et al., 2018, *ApJ*, 858, 110
- Predehl P., et al., 2020, *Nature*, 588, 227
- Proga D., Stone J. M., Kallman T. R., 2000, *ApJ*, 543, 686
- Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, *MNRAS*, 430, 2427
- Ramasawmy J., Stevens J., Martin G., Geach J. E., 2019, *MNRAS*, 486, 4320
- Ramesh R., Nelson D., 2024, *MNRAS*, 528, 3320
- Ramesh R., Nelson D., Fielding D., Brüggén M., 2024a, *arXiv e-prints*, p. arXiv:2407.00172
- Ramesh R., Nelson D., Fielding D., Brüggén M., 2024b, *A&A*, 684, L16
- Ramos Almeida C., et al., 2022, *A&A*, 658, A155
- Read J. I., Trentham N., 2005, *Philosophical Transactions of the Royal Society of London Series A*, 363, 2693
- Rees M. J., Ostriker J. P., 1977, *MNRAS*, 179, 541
- Rennehan D., Babul A., Moa B., Davé R., 2024, *MNRAS*, 532, 4793
- Renzini A., Peng Y. J., 2015, *Astrophysical Journal Letters*, 801, L29
- Revalski M., et al., 2018, *ApJ*, 867, 88
- Revalski M., et al., 2021, *ApJ*, 910, 139
- Richings A. J., Faucher-Giguère C.-A., 2018a, *MNRAS*, 474, 3673
- Richings A. J., Faucher-Giguère C.-A., 2018b, *MNRAS*, 478, 3100
- Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349
- Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780
- Richings A. J., Faucher-Giguère C.-A., Stern J., 2021, *MNRAS*, 503, 1568
- Riffel R. A., et al., 2023, *MNRAS*, 521, 1832
- Robertson B., et al., 2024, *ApJ*, 970, 31

- Rosario D. J., et al., 2012, *A&A*, 545, 45
- Rosario D. J., et al., 2013, *ApJ*, 763, 59
- Rosario D. J., et al., 2018, *MNRAS*, 473, 5658
- Rosas-Guevara Y., Bower R. G., Schaye J., McAlpine S., Vecchia C. D., Frenk C. S., Schaller M., Theuns T., 2016, *MNRAS*, 462, 190
- Rosdahl J., Teyssier R., 2015, *MNRAS*, 449, 4380
- Rosdahl J., Blaizot J., Aubert D., Stranex T., Teyssier R., 2013, *MNRAS*, 436, 2188
- Rosdahl J., Schaye J., Dubois Y., Kimm T., Teyssier R., 2017, *MNRAS*, 466, 11
- Rose M., Tadhunter C., Ramos Almeida C., Rodríguez Zaurín J., Santoro F., Spence R., 2018, *MNRAS*, 474, 128
- Roy N., et al., 2021, *ApJ*, 913, 33
- Rupke D. S. N., Veilleux S., 2013, *ApJ*, 768, 75
- Rupke D. S., Veilleux S., Sanders D. B., 2005, *ApJS*, 160, 115
- Rupke D. S. N., Gültekin K., Veilleux S., 2017, *ApJ*, 850, 40
- Ruszkowski M., Pfrommer C., 2023, *A&A Rev.*, 31, 4
- Salvestrini F., et al., 2022, *A&A*, 663, A28
- Schaal K., Springel V., 2015, *MNRAS*, 446, 3992
- Schaller M., et al., 2024, *MNRAS*, 530, 2378
- Schawinski K., Treister E., Urry C. M., Cardamone C. N., Simmons B., Yi S. K., 2011, *ApJ*, 727, L31
- Schawinski K., et al., 2014, *MNRAS*, 440, 889
- Schawinski K., Koss M., Berney S., Sartori L. F., 2015, *MNRAS*, 451, 2517
- Schaye J., et al., 2010, *MNRAS*, 402, 1536
- Schaye J., et al., 2015, *MNRAS*, 446, 521
- Schaye J., et al., 2023, *MNRAS*, 526, 4978
- Schneider E. E., Robertson B. E., 2017, *ApJ*, 834, 144
- Schneider E. E., Robertson B. E., Thompson T. A., 2018, *ApJ*, 862, 56
- Scholtz J., et al., 2018, *MNRAS*, 475, 1288
- Scholtz J., et al., 2020, *MNRAS*, 492, 3194
- Scholtz J., et al., 2021, *MNRAS*, 505, 5469
- Schreiber C., et al., 2015, *A&A*, 575
- Schulze A., et al., 2019, *MNRAS*, 488, 1180
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, 500, 33
- Shangguan J., Ho L. C., 2019, *ApJ*, 873, 90
- Shangguan J., Ho L. C., Xie Y., 2018, *ApJ*, 854, 158
- Shangguan J., Ho L. C., Bauer F. E., Wang R., Treister E., 2020a, *ApJS*, 247, 15
- Shangguan J., Ho L. C., Bauer F. E., Wang R., Treister E., 2020b, *ApJ*, 899, 112
- Sharma R. S., Brooks A. M., Tremmel M., Bellovary J., Quinn T. R., 2023, *ApJ*, 957, 16
- Shimizu T. T., Mushotzky R. F., Meléndez M., Koss M. J., Barger A. J., Cowie L. L., 2017, *MNRAS*, 466, 3161

- Sijacki D., Springel V., 2006, *MNRAS*, **366**, 397
- Sijacki D., Springel V., Di Matteo T., Hernquist L., 2007, *MNRAS*, **380**, 877
- Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, **452**, 575
- Silk J., Rees M. J., 1998, *Astron. Astrophys*, **331**, 1
- Sirressi M., et al., 2019, *MNRAS*, **489**, 1927
- Sądowski A., 2009, *ApJS*, **183**, 171
- Sądowski A., Narayan R., Tchekhovskoy A., Abarca D., Zhu Y., McKinney J. C., 2015, *MNRAS*, **447**, 49
- Smirnova-Pinchukova I., et al., 2022, *A&A*, **659**, A125
- Smith M. C., et al., 2024, *MNRAS*, **527**, 1216
- Soltan A., 1982, *MNRAS*, **200**, 115
- Somalwar J., Johnson S. D., Stern J., Goulding A. D., Greene J. E., Zakamska N. L., Alexandroff R. M., Chen H.-W., 2020, *ApJ*, **890**, L28
- Somerville R. S., Davé R., 2015, *ARA&A*, **53**, 51
- Somerville R. S., Primack J. R., 1999, *MNRAS*, **310**, 1087
- Somerville R. S., Hopkins P. F., Cox T. J., Robertson B. E., Hernquist L., 2008, *MNRAS*, **391**, 481
- Speagle J. S., Steinhardt C. L., Capak P. L., Silverman J. D., 2014, *ApJS*, **214**, 15
- Speranza G., et al., 2024, *A&A*, **681**, A63
- Springel V., 2005, *MNRAS*, **364**, 1105
- Springel V., 2010a, *ARA&A*, **48**, 391
- Springel V., 2010b, *MNRAS*, **401**, 791
- Springel V., Hernquist L., 2003, *MNRAS*, **339**, 289
- Springel V., Matteo T. D., Hernquist L., 2005a, *MNRAS*, **361**, 776
- Springel V., et al., 2005b, *Nature*, **435**, 629
- Springel V., et al., 2018, *MNRAS*, **475**, 676
- Stachenfeld K., et al., 2021, *Conference Paper, International Conference on Learning Representations*, p. [arXiv:2112.15275](https://arxiv.org/abs/2112.15275)
- Stanley F., Harrison C. M., Alexander D. M., Swinbank A. M., Aird J. A., Moro A. D., Hickox R. C., Mullaney J. R., 2015, *MNRAS*, **453**, 591
- Stanley F., et al., 2017, *MNRAS*, **472**, 2221
- Starkenburger T. K., Tonnesen S., Kopenhafer C., 2019, *ApJ*, **874**, L17
- Stevens I. R., Kallman T. R., 1990, *ApJ*, **365**, 321
- Sturm E., et al., 2011, *ApJ*, **733**, L16
- Su M., Slatyer T. R., Finkbeiner D. P., 2010, *ApJ*, **724**, 1044
- Sutherland R. S., Bicknell G. V., 2007, *ApJS*, **173**, 37
- Tacconi L. J., et al., 2018, *ApJ*, **853**, 179
- Tacconi L. J., Genzel R., Sternberg A., 2020, *ARA&A*, **58**, 157
- Tadhunter C., et al., 2018, *MNRAS*, **478**, 1558
- Talbot R. Y., Bourne M. A., Sijacki D., 2021, *MNRAS*, **504**, 3619
- Talbot R. Y., Sijacki D., Bourne M. A., 2022, *MNRAS*, **514**, 4535



- Talbot R. Y., Sijacki D., Bourne M. A., 2024, *MNRAS*, **528**, 5432
- Tan B., Fielding D. B., 2024, *MNRAS*, **527**, 9683
- Tanner R., Weaver K. A., 2022, *AJ*, **163**, 134
- Terrazas B. A., Bell E. F., Woo J., Henriques B. M. B., 2017, *ApJ*, 844, 170
- Terrazas B. A., et al., 2020, *MNRAS*, 493, 1888
- Teyssier R., 2002, *A&A*, **385**, 337
- Thacker R. J., MacMackin C., Wurster J., Hobbs A., 2014, *MNRAS*, **443**, 1125
- Thomas N., Davé R., Anglés-Alcázar D., Jarvis M., 2019, *MNRAS*, 487, 5764
- Thompson T. A., Fabian A. C., Quataert E., Murray N., 2015, *MNRAS*, **449**, 147
- Thorne K. S., 1974, *ApJ*, **191**, 507
- Tombesi F., Cappi M., Reeves J. N., Palumbo G. G. C., Braitto V., Dadina M., 2011, *ApJ*, **742**, 44
- Tombesi F., Cappi M., Reeves J. N., Braitto V., 2012, *MNRAS*, **422**, L1
- Tombesi F., Cappi M., Reeves J. N., Nemmen R. S., Braitto V., Gaspari M., Reynolds C. S., 2013, *MNRAS*, **430**, 1102
- Tombesi F., Meléndez M., Veilleux S., Reeves J. N., González-Alfonso E., Reynolds C. S., 2015, *Nature*, **519**, 436
- Toomre A., Toomre J., 1972, *ApJ*, **178**, 623
- Trayford J. W., Theuns T., Bower R. G., Crain R. A., del P. Lagos C., Schaller M., Schaye J., 2016, *MNRAS*, 460, 3925
- Treister E., Natarajan P., Sanders D. B., Urry C. M., Schawinski K., Kartaltepe J., 2010, *Science*, **328**, 600
- Trump J. R., et al., 2015, *ApJ*, **811**, 26
- Tumlinson J., Peebles M. S., Werk J. K., 2017, *ARA&A*, **55**, 389
- van der Vlugt D., Costa T., 2019, *MNRAS*, **490**, 4918
- van de Voort F., Springel V., Mandelker N., van den Bosch F. C., Pakmor R., 2019, *MNRAS*, **482**, L85
- Valentino F., et al., 2021, *A&A*, **654**, A165
- Vayner A., Zakamska N., Wright S. A., Armus L., Murray N., Walth G., 2021, *ApJ*, **923**, 59
- Veilleux S., Teng S. H., Rupke D. S. N., Maiolino R., Sturm E., 2014, *ApJ*, **790**, 116
- Veilleux S., Bolatto A., Tombesi F., Meléndez M., Sturm E., González-Alfonso E., Fischer J., Rupke D. S. N., 2017, *ApJ*, **843**, 18
- Veilleux S., Maiolino R., Bolatto A. D., Aalto S., 2020, *A&A Rev.*, **28**, 2
- Venturi G., et al., 2021, *A&A*, **648**, A17
- Verwilghen P., et al., 2024, *A&A*, **687**, A53
- Vietri G., et al., 2022, *A&A*, **659**, A129
- Villaescusa-Navarro F., et al., 2021, *ApJ*, **915**, 71
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Vogelsberger M., et al., 2014a, *MNRAS*, **444**, 1518
- Vogelsberger M., et al., 2014b, *Nature*, 509, 177–182
- Voit G. M., Oppenheimer B. D., Bell E. F., Terrazas B., Donahue M., 2024, *ApJ*, **960**, 28

- Wagner A. Y., Bicknell G. V., 2011, *The Astrophysical Journal*, 728, 29
- Wagner A. Y., Bicknell G. V., Umemura M., 2012, *ApJ*, 757, 136
- Wagner A. Y., Umemura M., Bicknell G. V., 2013, *ApJ*, 763, L18
- Wang J., Fabbiano G., Risaliti G., Elvis M., Mundell C. G., Dumas G., Schinnerer E., Zezas A., 2010, *ApJ*, 719, L208
- Wang J., et al., 2014, *ApJ*, 781, 55
- Wang W., et al., 2024a, *A&A*, 683, A169
- Wang C., Wang J., Dadina M., Fabbiano G., Elvis M., Bianchi S., Guainazzi M., 2024b, *ApJ*, 962, 188
- Ward S. R., Harrison C. M., Costa T., Mainieri V., 2022, *MNRAS*, 514, 2936
- Ward S. R., Costa T., Harrison C. M., Mainieri V., 2024, *MNRAS*, stae1816
- Weinberger R., Hernquist L., 2023, *MNRAS*, 519, 3011
- Weinberger R., et al., 2017, *MNRAS*, 465, 3291
- Weinberger R., et al., 2018, *MNRAS*, 479, 4056
- Wellons S., et al., 2023, *MNRAS*, 520, 5394
- Werhahn M., Pfrommer C., Girichidis P., Winner G., 2021a, *MNRAS*, 505, 3295
- Werhahn M., Pfrommer C., Girichidis P., 2021b, *MNRAS*, 508, 4072
- Whitaker K. E., et al., 2014, *ApJ*, 795, 104
- White S. D. M., Frenk C. S., 1991, *ApJ*, 379, 52
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- Wild V., Almaini O., Dunlop J., Simpson C., Rowlands K., Bowler R., Maltby D., Mclure R., 2016, *MNRAS*, 463, 832
- Willis A. G., Strom R. G., Wilson A. S., 1974, *Nature*, 250, 625
- Winner G., Pfrommer C., Girichidis P., Pakmor R., 2019, *MNRAS*, 488, 2235
- Wu P.-F., et al., 2018, *ApJ*, 868, 37
- Wylezalek D., Zakamska N. L., 2016, *MNRAS*, 461, 3724
- Xie Y., Ho L. C., Zhuang M.-Y., Shangguan J., 2021, *ApJ*, 910, 124
- Yang H., Yuan F., Yuan Y.-F., White C. J., 2021, *ApJ*, 914, 131
- Yesuf H. M., Ho L. C., 2019, *ApJ*, 884, 177
- Yirak K., Frank A., Cunningham A. J., 2010, *ApJ*, 722, 412
- York D. G., et al., 2000, *ApJ*, 120, 1579
- Yu Q., Tremaine S., 2002, *MNRAS*, 335, 965
- Yuan F., Narayan R., 2014, *ARA&A*, 52, 529
- Yuan F., Gan Z., Narayan R., Sadowski A., Bu D., Bai X.-N., 2015, *ApJ*, 804, 101
- Zakamska N. L., Greene J. E., 2014, *MNRAS*, 442, 784
- Zanella A., et al., 2018, *MNRAS*, 481, 1976
- Zhang R., Guo F., 2020, *ApJ*, 894, 117
- Zhang D., Thompson T. A., Quataert E., Murray N., 2017, *MNRAS*, 468, 4801
- Zhuang M.-Y., Ho L. C., 2019, *ApJ*, 882, 89
- Zhuang M.-Y., Ho L. C., Shangguan J., 2019, *ApJ*, 873, 103

- Zhuang M.-Y., Ho L. C., Shanguan J., 2021, *ApJ*, 906, 38
- Zinger E., et al., 2020, *MNRAS*, 499, 768
- Zolotov A., et al., 2015, *MNRAS*, 450, 2327
- ZuHone J. A., Hallman E. J., 2016, pyXSIM: Synthetic X-ray observations generator, Astrophysics Source Code Library, record ascl:1608.002
- ZuHone J. A., Vikhlinin A., Tremblay G. R., Randall S. W., Andrade-Santos F., Bourdin H., 2023, SOXS: Simulated Observations of X-ray Sources, Astrophysics Source Code Library, record ascl:2301.024
- Zubovas K., Bourne M. A., 2017, *MNRAS*, 468, 4956
- Zubovas K., King A., 2012, *ApJ*, 745, L34
- Zubovas K., Nardini E., 2020, *MNRAS*, 498, 3633
- Zubovas K., King A. R., Nayakshin S., 2011, *MNRAS*, 415, L21



# Publication List

During the duration of the doctoral degree, the author has undertaken the following work as the primary author which also inspired the thesis Chapters 3, 4 & 5:

- **Ward S. R.**, Harrison C. M., Costa T., Mainieri V. (2022), *Cosmological simulations predict that AGN preferentially live in gas-rich, star-forming galaxies despite effective feedback*, MNRAS, 514, 2936
- **Ward S. R.**, Costa T., Harrison C. M., Mainieri V. (2024), *AGN-driven outflows in clumpy media: multiphase structure and scaling relations*, MNRAS, stae1816
- **Ward S. R.**, Costa T., Harrison C. M., Mainieri V. (in prep.), *Mixing between AGN winds and ISM clouds produces luminous X-rays*.

In addition to the studies presented in this thesis, the author has also been involved in the following works:

- Almeida I., Costa T., Harrison C. M., **Ward S. R.** (in prep.), *The impact of AGN outflows on cold cloud formation in galactic simulations*
- Frias Castillo M., et al. [incl. **Ward S. R.**] (2024), *At the end of cosmic noon: Short gas depletion times in unobscured quasars at  $z \sim 1$* , A&A, 683, A211
- Bertola E., et al. [incl. **Ward S. R.**] (2024), *KASHz+SUPER: Evidence for cold molecular gas depletion in AGN hosts at cosmic noon*, accepted A&A
- Molyneux S., et al. [incl. **Ward S. R.**] (2024), *The Quasar Feedback Survey: characterising CO excitation in quasar host galaxies*, MNRAS, 527, 4420
- Girdhar A., et al. [incl. **Ward S. R.**] (2024), *Quasar Feedback Survey: molecular gas around radio lobes and in central outflows reveals dual feedback effects in ‘radio quiet’ quasars*, MNRAS, 527, 9322
- Harrison C. M., Girdhar A., **Ward S. R.** (2023), *Establishing the impact of luminous AGN with multi-wavelength observations and simulations*, IAU Symposium 378 Conference Proceedings
- Girdhar A., et al. [incl. **Ward S. R.**] (2022), *Quasar feedback survey: multiphase outflows, turbulence, and evidence for feedback caused by low power radio jets inclined into the galaxy disc*, MNRAS, 512, 1608
- Jarvis M. E., et al. [incl. **Ward S. R.**] (2021), *The quasar feedback survey: discovering hidden Radio-AGN and their connection to the host galaxy ionized gas*, MNRAS, 503, 1780

*You taught me the courage of stars before you left.  
How light carries on endlessly, even after death.  
With shortness of breath, you explained the infinite.  
How rare and beautiful it is to even exist.*

*I couldn't help but ask for you to say it all again.  
I tried to write it down, but I could never find a pen.  
I'd give anything to hear you say it one more time,  
that the universe was made just to be seen by my eyes.*

*–Saturn, Sleeping at Last*

# Acknowledgements

There aren't words enough to express my thanks to everyone who has helped and supported me on this PhD journey.

Firstly, to my family: Mum & Dad, I'm blaming my career choice on you thanks to all the 'educational' trips to the Greenwich Observatory and the Science Museum. Thank you for putting up with my probably incomprehensible physics lectures and the customs-defying care packages. To Lydia for always being around for a hug, even if mostly virtual. For Lego Star Wars & long discussions on how Andor is the best series ever made.

To my fantastic supervisors: Tiago, for your patience in answering my many questions. For setting up a numerical journal club (twice!), creating the scientific environment to dive deep into the weeds of the simulations. For setting me off on this exciting computational journey. To Vincenzo for always being a calming and helpful presence. For enabling and supporting my wild adventures, like securing me the funding to go to Chile, and then being so respectful of my weekend you didn't even tell me I'd got the grant! To Chris for always telling me to see the big picture – both scientifically, but also when it came to general advice about academia and the world beyond. For bringing me to the group at Newcastle and embodying what it means to be a leader who cares deeply about everyone in their group. I know finishing the PhD is the time to fly the nest, but I'm still looking forward to the fun science we can keep doing together (but I'll let you have a bit of a break from my Slack notifications for a bit!)

I have been incredibly lucky to spend three years of my PhD at ESO. This was truly a special place to stretch my scientific wings and learn to fly – thank you to the 'Old Folks' for setting the supportive and encouraging atmosphere that gave me the confidence to enjoy my science.

To Nelma, truly the glue who held it all together. For the kindness you showed to all the students, for always being on our side, and for the lockdown cat pictures that made opening Teams a little more bearable.

To my *didi* Aish for always looking out for me and for blazing a trail through the bureaucracy for my move to Newcastle. You have an infectious positivity that always made me smile no matter what I was going through.

To Tere, technically my 'first friend' at ESO after we started chatting when we missed a bus together on the interview day. The US isn't ready for our karaoke duet skills. Remember me, the cameraman who made you a star, when you receive your lifetime award for services to Tiktok.

To my absolutely fantastic office mates: To Keegan (you were my favourite, don't tell Tommy), my fellow goth. From board games to volleyball to hiking, I could always trust you had my back and your kindness and consideration were sorely missed when you left (as well as the valuable role of the Python Information Support Service). To Tommy (you were my favourite, don't tell Keegan) for introducing me to the greatest football club in the world (SC Freiburg) and getting me tickets to their greatest cup victory over Bayern

Munich. I would thank you for all the fun we had playing FIFA or shooting nerf guns in the office, but my supervisors might read this so probably better if they don't find out about that. . .

To Stephen, Mr ESO himself; table football legend, the most patriotically Wigan man alive, employed by the Lathums as their key promoter. You showed me just how fun science at ESO could be and I know you'll always be there when I need to express just how much I love Chicken Run.

To Jianhang for lockdown bike rides; Avinash for keeping my plants alive; Justyn (I still have the Band Camino on repeat); Marta for putting up the poster that launched our choir directing careers; Simon for being a legendary napper and the best chauffeur in all of Oz (don't worry, I'll pay you back at our UK road trip).

To the board gamers: Pierrick (thanks for being a fellow member of simulation gang, I'm so glad I got to see the Kevin character arc); Tadeja & Smaran (what jhumka?); *that guy* David (finally a worthy opponent: our battle will be legendary!); and Joshiwa for being the only person who makes me feel small. Extra thanks to you this year Joshi for looking after something very precious to me. I'm looking to many more online board games with all of you, but please stop ganging up on me – I'm telling you it's *David* who's the real enemy!

To the stars of the ESO volleyball team: Cami for being the smallest yet loudest member; our coach Linus; our supernova Paul; and to Malte for barking at everything for some reason. To dear friends Alice, Andrea, Luca, Gemma, Domi, Giulia, Román, Tuts, Ale, Juliet & Anita.

To all the other Lost Souls of ESO who I don't have space to mention, but bravely followed me up mountains, came to volleyball training, joined the choir, voted in countless polls I sent out, or wore a Christmas jumper – thank you. My crazy ideas stay in my head if I don't have a community to share them with and you are such a welcoming, adventurous and fun-loving group of truly wonderful people and it was an honour to be part of it. I wouldn't have made it this far without you.

Moving to a new institute for the last year of my PhD could've been tough, but I want to thank everyone in AstroObs for making me feel so welcome. To Vicky, the intrepid Space Investigator; Houda, for the coffee & chats; Beth & Carola, for all the Big Bang breakfasts; Rose (good luck sonifying my figures!); Ivan (I guess you're my ACDC bandmate now?); Devang, Karthika & Julie, for all the India travel tips; and to Henry, Connor and Yasin (who's the nicest guy I've ever met who could definitely beat me in a fight). A huge thank you to my Newcastle flatmates, Will and Alex, for taking a chance on a random stranger online and welcoming me into your home. I hope I've sufficiently paid back your generosity in board game nights.

To old friends: Jamie, here's to many more trips to random places (Austin too now apparently!?). To Mark and Josh, for many calls and hugs and chats. You are the most beautiful friends anyone could ask for.

And finally to Ivanna, my unintended joy. From the Atacama to Australia, from Picardie to Pondicherry, no matter what crazy adventure we're on, you always feel like home. Having you alongside me in the darkest days kept me going – thank you for all the zoom calls, the games of Wingspan, all the quiches and bubble teas, and trips to the bakery. And look – *we've done it!* Here's to the next crazy chapter.

*Chapter quotes were taken from the album Atlas: Space, by the artist Sleeping at Last – the album I listened to on my first flight to Munich in August 2020.*





