Magnetic Field Evolution in Black Hole Accretion Disks and Relativistic Jets



Payton Elyce Rodman

Supervisor: Prof. Christopher S. Reynolds

Institute of Astronomy
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Churchill College July 2024

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted, or, is being concurrently submitted, for any degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

The following Chapters contain material which has been published, or are in preparation for publication, in peer-reviewed journals:

- Chapter 2 is a reproduction of

  Evolution of the Magnetic Field in High- and Low-β Disks with Initially Toroidal Fields, **Rodman, P. E.**, Reynolds, C. S., 2024, ApJ, 960, 97

  with minor stylistic changes to maintain consistency in thesis dialect and structure, the addition of a statement on units at the end of Section 2.1, and the addition of an orbital time axis to all relevant plots.

- Chapter 3 presents previously unpublished analysis of the simulations of Chapter 2, exploring the connection to mean-field dynamo theory.

- Chapter 4 is based on ongoing work in collaboration with C. S. Reynolds and M. G. H. Krause.

This dissertation contains fewer than 60,000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Payton Elyce Rodman
July 2024

Magnetic Field Evolution in Black Hole Accretion Disks and Relativistic Jets

Payton Elyce Rodman

Supermassive black holes (SMBHs) catalyse change in their host galaxies, mediated through the inflow and subsequent outflow of material through the massive body. While an SMBH does not—indeed, *cannot*—hold its own magnetic field, the material it interacts with, feeds upon, and expels outwards is naturally magnetic, and this embedded magnetic field is expected to have a complex effect upon these processes. Due to the inherent difficulties in studying magnetic fields there is still much to learn about their role around SMBHs generally, and in accretion and the launching of relativistic jets in particular. This thesis aims to tackle both topics using three-dimensional magnetohydrodynamics (MHD) and relativistic magnetohydrodynamics (RMHD) simulations.

In the first half (Chapters 2 and 3), I focus on the role of magnetic fields in simulations of thick accretion disks. At stake is a core question within the field of disk simulations: can a simulated accretion disk self-generate the strong, large-scale poloidal field needed to launch relativistic jets? To assess this question, I simulate disks with initially toroidal fields of both high (plasma beta $\beta = P_{\rm gas}/P_{\rm mag} = 5$) and low ($\beta = 200$) field strength and at high and moderate resolution, using the code *Athena++*. I find that both weak and strong initially toroidal field disks can generate a poloidal field, with the initially strong case producing a more ordered and large-scale field. These findings are tempered by caution on the resolution, as an $m = 1$ mode overdensity is created at the crossroads of a weak field, moderate resolution, and resolution step-change. I additionally perform a minor study of the disk dynamo, covered in Chapter 3. Using mean-field dynamo theory, I calculate the diagonal, symmetric α_{ij} components under both a global azimuthal average and a local Gaussian filter average, finding that local averaging returns more robust numerical signals. We retrieve a northern hemisphere $\alpha_{\phi\phi}$ sign that is consistent with previous works and is in line with expectations of an $\alpha\Omega$-type dynamo. The mean-field formalism returns a better fit later in the disk evolution, suggesting that additional dynamo factors are present in the early phases of field evolution and growth.

In the second half of the thesis, in Chapter 4, I move my focus to relativistic jets and investigate whether magnetised head-tail (HT) jets launched in a realistic turbulent magnetised background can form magnetic linkages between the lobes as have been observed in recent high-resolution radio images. To assess this, I simulate a suite of low density

(density ratio $\eta = \rho_{\rm jet}/\rho_{\rm amb} = 10^{-3}$), pressure-matched jets at both low ($v_j = 0.24c$) and high ($0.94c$) velocity into pre-evolved turbulent backgrounds generated using the *PLUTO* code and subject them to strong, transonic crosswinds. I find that magnetic linkages do occur, but are between the jet lobes and the ambient medium, with the region between lobes being too turbulent to support long filaments. I additionally find that the location and direction of the existing filaments depend on the particular turbulent background realisation, implying that observations of filaments could be used to constrain the ICM field.

Acknowledgements

First and foremost, I would like to express my incredible gratitude to my supervisor, Chris Reynolds. These last five years have had more than enough challenges and hurdles, and yet, despite a pandemic and changing jobs, you always managed to make time to discuss simulations with me.

Secondly, to Martin Krause, my colleague and honorary jets co-supervisor, without whose guidance the final chapter of this thesis would surely be half as long and a quarter as good.

Thanks must also go to the wider THEAPA and X-ray groups, past and present, for all the insightful discussions, to Paul Hewett for being my department contact over the last year and answering my many trivial questions, and to the Gates Cambridge Foundation for providing the funding for this work.

I would not be where I am today without the support of my family—thank you for always believing in me and, in particular, for hosting me during the pandemic. Special thanks go to my brother, Saxon, for designing the chapter artwork for this thesis

And finally, to Elie. Words cannot express my love and gratitude. You have supported me in every way, big and small, on this journey, and that this thesis is finished at all is in no small part due to you. I don't know that I can ever repay you, but I will always try.

To Bruce Rodman.

Table of contents

List of figures **xv**

List of tables **xix**

Nomenclature **xxi**

1 Introduction **3**

1.1 Introduction to Active Galactic Nuclei 3

1.1.1 Active Galactic Nuclei feedback 6

1.2 Accretion disks 12

1.2.1 Angular momentum transport in hydrodynamics 14

1.2.2 Adding magnetic fields to the mix 16

1.2.3 The Magnetorotational Instability 18

1.2.4 Summary 21

1.3 Characterising the large-scale dynamo 23

1.3.1 Mean-field theory 24

1.3.2 First-Order Smoothing Approximation 26

1.3.3 Summary 36

1.4 Black hole jets 36

1.4.1 Jet morphology 40

1.4.2 Jet magnetic field structure 45

1.4.3 Radio filaments and threads 48

1.4.4 Summary 51

1.5 Numerical Challenges 51

1.6 Thesis outline 54

2 Evolution of the magnetic field in high- and low-β disks with initially toroidal fields **57**
2.1 Introduction . . . 57
2.2 Numerical model . . . 62
2.2.1 Foundational equations . . . 62
2.2.2 Mesh, boundary, and initial conditions . . . 63
2.2.3 Diagnostics and convergence . . . 65
2.3 Results . . . 70
2.3.1 Disk structure and evolution . . . 70
2.3.2 Magnetic field strength, morphology, and evolution . . . 72
2.4 Discussion and conclusion . . . 84

3 Mean-field dynamo coefficients in thick accretion disks with initially toroidal fields **89**
3.1 Introduction . . . 89
3.2 Calculating the mean-field dynamo coefficients . . . 92
3.2.1 Fluid variable averaging . . . 93
3.2.2 Fitting for α_{ii} . . . 95
3.3 Results and discussion . . . 96
3.4 Conclusion . . . 100

4 3D RMHD simulations of magnetically-linked relativistic jets **111**
4.1 Introduction . . . 111
4.2 Numerical model . . . 114
4.2.1 Foundational equations . . . 114
4.2.2 Mesh and boundary construction . . . 116
4.2.3 Initial conditions . . . 116
4.2.4 Jet termination . . . 126
4.2.5 Visualisation . . . 126
4.3 Results . . . 128
4.3.1 Jet dynamics and evolution . . . 129
4.3.2 Magnetic field interaction . . . 133
4.3.3 Synthetic observational signatures . . . 135
4.4 Discussion and conclusion . . . 146

5 Conclusion and Outlook **149**
5.1 Conclusion 149
5.2 Future Outlook 151

References **153**

List of figures

1.1 The $M_{\rm BH}$–σ relation. 5
1.2 The unified model of AGN, visualised. 7
1.3 Viscous broadening of an initial ring of matter of mass m at initial radius R_0. 14
1.4 The spring analogue of the magnetorotational instability (MRI). 19
1.5 Schematic representation and Kolmogorov energy spectrum of the turbulent cascade. 22
1.6 Schematic representation of the Stretch-Twist-Fold (STF) method of field amplification 33
1.7 Schematic of the two pathways of the Parker instability. 34
1.8 A multi-wavelength view of Cygnus A. 37
1.9 The basic jet components. 38
1.10 HST observations of the main jet of M87, with component knots labelled. . 39
1.11 Diagram of apparent superluminal motion of an emitting source. 39
1.12 Radio jet power vs host galaxy optical magnitude. 42
1.13 Examples of FR-I and FR-II Type radio jets. 43
1.14 Examples of WAT and NAT jets. 44
1.15 Examples of jet filaments. 49

2.1 Visualisation of the fiducial-resolution disk simulation mesh in 3D 64
2.2 Disk convergence metrics: quality factors in the θ- and ϕ-directions, and magnetic tilt angle, θ_B. 69
2.3 Three-dimensional side-on volume rendering of the simulation domain for `b200`. 70
2.4 Mass flux $\dot{M}$ through the surface at $r = 6r_g$, spatially-averaged magnetic plasma $\langle\beta\rangle$, spatially-averaged geometric scale height $\langle H\rangle$, and spatially-averaged Shakura-Sunyaev alpha $\langle\alpha_{\rm visc}\rangle$. 75

2.5 Late-time-averaged profiles for the density, radial velocity, temperature, and Shakura-Sunyaev alpha as a function of disk radius. 76
2.6 Late-time snapshots of midplane density for `b200` and `b5` 77
2.7 Density along a midplane slice over time, for `b200` and `b200_hi`. 78
2.8 Velocity streamlines for the whole domain and zoom-in on the $r < 100r_g$ region overlaid onto density for `b200` and `b5`. 79
2.9 Magnetic flux, $\Phi_B = \int B_r \cdot dS$, through the ISCO ($r = 6r_g$) and the ratio of signed and unsigned magnetic fluxes. 80
2.10 As in Figure 2.8, but with magnetic field streamlines. 81
2.11 Butterfly diagrams of the azimuthally-averaged B_ϕ at $r = 25r_g$. 82

3.1 The effect of different azimuthal averaging methods upon the variable B_ϕ in the upper hemisphere of `b5` at late time. 94
3.2 Sampling in α and $\tan(\theta_\alpha)$ space. 95
3.3 Instantaneous $\langle \boldsymbol{\varepsilon}'_\phi \rangle$-$\langle \boldsymbol{B}_\phi \rangle$ fit for the `b5` simulation under both averaging methods. 103
3.4 Fitted α_{ii}, C_i values for the `b200` simulation with azimuthal averaging. . . 104
3.5 Fitted α_{ii}, C_i values for the `b5` simulation with azimuthal averaging. 105
3.6 Fitted α_{ii}, C_i values for the `b200` simulation with Gaussian filter averaging. 106
3.7 Fitted α_{ii}, C_i values for the `b5` simulation with Gaussian filter smoothing. . 107
3.8 Variation of $\langle \alpha_{ii} \rangle$ and $\langle C_i \rangle$ with time for the `b200` simulation in the Northern Hemisphere. 108
3.9 Variation of $\langle \alpha_{ii} \rangle$ and $\langle C_i \rangle$ with time for the `b5` simulation in the Northern Hemisphere. 109

4.1 Source ESO 137-006 imaged at 1030 MHz using MeerKAT. 112
4.2 Final snapshots of one realisation of the turbulent background. 121
4.3 Evolution of the turbulent background simulation in time. 122
4.4 The energy power spectra $P(k)$ and histogram of B_y for the turbulent simulation. 122
4.5 Profile of the added jet magnetic field. 124
4.6 A still of simulation `v0.24w0` with components annotated. 128
4.7 Time evolution of the jet density. 130
4.8 Midplane slices of selected jet variables at the simulation end. 131
4.9 Jet length over time. 132
4.10 Flapping at the jet ends of simulation `v0.24` over time. 133

4.11 Flapping at the jet ends of simulation `v0.24w0` over time. 134
4.12 Time evolution of the magnetic field strength. 138
4.13 3D magnetic field streamlines for `v0.94` at $t = 11,300$. 139
4.14 Magnetic current density around the `v0.24` upper jet. 140
4.15 Stokes parameters I, Q, U, and P for each simulation. 141
4.16 Simulated Rotation Measure maps. 142
4.17 The effect of beam size on simulated rotation measure. 143
4.18 The Rotation Measure Structure Function S(r) 144

List of tables

2.1 Disk simulation parameters. 65

3.1 Mean values of α_{ii} and C_i with 95 per cent CI, using azimuthal averaging. . 97
3.2 Mean values of α_{ii} and C_i with 95 per cent CI, using the Gaussian filter method. 97
3.3 Values of α_{ii} from literature. 98

4.1 Code units for the jet simulation. 114
4.2 Initial parameter values for the jet and turbulent background. 125
4.3 Initial parameter values for the `v0.94` and `v0.24` jet simulations. 126

Nomenclature

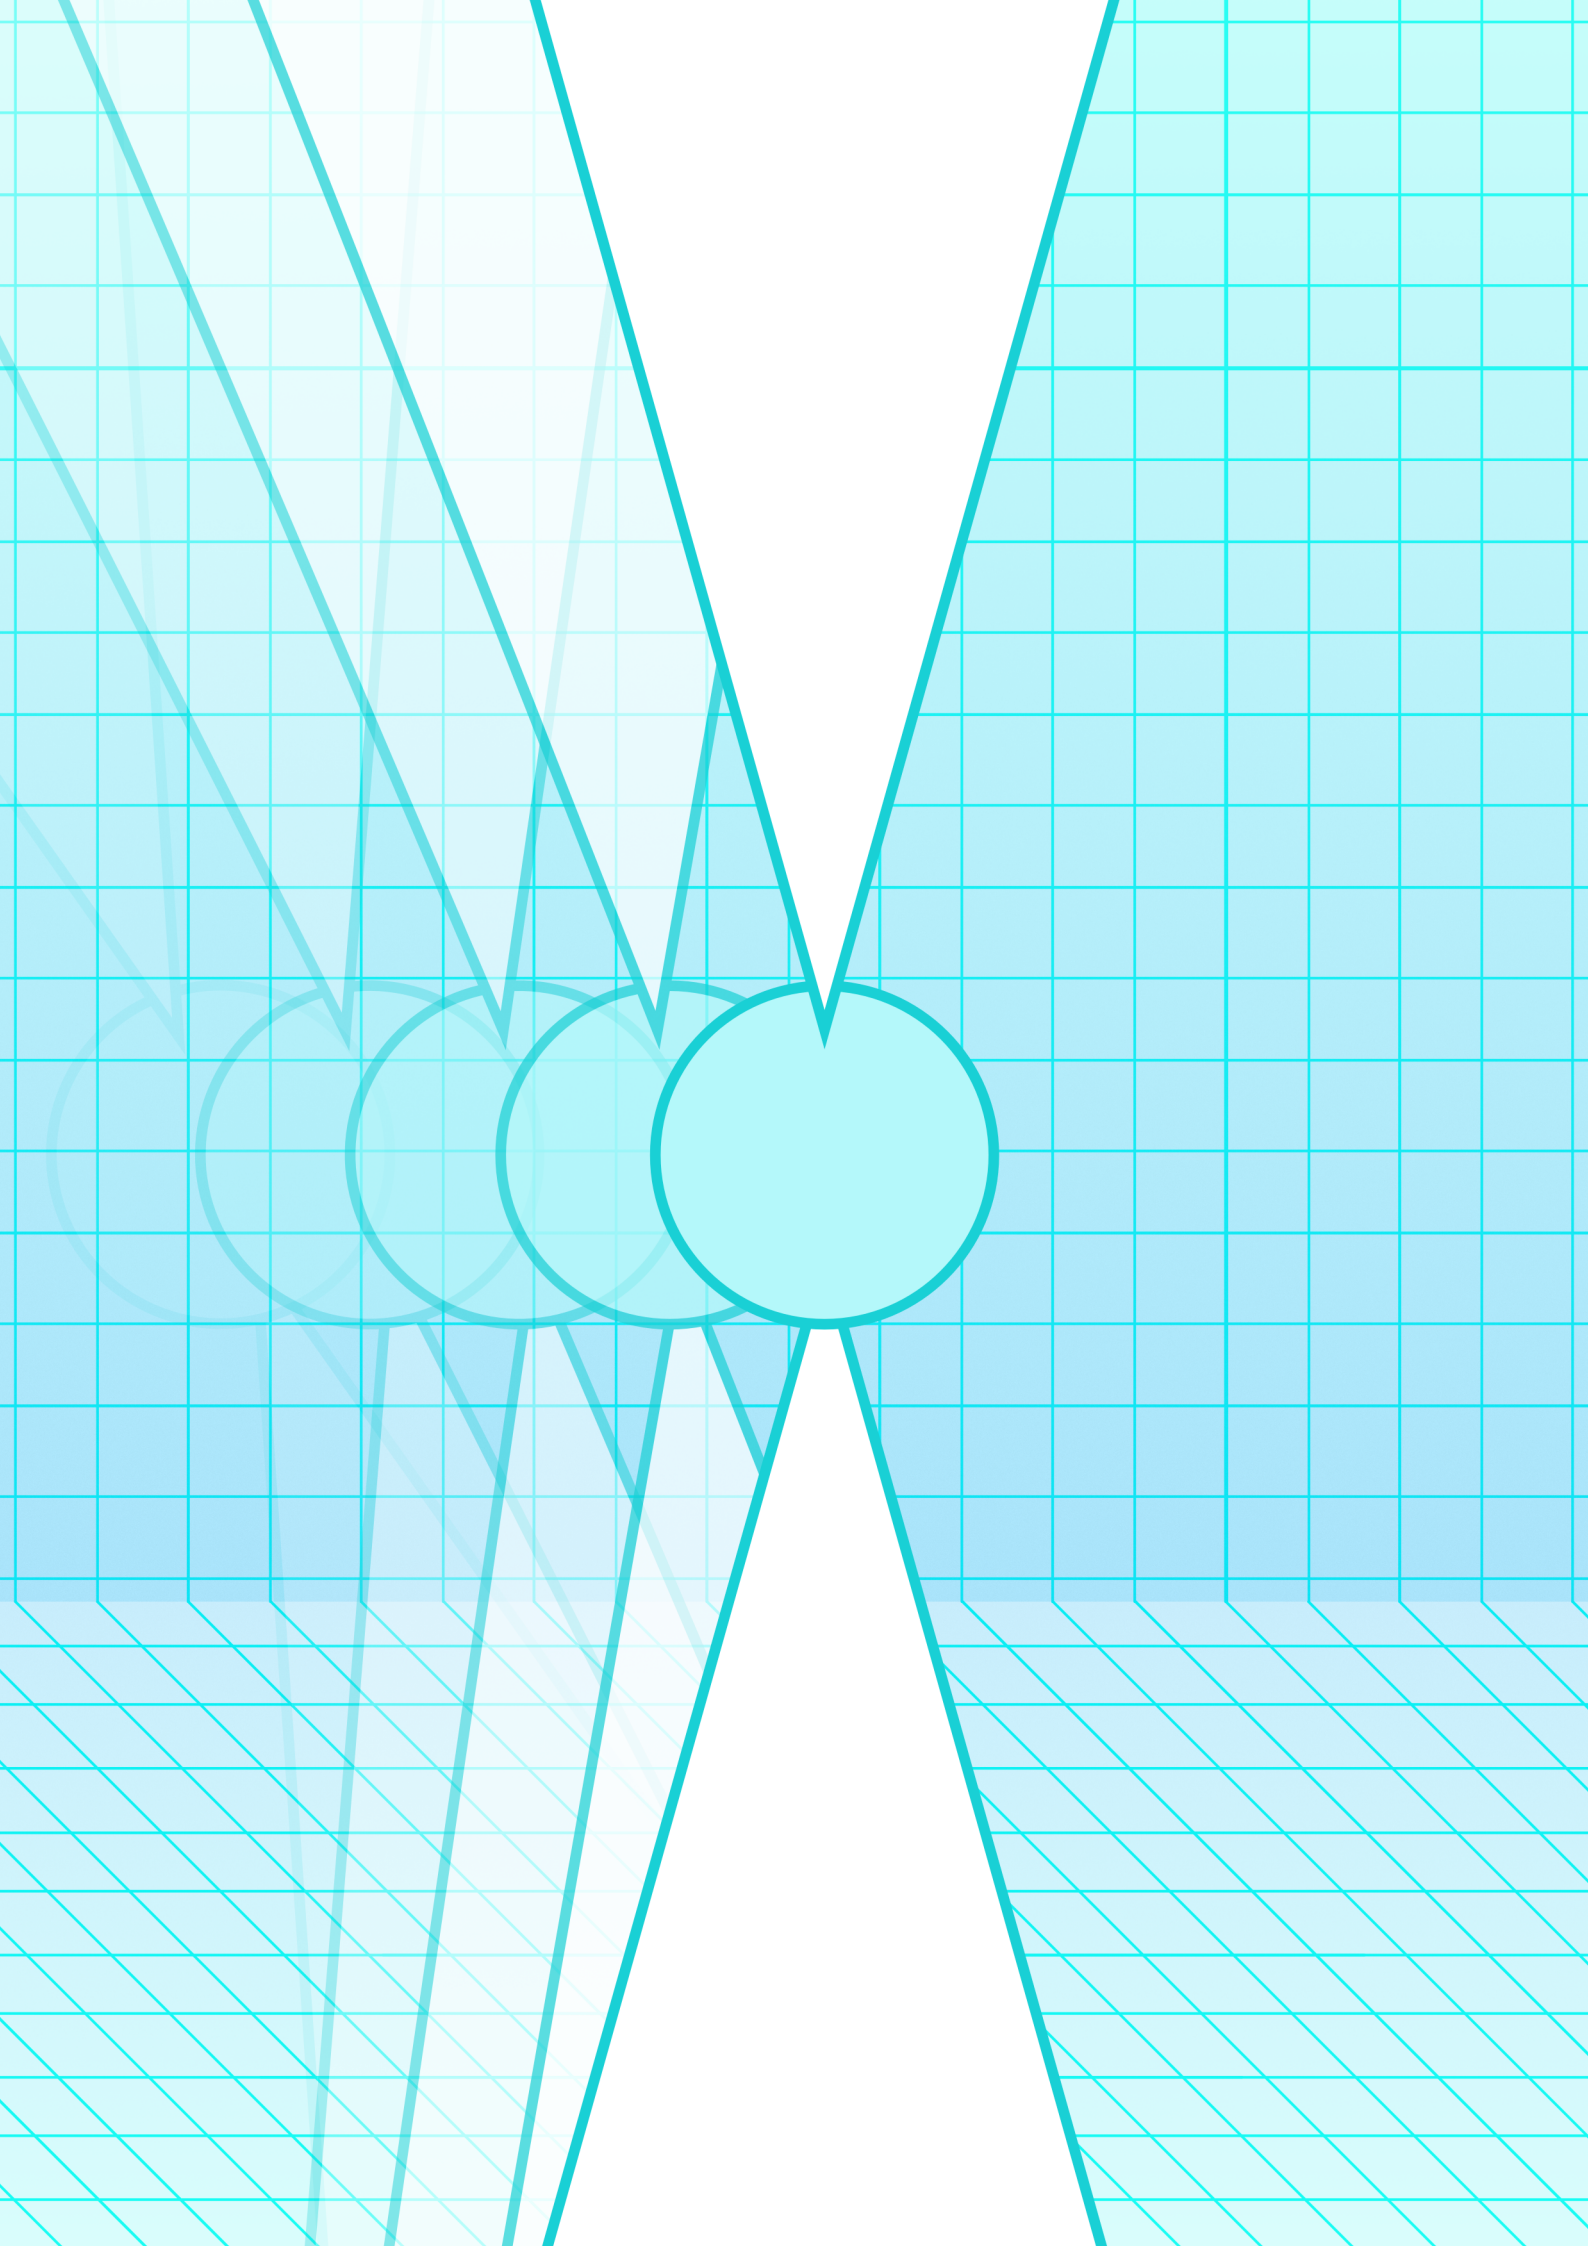
Acronyms / Abbreviations

AGN	Active Galactic Nucleus/Nuclei
ASKAP	Australian Square Kilometre Array Pathfinder
BC	Boundary Condition
BCG	Brightest Cluster Galaxy
BH	Black Hole
BL	Broad Line
BP	Blandford-Payne
BZ	Blandford-Znajek
CFL	Courant, Friedrichs, & Lewy
CSD3	Cambridge Service for Data Driven Discovery
CST	Collimated Synchrotron Threads
CT	Constrained Transport
DCBH	Direct Collapse Black Hole
EMF	Electromotive Force
EOS	Equation of State
FOSA	First Order Smoothing Approximation
FR-I	Fanaroff-Riley Type I

FR-II	Fanaroff-Riley Type II
GRMHD	General Relativistic Magnetohydrodynamics
GRRMHD	General Relativistic Radiative Magnetohydrodynamics
HD	Hydrodynamics
HLLC	Harten–Lax–van Leer Contact
HLLD	Harten-Lax-van Leer Discontinuities
HST	Hubble Space Telescope
IC	Inverse Compton
ICM	Intracluster Medium
IGM	Intergalactic Medium
ISCO	Innermost Stable Circular Orbit
ISM	Interstellar Medium
JWST	James Webb Space Telescope
KH	Kelvin-Helmholtz
LOFAR	Low Frequency Array
LRD	Little Red Dot
MAD	Magnetically Arrested Disk
MCMC	Markov Chain Monte Carlo
MHD	Magnetohydrodynamics
MPI	Message Passing Interface
MRI	Magnetorotational Instability
NAT	Narrow Angle Tail
NH	Northern Hemisphere

NVF	Net-vertical-flux
OU	Ornstein-Uhlenbeck
PI	Parker Instability
PLM	Piecewise Linear Method
POSSUM	Polarisation Sky Survey of the Universe's Magnetism
PPI	Papaloizou-Pringle Instability
RIAF	Radiatively-inefficient Accretion Flow
RK2	2nd-order Runge-Kutta
RL/RQ	Radio Loud / Radio Quiet
RMHD	Relativistic Magnetohydrodynamics
RM	Rotation Measure
RMSF	Rotation Measure Structure Function
RWI	Rossby Wave Instability
SANE	Standard and Normal Evolution
SDE	Stochastic Differential Equation
SH	Southern Hemisphere
SKA	Square Kilometre Array
SMBH	Supermassive Black Hole
SMR	Static Mesh Refinement
SPH	Smoothed Particle Hydrodynamics
STF	Stretch-Twist-Fold
SZ	Sunyaev-Zeldovich
TM	Taub-Mathews

UV	Ultraviolet
WAT	Wide Angle Tail
ZNVF	Zero-net-vertical-flux



Chapter 1

Introduction

1.1 Introduction to Active Galactic Nuclei

In most galaxies the majority of observed light is starlight, contributed from the millions to trillions of individual stars within. But in a select few cases (a few per cent of galaxies in the local Universe), the majority of light does not come from stars but from something else: a highly compact region less than one light day across at the centre of the galaxy, producing radiation across all wavelengths.

These galaxies are termed Active Galaxies, and their nucleus is an Active Galactic Nucleus (AGN). At their heart is a supermassive black hole (SMBH) which exists within a local system composed of the black hole itself, a rotating circumnuclear disk of material which feeds it (or may starve it), and, occasionally, jets which emerge from near the poles of the black hole (BH) and deposit energy far from the nucleus. In this way, the AGN is connected both to its galaxy and the wider environment through the intake and discharge of material and energy. Hence, the classification of AGN as galaxies with bright, non-stellar emission is the manifestation of a growing supermassive black hole. As with any complex system, what affects one portion is likely to generate knock-on effects throughout the rest. It is for these reasons (and others we will see later) that AGN are sometimes referred to as a galactic 'thermostat'.

Supermassive black hole mass is closely correlated with host galaxy bulge mass in the nearby Universe through the $M_{\rm BH}$–$M_{\rm bulge}$ relation (Häring & Rix, 2004; Magorrian et al., 1998) along with relations to bulge luminosity $L_{\rm bulge}$ (Magorrian et al., 1998) and stellar velocity dispersion within the bulge σ (the $M_{\rm BH}$–σ relation; Ferrarese & Merritt, 2000; Gebhardt et al., 2000), strongly suggesting that supermassive black holes evolve alongside their host galaxies. Such correlations are not surprising; despite the hole's comparatively

tiny mass, it punches far above its weight in terms of energy liberated. If the binding energy of the galaxy bulge goes as $E_{\rm gal} \sim M_{\rm gal}\sigma^2$ and the energy liberated by mass growth of the black hole is $E_{\rm BH} = \eta M_{\rm BH}c^2$, then for a typical galaxy mass $M_{\rm gal} \sim 1.4 \times 10^{-3} M_\odot$ and dispersion $\sigma < 400$ km/s and a typical accretion radiative efficiency $\eta \sim 10\%$, this gives (Fabian, 2012)

$$\frac{E_{\rm BH}}{E_{\rm gal}} \gtrsim 80. \tag{1.1}$$

Only a small fraction of this energy coupled to the cold gas is needed to impact the host galaxy significantly.

There are two main historical classes of AGN, the Seyfert galaxies (first described by Carl Seyfert in 1943) and the quasars (first described in a series of papers by Hazard et al., Schmidt, Oke, and Greenstein & Matthews in 1963[1]). They can be broadly distinguished by their luminosity, with Seyfert galaxies emitting visible light at a luminosity comparable to the stars in the galaxy ($\sim 10^{11} L_\odot$) and quasars being much more luminous ($\sim$ 100+ times larger than galactic starlight). Roughly 10–20 per cent of quasars are radio-loud ($L_{\rm 6GHz} \sim 10^{23}$–$10^{27}$ WHz^{-1}) and are associated with collimated outflows known as 'jets' (Fabian, 2012; Urry & Padovani, 1995), with the remainder being radio-quiet ($L_{\rm 6GHz} \sim 10^{21}$–$10^{23}$ WHz^{-1}), with their radio emission attributed to starbursts within the host galaxy. When they were first discovered and described, it was not immediately apparent that Seyfert galaxies and quasars were different presentations of the same object. Seyfert galaxies, being lower luminosity, were observed in the nearby Universe and hence were more readily associated with galaxies. On the contrary, quasars are more common in the early Universe (high redshift), and this, mixed with their high luminosity, obscured the existence of a host galaxy, making them appear as point sources and leading to the name 'quasi-stellar radio source'. It was not until observations with the Hubble Space Telescope in the 1990s that quasars could be definitively linked to galaxies (e.g. Bahcall et al., 1996; Boyce et al., 1996, and references therein).

A great diversity is observed within the AGN population, with individual objects separated into groups defined by the presence or absence of radio emission, the luminosity, and the existence and strength of various broad and narrow lines. Figure 1.2 shows a subset of the most common classifications, namely (for review, see e.g. Carroll & Ostlie, 2017; Padovani et al., 2017),

[1]There are previous mentions of such 'quasi-stellar radio sources', although they were incorrectly assumed to be weird stars within our own Galaxy. The measurement of large redshifts for sources 3C 48 ($z = 0.37$) and 3C 273 ($z = 0.16$) in 1963 was the first confirmation that quasars were associated with distant galaxies, and hence being AGN.

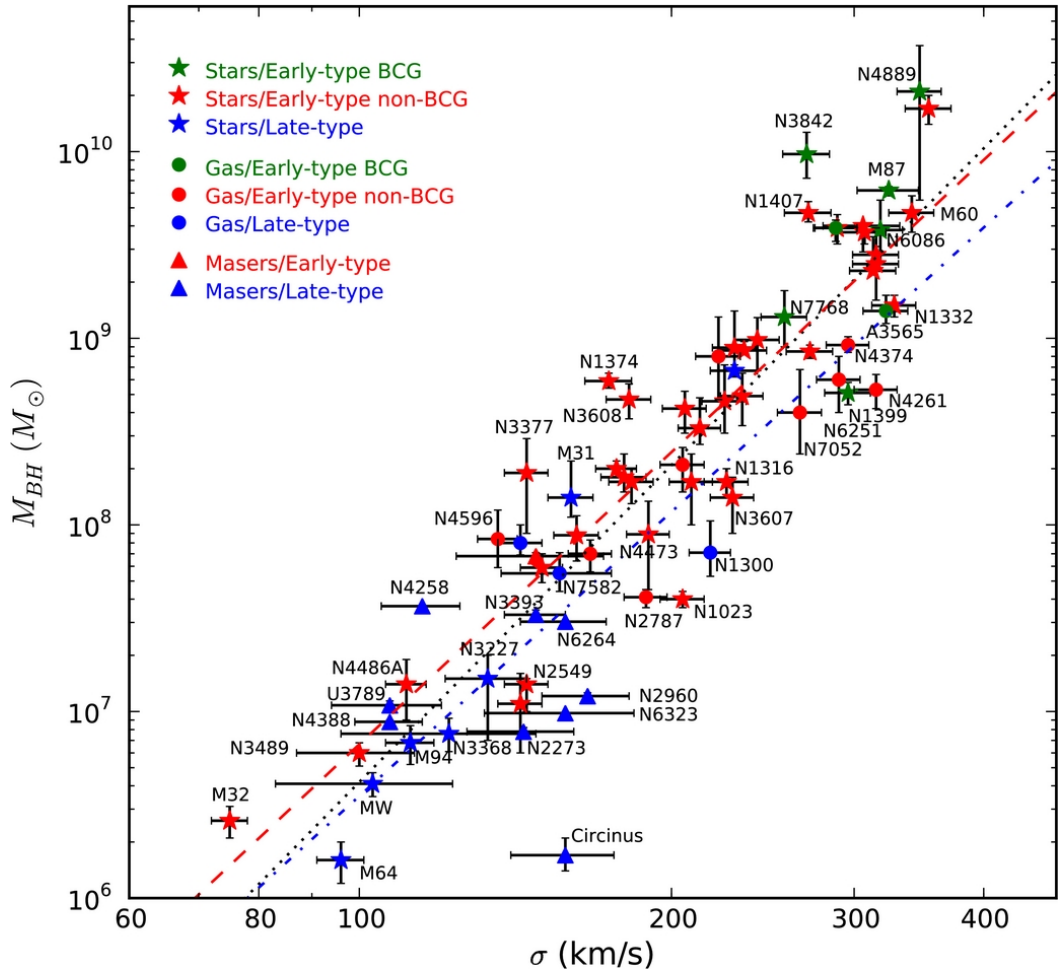


Fig. 1.1 The relation between black hole mass $M_{\rm BH}$ and stellar velocity dispersion σ for a range of sources. Brightest Cluster Galaxies (BCGs) are plotted in green, non-BCG elliptical and S0 galaxies are in red, and spiral galaxies are in blue. Symbol type denotes the method of measurement of $M_{\rm BH}$: stellar dynamics (star), gas dynamics (circle), or maser (triangle). From McConnell & Ma (2013) with permission.

- Seyfert galaxies: quasars whose host galaxy (usually spiral) is also clearly visible, further classified as either type 1 or 2 depending on the presence of broad emission lines, notably Hα. The Seyferts typically have a lower luminosity than quasars and are observed at lower redshift (per the requirement to resolve the host galaxy).

- Radio galaxies: galaxies with measurable radio emission, including a radio nucleus coincident with the galaxy core and commonly observed with extended radio features associated with relativistic jets (see Section 1.4 for discussion on jets). As with Seyferts, radio galaxies can be further classified according to the presence of both broad and narrow emission lines (Broad Line radio galaxies), or of narrow only (Narrow Line radio galaxies).

- Quasars: extremely luminous, point-like AGN, generally observed at high redshift. They may be further classified according to radio loudness, measured either by radio flux density/luminosity or by the ratio of radio-to-optical flux density/luminosity. Radio-loud quasars are thought to host relativistic jets, albeit viewed at a smaller angle than those seen in radio-loud galaxies, making the jets appear shorter in quasars compared to radio-loud galaxies.

- Blazars: galaxies featuring relativistic jets aligned close to the line of sight, as a result often associated with fast variability and high energy γ-radiation.

Later unification schemes (e.g. Antonucci, 1993; Urry & Padovani, 1995) proposed that observational differences between some of these classes may come from the observing angle on the source, although this argument does not appear sufficient on its own to distinguish between different classes of Seyfert galaxies (e.g. Ricci et al., 2011). Further, as many of the classifications covered here have their basis in early observational studies at comparatively low resolution (and, in some cases, before the existence of black holes was widely accepted), there is significant overlap between classes, and the definition of these terms is continually evolving.

1.1.1 Active Galactic Nuclei feedback

The process through which an AGN liberates energy and modulates galaxy evolution is by *AGN feedback*, and while the topic remains highly contentious, it can be broadly categorised into two streams: the *radiative* or *quasar* mode, in which radiation or wide-angle outflows from the process of accretion directly heat the gas surrounding the BH, and the *maintenance* or *kinetic* mode, characterised by the presence of relativistic jets launched near the poles of

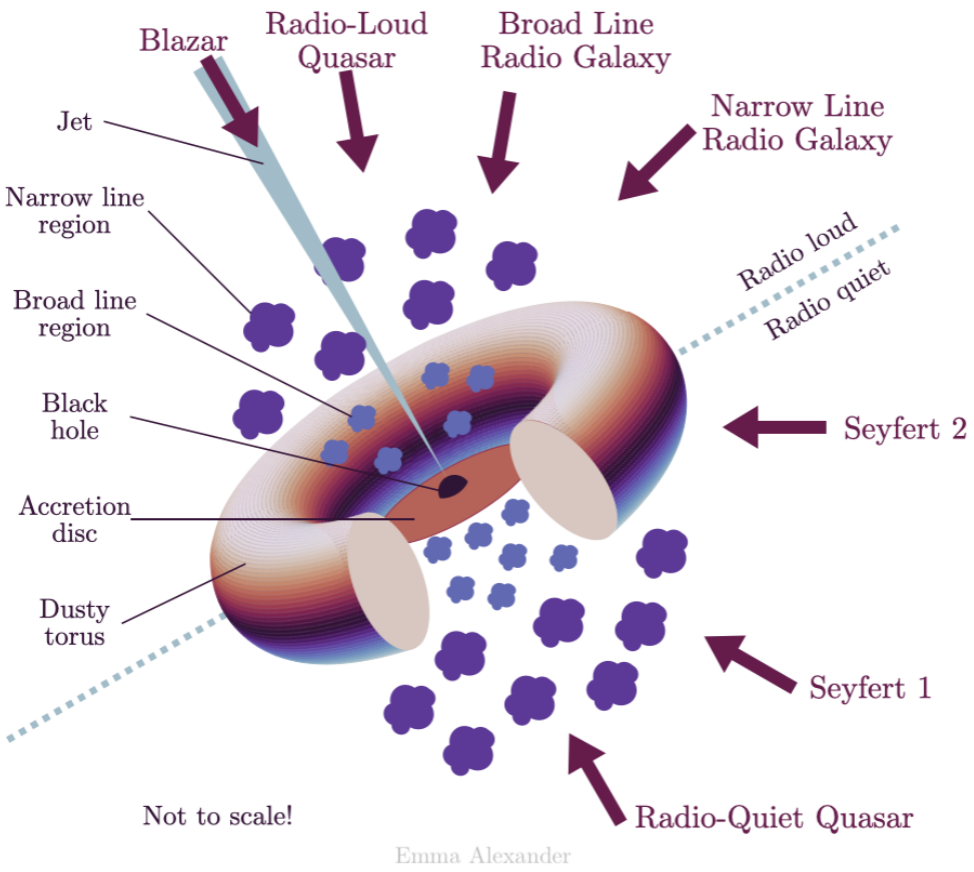


Fig. 1.2 The unified model of AGN, visualised with observational descriptors on the right and components labelled on the left. Figure credit: ©Emma Alexander / CC-BY-SA-4.0

the BH which expand into the interstellar and intracluster media (ISM and ICM) and heat the external gas through weak shocks (Fabian, 2012).

Regardless of terminology, feedback is rather simply conceptualised: activity from the AGN somehow heats the gas, preventing it from cooling and forming stars. As a result, we observe galaxies with less cool gas than expected and older star populations. As such, feedback is a convenient concept to explain why galaxy evolution appears to deviate from prediction. What remains disputed is the numerous processes by which feedback may occur, their strength, timescales, and the area or location of impact.

A full review of the many proposed forms of feedback, their evidence, and their detractors would take a whole thesis in its own right and has been covered in numerous lengthy reviews (e.g. Alexander & Hickox, 2012; Fabian, 2012; King & Pounds, 2015; Kormendy & Ho, 2013; McNamara & Nulsen, 2007, 2012) and in many individual studies. For the sake of brevity, we will focus our attention on feedback as it relates to disk and jet activity.

Disks, while of central importance to the supermassive black hole, are minuscule compared to the scale of the host galaxy, posing immediate difficulty when considering possible feedback processes. Additionally, the rate at which the disk can process material should, in theory, greatly limit how much energy can be output. However, the disk is not entirely isolated. Accretion disks are themselves fed material from a much larger reservoir of galaxy gas, and if a large surplus of material and new, more efficient modes of accretion were to suddenly become available, the theoretical energy output of the disk could be amplified.

In the simplest conception, energy output is capped at the Eddington luminosity,

$$L_{\rm Edd} = \frac{4\pi G M_{\rm BH} m_p c}{\sigma_T} \tag{1.2}$$

where m_p is the proton mass, σ_T is the Thomson cross-section of an electron, and the rest having their usual meaning. A corresponding Eddington mass accretion rate $\dot{M}_{\rm Edd}$ can also be derived assuming $L_{\rm Edd} = \epsilon \dot{M}_{\rm Edd} c^2$ for efficiency ϵ (generally taken as $\epsilon \sim 0.1$). When the bolometric luminosity of the AGN $L_{\rm bol}$ exceeds $L_{\rm Edd}$, it indicates that the outward pressure from radiation exceeds the inward force of gravity, and all accreting gas would be blown away. Whilst some AGN can maintain a high $L_{\rm bol}/L_{\rm Edd}$ ratio, the process should be self-limiting above $L_{\rm bol}/L_{\rm Edd} \approx 1$ and such 'super-Eddington' accretion was previously thought to be rare and short-lived (Alexander & Hickox, 2012).

This picture is made more complicated by the recent *JWST* observations of so-called Little Red Dots (LRDs), a newly discovered population of $z \sim 5$ Broad-Line (BL) AGN with red *JWST*/NIRCam colour and which exhibit a paucity of X-ray emission expected

from accretion processes (Lambrides et al., 2024). Many of these LRDs, over 70 per cent, have inferred BH masses at least an order of magnitude higher than expected from the local M_{BH}–$M_\star$ relation, where $M_\star$ is the stellar mass of the host galaxy (Lambrides et al., 2024; Maiolino et al., 2023; Pacucci et al., 2023). Some sources even approach $M_{BH} \sim M_\star$, in line with theoretical predictions for the evolution of Direct Collapse Black Holes (DCBHs), a kind of 'heavy seed' (e.g. Visbal & Haiman, 2018). Although estimating the BH mass at such high redshift is a difficult task, Pacucci & Loeb (2024) found that regardless of seed size, central BHs at $z \sim 5$ must accrete at or near the Eddington rate for much of the age of the Universe, suggesting that new modes of accretion—particularly super-Eddington accretion—may be needed. General Relativistic Radiative Magnetohydrodynamical (GRRMHD) simulations by Pacucci & Narayan (2024) have additionally found that mildly super-Eddington accretion onto a slowly spinning BH viewed from > 30 deg to the poles could explain ~ 50 per cent of observed LRDs if the masses are overestimated by a factor of 3, which is the approximate reported error at 1σ (e.g. Maiolino et al., 2023). Extrapolating from data at $z \sim$ 1–2, Begelman & Volonteri (2017) place the fraction of $z \gtrsim 5$ AGN accreting at super-Eddington rates above 50 per cent. Additionally, there may be evidence of past super-Eddington accretion in the BH spin, with GRRMHD simulations by Ricarte et al. (2023) finding that super-Eddington accretion in the range $1.4 < \dot{M}/\dot{M}_{\rm Edd} < 13.4$ would lead to a relatively rapid spin down over 3–30 Myrs. A limited sample of ∼20 high-z sources with reliable spin measurements (Reynolds, 2019; Vasudevan et al., 2016) suggests higher mass BHs have lower spins ($a \sim 0.5$–0.7), further supporting the possibility of significant super-Eddington accretion at earlier times.

At moderate-to-high redshift $z \gtrsim 2$, the merging of gas-rich halos leads to the rapid growth of the SMBH and the bulge (Kormendy & Ho, 2013). At some point in this process, accretion becomes radiatively efficient, and 'radiative' or 'quasar' mode feedback is triggered. Outflows—perhaps winds driven by radiation pressure (King & Pounds, 2015)—expel the cold gas that would otherwise have formed stars, leading to a suppression of star formation and causing galaxies to begin moving from being blue and star-forming to 'old, red, and dead', with the majority of AGN appearing in the in-between 'green valley' (Alexander & Hickox, 2012; Nandra et al., 2007). However, some contention remains concerning the spatial range of this effect. Analytic models (e.g. King et al., 2011) suggest that such outflows can act over tens of kpc (galaxy scales) and produce the desired results. Observations, however, are more mixed, with some (e.g. Harrison et al., 2014, 2012; Liu et al., 2013) finding evidence of outflows out to ~ 10 kpc, while others (e.g. Fischer et al., 2018; Husemann et al., 2016; Villar-Martín et al., 2016) find no evidence of outflows beyond 1–2 kpc.

At low redshift ($z < 2$), feedback operates primarily through a different mechanism, the 'maintenance' or 'kinetic' mode. In this process, relativistic jets impart energy onto the incident medium through which they travel and impact. As the jet moves out, it excavates a 'bubble' or 'cavity' in the X-ray emitting media close to the AGN, which buoyantly rises through the intracluster or intragroup medium, bounded by a high-pressure region and a weak shock (Fabian, 2012). In most known cases, the X-ray cavity coincides with a Fanaroff-Riley Type I jet (Section 1.4, Fabian, 2012), although in some cases there is no jet visible due to spectral ageing of electrons—so-called 'ghost' bubbles.

To estimate the energy available for feedback through the cavity, it makes sense to first consider the energy required to form and maintain it, which can be expressed through the enthalpy H of the radio lobe which forms the cavity. H is the sum of the work required to displace the ambient X-ray emitting gas (pV) plus the thermal energy (E) of the lobe contents itself,

$$H = pV + E \tag{1.3}$$

where p is the pressure in the radio lobe and V is the radio lobe volume.

In the case of a constant ratio of specific heats, Γ, the above can be expressed as (McNamara & Nulsen, 2007, 2012)

$$H = \frac{\Gamma}{\Gamma - 1} pV = \begin{cases} 2.5pV \text{ if } \Gamma = 5/3 \\ 4pV \text{ if } \Gamma = 4/3 \end{cases} \tag{1.4}$$

or, if magnetically dominated, $H = 2pV$. Values may be lower limits if strong shocks are also present. Estimates of the thermal energy liberated within the weak shock boundaries of bubbles in the Perseus cluster are 3.5 times pV (Graham et al., 2008), close to the $\Gamma = 4/3$ estimates. Cavity measurements by Rafferty et al. (2006), assuming $H = 4pV$, find a mean cavity power of 1.01×10^{45} erg s^{-1}, enough to offset the mean cooling power of 6.45×10^{44} erg s^{-1}.

The time-averaged jet power required to inflate a pair of cavities is (Rafferty et al., 2006)

$$P_{\text{cav}} = \frac{E_{\text{cav}}}{t} \tag{1.5}$$

where t is the average time between outbursts and E_{cav} is the energy required to create the cavity. t is difficult to estimate since many observed cavities appear as singletons, although measurements from Perseus and Abell 2597 give $t = 60$–100 Myr (McNamara & Nulsen, 2007). It is not yet clear how representative this value is.

As the cavity rises, ambient material will be entrained in its wake, effectively converting gravitational potential energy into kinetic energy, which can then be dissipated as heat into the local ICM. Energy may also be thermalised from shocks. As cavities have a high sound speed, shocks will travel through them faster than they travel around them, forming a vortex ring that allows for greater transfer of shock energy to kinetic energy. Relatedly, energy may be dissipated through sound waves. However, its importance is less clear as it depends on unknown transport coefficients and the nature of both viscosity and magnetic fields in the surrounding medium (McNamara & Nulsen, 2007).

The energy liberated from weak shocks and sound dissipation will likely be most important close to the AGN. At the same time, cavity heating through wake entrainment will be more significant where radio lobes are formed.

Buoyant cavities should be Rayleigh-Taylor and shear unstable; however, most observed cavities appear regular, implying greater stability than simulations suggest. One possible stabilising force is magnetic fields, which, when draped across the active cavity surface, suppress hydrodynamic instabilities (Dursi & Pfrommer, 2008; Lyutikov, 2006). Such a magnetic sheath likely also suppresses the formation of wake turbulence (Bambic et al., 2018a), through reducing the amount of AGN energy that can be thermalised.

Altogether, we have a series of mechanisms or structures which, in theory, contain the necessary energy to offset cooling and create the galaxies we see across cosmic time. However, many questions remain about how this energy is liberated and transmitted throughout a galaxy. Answering such questions is particularly difficult in light of observational constraints, like the lack of strong shocks or steep temperature gradients that would be associated with highly efficient heating. To make the issue even more complex, AGN feedback is not always necessarily *negative*. Positive AGN feedback may occur when a powerful wind or relativistic jet impacts on the surrounding ambient medium, producing shocks and turbulence which can locally induce the gravitational collapse of the material into stars (e.g. Feain et al., 2007; Silk, 2013; Silk & Norman, 2009; Zinn et al., 2013).

To improve our understanding of AGN feedback and galaxy evolution, we must look at the fundamental components in greater detail: the inflow and outflow of material and energy around the central engine. Inflow covers the movement of gas from large galactic radii down to the critical inner accretion radius, while outflows are generally split into collimated jets and uncollimated winds of various origins. Accretion disks and collimated jets are intrinsically linked through the BH, and so we begin first with a review of disks in Section 1.2 before moving on to a review of jets in Section 1.4. We close with a discussion of the numerical challenges in simulating BH processes such as disks and jets.

1.2 Accretion disks

As material flows towards a black hole, it naturally circularises and forms into a rotating disk termed the *accretion disk*. Accretion of material from this disk is the primary method through which an SMBH gains mass outside of infrequent BH-BH merger events and is the likely fuel for the production of relativistic jets. As such, understanding BH accretion is central to any discussion of AGN energetics or SMBH growth, and so we are naturally motivated to understand this component of the BH system. In doing so, we are forced to confront cutting-edge issues in plasma turbulence and computational fluid dynamics.

The field of accretion disk study found its footing in the 1970s with the seminal works of Shakura & Sunyaev (1973) and Lynden-Bell & Pringle (1974), authors whose names would go on to define foundational concepts and terms within the field decades later. These early studies were primarily limited to analytical and 1D models due to the limited computer power available, and as a consequence, the first distinct disk morphology was created and named the Shakura-Sunyaev thin disk, based on the work of Shakura & Sunyaev (1973). In the thin disk model, which is generally defined to have a height-to-radius ratio of $H/R \lesssim 0.1$, the radial and vertical behaviour can be cleanly separated, and hence, the vertical dimension can be safely integrated over with little loss of information, reducing the problem to simple 1D: properties as a function of radius alone.

Shakura & Sunyaev (1973) treat accretion using viscous fluid hydrodynamics (HD). In a non-magnetised disk, the flow of material within the disk can be described by the Navier-Stokes equations of classical fluid hydrodynamics (e.g. Papaloizou & Lin, 1995), expressed here in Eulerian form,

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \boldsymbol{u}) = 0, \tag{1.6}$$

$$\rho \frac{\partial \boldsymbol{u}}{\partial t} + (\rho \boldsymbol{u} \cdot \boldsymbol{\nabla}) \boldsymbol{u} = -\frac{1}{\rho} \nabla P - \nabla \Phi + \frac{\eta}{\rho} \left[\nabla^2 \boldsymbol{u} + \frac{1}{3} \boldsymbol{\nabla} (\boldsymbol{\nabla} \cdot \boldsymbol{u}) \right], \tag{1.7}$$

$$\frac{\partial \rho \epsilon}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \epsilon \boldsymbol{u}) = -P \boldsymbol{\nabla} \cdot \boldsymbol{u}, \tag{1.8}$$

namely the continuity (mass conservation), momentum, and energy equations, for mass density ρ, fluid velocity $\boldsymbol{u}$, gas pressure P, gravitational potential Φ, coefficient of shear viscosity η, and total specific internal energy ϵ. When combined with an Equation of State (EOS) $P = P(\rho, \epsilon)$ and an equation to describe the gravitational potential Φ, the system of equations is closed and can be solved for a given time t. It is important to note that the above equations can be written in many different forms to suit the particular problem. The

equations can be expressed in the mathematically equivalent Lagrangian form using the convective derivative D/Dt. Expanding the above equations by adding external source terms and magnetic fields (leading to magnetohydrodynamics, MHD, shown in detail later) is also possible.

We naturally expect an accretion disk to accrete material onto the central body. However, the process is more complicated than it might first appear. To move material from the disk's outer regions to the inner regions, we must transport angular momentum in the opposite direction. In assessing this problem, we return to our disk in cylindrical coordinates under the thin disk assumption.

To assess the flow of material, we focus on Equations (1.6) and (1.7), the mass continuity and momentum equations (for a similar derivation, see e.g. Papaloizou & Lin, 1995). Given the natural azimuthal symmetry of the problem, it is useful to recast these in cylindrical polar coordinates (R, ϕ, z). Assuming axial symmetry ($\partial/\partial\phi = 0$), hydrostatic equilibrium in the z-direction ($u_z = 0$), and integrating over z yields a new form of the continuity equation,

$$\frac{\partial \Sigma}{\partial t} + \frac{1}{R}\frac{\partial}{\partial R}\left(R\Sigma u_R\right) = 0, \tag{1.9}$$

where $\Sigma(R,t) \equiv \int_{\infty}^{\infty} \rho(R,z,t)dz$ is the surface density. We can similarly express the ϕ-component of the momentum equation (1.7) as

$$\frac{\partial}{\partial t}\left(R\Sigma u_\phi\right) = -\frac{1}{R}\frac{\partial}{\partial R}\left(\Sigma R^2 u_\phi u_R\right) + \frac{1}{R}\frac{\partial}{\partial R}\left(\nu\Sigma R^3 \frac{d\Omega}{dR}\right). \tag{1.10}$$

If the flow is approximately Keplerian then $\partial u_\phi/\partial t = 0$, and one can combine and rearrange Equations (1.9) and (1.10) to get the radial velocity

$$u_R = \frac{\frac{\partial}{\partial R}\left(\nu\Sigma R^3 \frac{d\Omega}{dR}\right)}{R\Sigma\frac{\partial}{\partial R}\left(R\Omega^2\right)} \tag{1.11}$$

which, when substituted back into the continuity equation (1.9) gives a single diffusion equation (Lin & Papaloizou, 1985; Papaloizou & Lin, 1995),

$$\frac{\partial \Sigma}{\partial t} = \frac{1}{R}\frac{\partial}{\partial R}\left[3R^{1/2}\frac{\partial}{\partial R}\left(\nu\Sigma R^{1/2}\right)\right]. \tag{1.12}$$

Hence, an initial ring of matter will broaden over time, with most material moving to lower radii while an extended tail carries angular momentum to higher radii (Fig. 1.3).

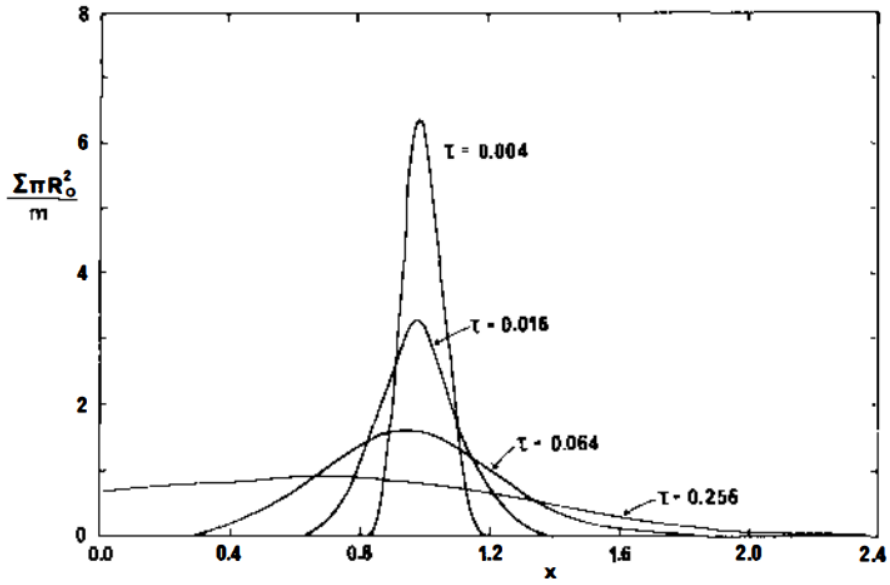


Fig. 1.3 Viscous broadening of an initial ring of matter of mass m at initial radius R_0. Lines are labelled for dimensionless times $\tau = 12\nu t/R_0^2$, as a function of dimensionless radius $x = R/R_0$. Reproduced from Pringle (1981) with permission.

1.2.1 Angular momentum transport in hydrodynamics

In the absence of any source terms in the HD equations, the bracketed section of Equation (1.12) ascribes angular momentum transport to the viscosity ν. Shakura & Sunyaev (1973) recognised that this viscosity term could include contributions from magnetic fields, turbulence, molecular viscosity, and radiative viscosity, although assessing their relative importance would not be an easy task.

Random molecular motion can be dismissed as a significant contributor as it acts over a distance l on timescales $\tau_{\rm mol} \sim l/\nu_{\rm mol}$ for molecular viscosity $\nu_{\rm mol}$. For a molecular viscosity $\nu_{\rm mol} \sim 10^5\ {\rm cm}^2\ {\rm s}^{-1}$, it would take $\sim 3\times10^7$ years for a disturbance to propagate $l \sim 10^{10}$ cm by viscous diffusion alone, far too slow to account for observed variability in AGN accretion disks (Balbus, 2003), and would lead to an accretion rate far too low to explain their observed luminosity. Radiative viscosity can be similarly dismissed as thin Keplerian disks cannot maintain the high temperatures needed for steady accretion driven by this form of viscosity (Loeb & Laor, 1992).

With molecular and radiative viscosity ruled out, only magnetic fields and fluid turbulence remain as possible contributions to angular momentum transport. Unfortunately for Shakura & Sunyaev (1973), the computational resources available at the time did not allow for the simulations needed to measure these contributions (having access to neither magnetic

fields nor 2/3D), and while the existence of magnetic fields could be reasonably assumed throughout the Universe, the presence of turbulence within disks and its nature was an unresolved question. Shakura & Sunyaev (1973) avoided this issue by parameterising away the unknowns under a new non-dimensional constant α. This 'α-prescription' was based on dimensional arguments. Viscous processes introduce stress to the system of form

$$t_{r\phi} \sim -\rho c_s^2 \left(\frac{u_{\rm turb}}{c_s}\right) - \rho c_s^2 \left(\frac{B^2}{4\pi\rho c_s^2}\right) = \alpha\rho c_s^2 \tag{1.13}$$

for density ρ, sound speed c_s, turbulent velocity $u_{\rm turb}$, and magnetic field strength B. The first component of $t_{r\phi}$ above is the contribution from turbulent eddies; the latter is the contribution from magnetic fields within the disk.

The presence of turbulence within the disk did have some basis by appealing to the same arguments that ruled out molecular viscosity: a predicted high Reynolds number of the flow. The Reynolds number, Re, is the ratio of fluid inertia to viscous dissipation,

$$\text{Re} = \frac{vL}{\nu}, \tag{1.14}$$

where v and L are the characteristic flow speed and length. A dimensionless value, the Reynolds number is often associated with a transition from laminar to turbulent flow around some critical value ~ 1000 in laboratory flows. If we take ν to be the molecular viscosity $\nu_{\rm mol}$—a value so small we previously justified dismissing it—we are left with a very high Reynolds number, meaning accretion flows may be unstable to turbulent eddies, and such motions could contribute to the viscous stress tensor in the form of anomalous turbulent viscosity (Balbus & Hawley, 1998), which Lynden-Bell & Pringle (1974) expressed as

$$t_{r\phi} = \rho\nu_{\rm turb}\frac{d\Omega}{dR}. \tag{1.15}$$

Pringle (1981) found that (1.13) and (1.15) can be made equivalent by appropriately defining the turbulent viscosity $\nu_{\rm turb} \sim l_{\rm turb}u_{\rm turb}$, where $l_{\rm turb}$ is an appropriate turbulent length scale and $u_{\rm turb}$ an appropriate turbulent velocity. The largest eddies can be no larger than the height of the disk H and the motion of turbulence is assumed to be subsonic (as any eddies with $u_{\rm turb} > c_s$ would develop shocks and fall back below c_s), so we can express the kinematic (turbulent) viscosity as $\nu_{\rm turb} \sim c_s H \equiv \alpha c_s H$ where $0 < \alpha < 1$, making (1.13) and (1.15) equivalent.

Although it is now widely accepted that Reynolds numbers are high within accretion disks, earlier authors were right to be sceptical of these arguments, as an analysis of disk stability raises issues with the purely HD turbulence model. The assumption of turbulence within a high Reynolds number accretion flow is primarily based on the study of classical planar Couette flow, in which a fluid is sheared against a rigid boundary layer. Balbus et al. (1996) noted that this assumption was invalid as the dynamics of planar sheared and differentially rotating flows are fundamentally different and that local approximations to Cartesian geometry within the disk do not lead to the emergence of Cartesian dynamics on the same scale. Through a Navier-Stokes code, the authors found that hydrodynamical Keplerian disks were both linearly and nonlinearly stable. Even when turbulence is forced, the Reynolds stress decays rapidly and oscillates around zero.

Linear stability arises from the Rayleigh criterion,

$$\frac{d(R^2\Omega)^2}{dR} > 0 \ \ [\text{STABILITY}] \tag{1.16}$$

and is trivially satisfied for astrophysical disks, while Coriolis and tidal forces in differentially-rotating Keplerian flows stabilise the nonlinear instabilities seen in planar shear flows (Hawley et al., 1999). Hence, simple HD accretion flows should be both linearly and nonlinearly *stable* and cannot produce the necessary turbulence or other effective viscosity needed for accretion.

Convectively unstable vertical stratification in disks of finite height is destabilising but would lead to angular momentum transport in the wrong direction (Lesur & Longaretti, 2005). More recent work by Ghosh & Mukhopadhyay (2020) found that hydrodynamical flows could theoretically be made nonlinearly unstable despite the Coriolis force if subjected to sufficiently strong perturbations; however, previous simulations done by Hawley (1990) found that large-amplitude perturbations on the order of 10 per cent or more can act like an artificial viscosity in the form of numerical diffusion (i.e. an error term). In addition, Lesur & Longaretti (2005) found that while a sufficiently high Re flow can transition to turbulence, the resolution and time required is prohibitive and the subsequent turbulent transport too inefficient for Keplerian disks, a finding supported by Ji et al. (2006) in laboratory flows for Reynolds numbers up to Re$\sim 10^6$.

1.2.2 Adding magnetic fields to the mix

With hydrodynamical turbulence ruled out, we are left only with magnetic fields. For this, we must now extend our fluid equations to include contributions from the magnetic field $\boldsymbol{B}$,

yielding magnetohydrodynamics (MHD; additions in magenta),

$$\begin{aligned}
\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \boldsymbol{u}) &= 0, \\
\rho \frac{\partial \boldsymbol{u}}{\partial t} + (\rho \boldsymbol{u} \cdot \boldsymbol{\nabla}) \boldsymbol{u} &= -\frac{1}{\rho} \nabla \left(P + \frac{B^2}{8\pi} \right) - \nabla \Phi + \left(\frac{\boldsymbol{B}}{4\pi} \cdot \boldsymbol{\nabla} \right) \boldsymbol{B} \\
&\quad + \frac{\eta_V}{\rho} \left[\nabla^2 \boldsymbol{u} + \frac{1}{3} \boldsymbol{\nabla} (\boldsymbol{\nabla} \cdot \boldsymbol{u}) \right], \qquad (1.17) \\
\frac{\partial \rho \epsilon}{\partial t} &= -\boldsymbol{\nabla} \cdot \left[(\rho \epsilon + P) \boldsymbol{u} - (\boldsymbol{u} \cdot \boldsymbol{B}) \boldsymbol{B} - \eta_B (\boldsymbol{J} \times \boldsymbol{B}) \right], \qquad (1.18)
\end{aligned}$$

along with the addition of the induction equation,

$$\frac{\partial \boldsymbol{B}}{\partial t} = \boldsymbol{\nabla} \times \left[\boldsymbol{u} \times \boldsymbol{B} - \eta_B \boldsymbol{\nabla} \times \boldsymbol{B} \right], \qquad (1.19)$$

where $\eta_B = c^2/(4\pi\sigma)$ is the magnetic diffusivity and σ the electrical conductivity. If η_B is spatially-constant, then the above simplifies to

$$\frac{\partial \boldsymbol{B}}{\partial t} = \boldsymbol{\nabla} \times [\boldsymbol{u} \times \boldsymbol{B}] + \eta_B \nabla^2 \boldsymbol{B}. \qquad (1.20)$$

We are also beholden to the solenoidal constraint,

$$\boldsymbol{\nabla} \cdot \boldsymbol{B} = 0, \qquad (1.21)$$

which excludes the existence of monopoles. This requirement does not factor into the evolution of the system of equations but constrains the choices of initial $\boldsymbol{B}$.

In ideal MHD, we assume $\eta_B = \eta_V = 0$. Extending the MHD equations to include other non-ideal effects like ambipolar diffusion and the Hall effect is also possible. In the context of this thesis, ideal MHD will mostly suffice.

It is also useful at this point to introduce the magnetic plasma β, which is simply the ratio of magnetic pressure to gas pressure,

$$\beta = \frac{P_{\text{gas}}}{P_{\text{mag}}} = \frac{n k_B T}{B^2/2\mu_0} \qquad (1.22)$$

where n is the number density of the gas/fluid, k_B is Boltzmann's constant, T is the temperature of the gas/fluid, and μ_0 is the vacuum permeability.

In the spirit of Shakura & Sunyaev (1973), the vertically-averaged stress tensor $T_{r\phi} = \int t_{r\phi} dz$ can alternatively be expressed as the sum of magnetic and non-magnetic terms

$$T_{r\phi} = \langle R_{r\phi} + M_{r\phi} \rangle \tag{1.23}$$

where $R_{r\phi} = \rho u_r \delta u_\phi$ is the turbulent Reynolds stress and $M_{r\phi} = -(B_r B_\phi)/4\pi$ is the magnetic Maxwell stress. Here, δu_ϕ is the difference between the local azimuthal velocity u_ϕ and the Keplerian velocity $R\Omega_K$. As we have already seen that the Reynolds (i.e. hydrodynamic) stress alone is insufficient to drive transport angular momentum outwards, we will consider the magnetic stresses in greater detail.

Velikhov (1959) and Chandrasekhar (1960) noted that the Rayleigh stability criterion (1.16) could not be recovered in the limit that the magnetic field strength went to zero in an incompressible Couette flow threaded by a vertical magnetic field, indicating that magnetic fields played an essential role in destabiliting fluid flow and that this role did not scale simply with the strength of the field. Fricke (1969) derived a form of this criterion in the local Boussinesq approximation[2] for differentially rotating stars threaded by poloidal magnetic fields, but it wasn't until the work of Balbus & Hawley (1991) that the nature of this magnetic instability was explored for differentially-rotating astrophysical disks. It has subsequently come to be known as the magnetorotational instability (MRI), although some works from the 1990s (e.g. Brandenburg et al., 1995; Curry et al., 1994) alternatively refer to it as the Velikhov-Chandrasekhar or Balbus-Hawley instability, per the usual naming conventions. For consistency here, I will refer to it solely as the MRI.

1.2.3 The Magnetorotational Instability

The MRI can be intuitively understood via analogy by imagining two point masses embedded in a differentially rotating flow, connected via a weak spring (representing a weak vertical magnetic field; Fig. 1.4). If one of the masses is displaced, it will stretch the spring and generate tension, transferring angular momentum L from the innermost mass m_{in} to the outermost m_{out}. With a lower angular momentum, m_{in} falls to a lower radius as $L \propto \sqrt{r}$, generating more tension which leads to further inwards movement. At the same time, m_{out} moves continuously in the other direction, gaining more angular momentum as it moves outwards. Hence, small perturbations within a magnetised accretion disk lead to runaway

[2]The Boussinesq approximation is one in which the velocity field is nearly incompressible over all scales of interest, where the inverse characteristic timescale $(t_{\text{char}})^{-1} \equiv \boldsymbol{\nabla} \cdot \boldsymbol{u}$ is long compared to the characteristic turnover time of turbulence, but not necessarily long compared to other timescales in the flow (Balbus & Hawley, 1998).

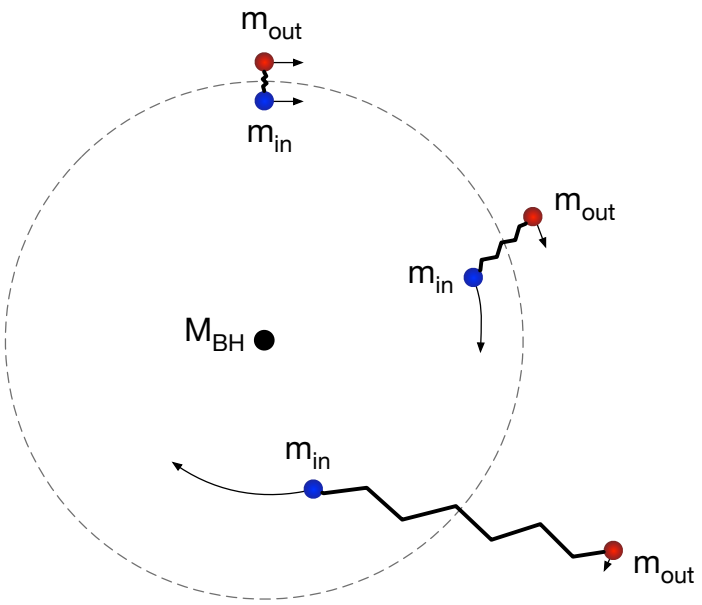


Fig. 1.4 The spring analogue of the magnetorotational instability (MRI) for equal point masses $m_{\rm in}$ and $m_{\rm out}$ within the gravitational potential of a central mass $M_{\rm BH}$.

radial movement of material through the outwards transport of angular momentum—an instability.

Expressed mathematically, the local equations of motion for this toy model for which R is the radius, $\hat{\boldsymbol{x}}$ points radially outwards, and $\hat{\boldsymbol{y}}$ points in the azimuthal direction, read (Balbus & Hawley, 1998):

$$\ddot{x} - 2\Omega\dot{y} = -x\frac{d\Omega^2}{d\ln R} - Kx \tag{1.24}$$

$$\ddot{y} + 2\Omega\dot{x} = -Ky, \tag{1.25}$$

where 2Ω is the contribution from the Coriolis force, $d\Omega^2/d\ln R$ is the tidal force, and Kx and Ky are external restoring forces with spring constant K.

Equations (1.24) and (1.25) will only have solutions with a time dependence of the form $e^{-i\omega t}$ provided that the angular frequency ω satisfies the *dispersion relation*

$$\omega^4 - \omega^2\left[\kappa^2 + 2K\right] + K\left[K + \frac{d\Omega^2}{d\ln R}\right] = 0. \tag{1.26}$$

More formally, Balbus & Hawley (1991) derive the stability conditions of a magnetised accretion disk via the dispersion relation for the angular frequency ω of perturbations of the form $e^{i(k_r r + k_z z - \omega t)}$, which in the case of a Keplerian disk can be expressed as (Balbus, 2003;

Balbus & Hawley, 1998)

$$\omega^4 - \omega^2 \left[\kappa^2 + 2\left(\boldsymbol{k}\cdot\boldsymbol{v}_A\right)^2\right] + \left(\boldsymbol{k}\cdot\boldsymbol{v}_A\right)^2 \left[\left(\boldsymbol{k}\cdot\boldsymbol{v}_A\right)^2 + \frac{d\Omega^2}{d\ln R}\right] = 0, \tag{1.27}$$

with wavenumber $\boldsymbol{k} \equiv k_r\hat{\boldsymbol{r}} + k_z\hat{\boldsymbol{z}}$, Alfvén velocity

$$\boldsymbol{v}_A = \frac{\boldsymbol{B}}{\sqrt{4\pi\rho}}, \tag{1.28}$$

and epicyclic frequency κ

$$\kappa^2 = 4\Omega^2 + \frac{d\Omega^2}{d\ln R}. \tag{1.29}$$

It should be noted that Equation (1.27) above is functionally identical to the spring-like case of (1.26) with a spring constant $K = \left(\boldsymbol{k}\cdot\boldsymbol{v}_A\right)^2$.

The dispersion relation (1.27) is quadratic in ω^2, and so it will have unstable roots if the discriminant is negative, i.e.

$$\left[\kappa^2 + \left(\boldsymbol{k}\cdot\boldsymbol{v}_A\right)^2\right] - 4\left(\boldsymbol{k}\cdot\boldsymbol{v}_A\right)^2 \left[\left(\boldsymbol{k}\cdot\boldsymbol{v}_A\right)^2 + \frac{d\Omega^2}{d\ln R}\right] < 0. \tag{1.30}$$

There are no unstable roots for (1.27) so long as

$$\frac{d\Omega^2}{d\ln R} \geq 0, \tag{1.31}$$

but this is almost universally violated in accretion disks. This may be recognised as being similar to the Rayleigh stability criterion (1.16), but with angular velocity in place of angular momentum. If (1.31) is not satisfied, and the field is sufficiently weak, then an unstable mode exists where ω is both positive and imaginary. For wave solutions of the form $e^{-i\omega t}$, the unstable mode leads to an exponential growth of small perturbations—the MRI. From (1.30), this will occur for wavenumbers below

$$k_{\text{crit}} = \frac{1}{v_A}\left|\frac{d\Omega^2}{d\ln R}\right|^{1/2}, \tag{1.32}$$

or equivalently, for wavelengths above

$$\lambda_{\text{MRI}} = 2\pi\frac{v_A}{\Omega}, \tag{1.33}$$

while the fastest-growing mode occurs nearby at (Balbus, 2003)

$$(\boldsymbol{k} \cdot \boldsymbol{v}_A)^2_{\text{max}} = \left(\frac{1}{4} + \frac{\kappa^2}{16\Omega^2}\right) \left|\frac{d\Omega^2}{d \ln R}\right|. \tag{1.34}$$

It should be noted that the new stability criterion, which replaces the Rayleigh criterion (1.31), is still not a function of magnetic field strength, indicating that the MRI should be present even for arbitrarily weak fields. This can also be seen in (1.27): the magnetic field strength (through the Aflvén velocity $\boldsymbol{v}_A$) is always multiplied by the wavenumber $\boldsymbol{k}$, which grows as the field shrinks.

From this point, there are a few possible limiting cases for the MRI. Firstly, the instability will saturate when λ_{MRI} exceeds the height of the disk ($2H$), which occurs for $v_{A,z} \sim c_s$ or when the plasma β falls below ~ 3 (Balbus & Hawley, 1991). Alternatively, the MRI may be regulated by the growth of 'parasitic modes' which feed off of the extra energy generated by the MRI (Goodman & Xu, 1994). Finally, the instability may be regulated by the process of *turbulent cascade*, in which energy 'cascades' from the largest turbulent structures down to the smallest, at which point it is dissipated via molecular viscosity (Fig. 1.5). The transition from energy cascade to viscous dissipation occurs at the Kolmogorov length scale,

$$\eta = \left(\frac{\nu^3_{\text{turb}}}{\epsilon}\right)^{1/4}, \tag{1.35}$$

where ν_{turb} is the kinematic turbulent viscosity, and ϵ is the rate of turbulent energy dissipation per unit mass.

The MRI is considered a *local* linear instability as it does not care for the global behaviour of the disk nor its boundary conditions and arises from small-scale dynamics. The global mode structure is still of interest, though. The inner radius of an accretion disk may be subject to a different magnetic environment as in young stellar objects (e.g. Königl, 1991), while large vertical wavelengths define the saturation limit of the MRI.

1.2.4 Summary

To summarise, accretion disks are the primary method by which the SMBH accretes material and grows mass outside of merger events, and so is likely the dominant process within galaxies most of the time. The process of accretion is now known to be governed by an instability that arises from the interaction of weak magnetic fields and differential rotation of the disk called the magnetorotational instability, or MRI. However, there are still many open

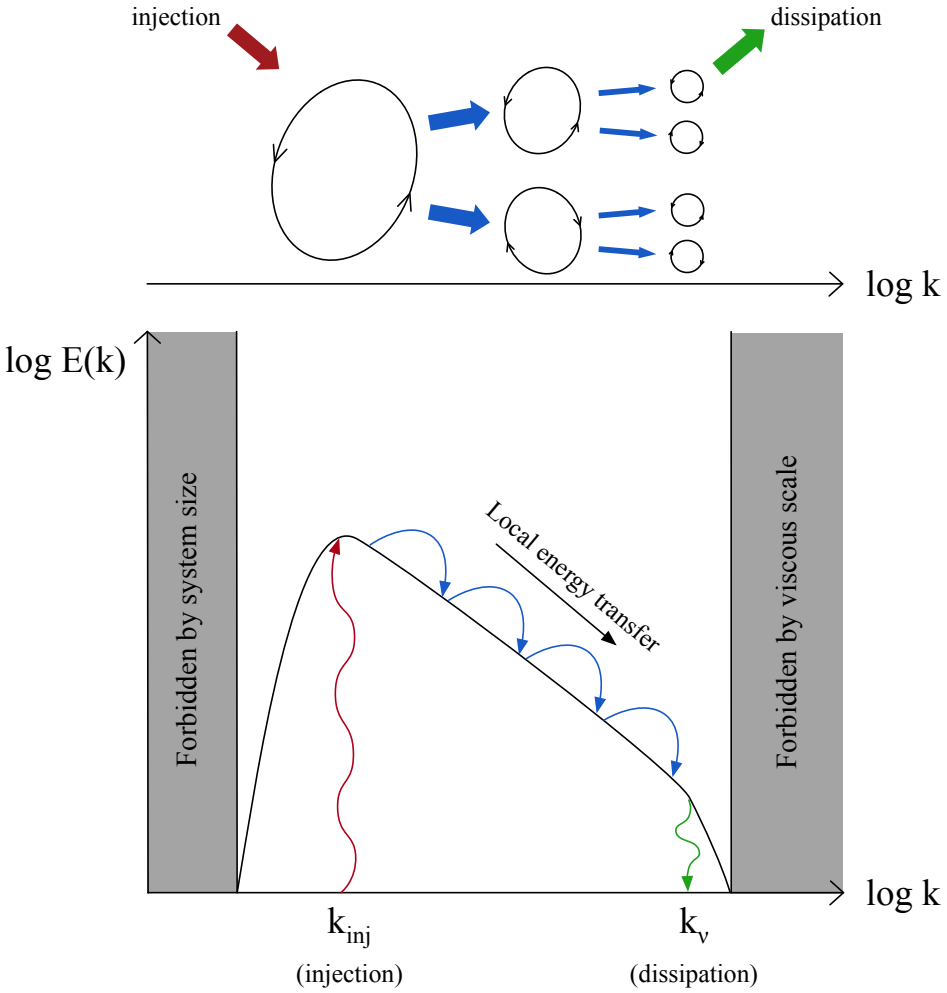


Fig. 1.5 (*Top*) Schematic representation of the turbulent cascade, with energy injected at the largest scales and transferred through subsequent smaller scales until it is ultimately dissipated at the viscous scale. (*Bottom*) The associated Kolmogorov energy spectrum of the turbulent cascade.

questions remaining. Upon saturation the MRI is highly non-linear and not easily amenable to analytic study, especially in its complex interactions with turbulence in the disk. In a literal sense, an 'accretion' disk does not exist without its magnetic field, and so our understanding of accretion processes and how they link to large-scale dynamics (e.g. jets, AGN feedback) is dependent upon our understanding of magnetic field evolution within the disk.

1.3 Characterising the large-scale dynamo

Accretion disks are expected to be initially weakly magnetised; however, without some sustaining dynamics this field should eventually decay away via Ohmic dissipation through the modified MHD induction equation (1.20),

$$\frac{\partial \boldsymbol{B}}{\partial t} = \boldsymbol{\nabla} \times (\boldsymbol{u} \times \boldsymbol{B}) + \eta \nabla^2 \boldsymbol{B}.$$

The first term on the right-hand side is the inductive term, which acts as a source term. The final term is a sink term for magnetic energy through microphysical Ohmic dissipation. The balance between induction and Ohmic dissipation—and hence the growth or decay of the field—is captured by the *magnetic Reynolds number*,

$$\mathrm{Rm} = \frac{ul}{\eta}, \tag{1.36}$$

where u is the characteristic flow speed and l is a typical length over which the magnetic field varies, akin to that of (1.14). The ratio of Rm and Re is the magnetic Prandtl number,

$$\mathrm{Pm} = \frac{\mathrm{Rm}}{\mathrm{Re}} = \frac{\nu}{\eta}, \tag{1.37}$$

which indicates the relative strength of kinematic viscosity to magnetic diffusivity. In a direct turbulent cascade, Pm dictates whether the cascade hits the viscous dissipation scale ($\mathrm{Pm} > 1$) or the magnetic diffusivity scale ($\mathrm{Pm} < 1$) first.

Without some external driving mechanism to act as a source term, the magnetic field will naturally decay away via Ohmic dissipation. To sustain and amplify the field, we need a viable *dynamo model*.

The classical dynamo finds its basis in Faraday's law: when a conductor moves through a magnetic field $\boldsymbol{B}$ with velocity $\boldsymbol{u}$, it induces an electromotive force (EMF) $(\boldsymbol{u} \times \boldsymbol{B})$ which drives a current density $\boldsymbol{J}$—the dynamo. By Ampere's law ($\boldsymbol{J} = \mu^{-1} \boldsymbol{\nabla} \times \boldsymbol{B}'$), this current density then induces a magnetic field $\boldsymbol{B}'$. If $\boldsymbol{B}$ and $\boldsymbol{B}'$ are the same, then the dynamo is said to be self-excited. A more liberal definition of a dynamo follows from Moffatt (1978): a given velocity field $\boldsymbol{u}$ is said to act as a dynamo if the total magnetic energy remains above zero as $t \to \infty$.

1.3.1 Mean-field theory

The dynamo, as we have defined it, is a vague concept open to many different forms of study and methods of characterisation. In general, to characterise a dynamo we must look for solutions to the MHD equations, including the induction equation (1.20) describing the evolution of the magnetic field $\boldsymbol{B}$ in response to a velocity $\boldsymbol{u}$, and the momentum equation (1.17) which includes the back-reaction of the magnetic field onto the fluid flow via the Lorentz force

$$\boldsymbol{J} \times \boldsymbol{B} = \frac{(\boldsymbol{B} \cdot \nabla)\boldsymbol{B}}{4\pi} - \nabla \left(\frac{B^2}{8\pi} \right), \tag{1.38}$$

along with the rest of the MHD equations.

In practice, this means solving a set of coupled, nonlinear partial differential equations—a process that is both mathematically and computationally intensive, if not intractable.

To make the problem more amenable to analysis, a number of simplifying assumptions can be made. The most widespread of these assumptions is that of the *kinematic* model. In kinematic dynamo theory, we assume that the magnetic field has no back-reaction on the fluid flow; that is, we assume no Lorentz force is acting. If the velocity field is prescribed, then the induction equation can be solved by itself. Still, the problem is not simple. Early work by Cowling (1933) showed that axisymmetric fields could not be sustained by dynamo action, while Zel'dovich (1957) showed that dynamos could not operate in anything less than 3D. The kinds of simple flows that would make the induction equation analytically solvable are not those that generate or sustain dynamo action.

Another further simplifying assumption may be made: we may assume that our magnetic field and fluid flow can be separated into large-scale and small-scale components. This method of analysis is known as *mean-field electrodynamics* and was first formulated by Steenbeck et al. (1966) (although see also Charbonneau, 2010; Gilbert, 2002; Hughes & Tobias, 2010; Krause & Rädler, 1980; Moffatt, 1978; Parker, 1955; Roberts & Soward, 1992, for discussion and review). Here we additionally rely on the kinematic assumption, so the theory is sometimes referred to as the *kinematic mean-field theory*. Whilst the theory is widely applied, it is not always applicable. In particular, mean-field theory relies on a separation of scale between small-scale (characteristic length l) and large-scale ($L \gg l$), and so cannot be applied if such a separation does not exist ($L \sim l$). Further, the theory technically only applies for flows with magnetic Reynolds numbers $\mathrm{Rm} \gg 1$ or Strouhal number $S \ll 1$, neither of which is typically satisfied in an accretion disk and leads to consequences which we will explore later. Despite these setbacks, kinematic mean-field theory remains dominant in astrophysical fluid dynamics.

The kinematic mean-field theory begins by expressing $\boldsymbol{B}$ and $\boldsymbol{u}$ as a sum of mean components $\overline{\boldsymbol{B}}$, $\overline{\boldsymbol{u}}$ and small-scale fluctuations $\boldsymbol{b}'$, $\boldsymbol{u}'$:

$$\begin{aligned} \boldsymbol{B} &= \overline{\boldsymbol{B}} + \boldsymbol{b}' \\ \boldsymbol{u} &= \overline{\boldsymbol{u}} + \boldsymbol{u}', \end{aligned} \tag{1.39}$$

where we assume that $\langle \boldsymbol{b}' \rangle = \langle \boldsymbol{u}' \rangle = 0$. In the accretion disk setting, the mean-field components are often taken to be axisymmetric while $\boldsymbol{b}'$ and $\boldsymbol{u}'$ are likely to be asymmetric.

This decomposition method can also be applied to the MHD induction equation (1.20),

$$\frac{\partial \overline{\boldsymbol{B}}}{\partial t} = \boldsymbol{\nabla} \times (\overline{\boldsymbol{u}} \times \overline{\boldsymbol{B}}) + (\boldsymbol{\nabla} \times \overline{\boldsymbol{\varepsilon}}) + \eta \nabla^2 \overline{\boldsymbol{B}} \tag{1.40}$$

$$\frac{\partial \boldsymbol{b}'}{\partial t} = \boldsymbol{\nabla} \times (\overline{\boldsymbol{u}} \times \boldsymbol{b}') + \boldsymbol{\nabla} \times (\boldsymbol{u}' \times \overline{\boldsymbol{B}}) + \boldsymbol{\nabla} \times \boldsymbol{G} + \eta \nabla^2 \boldsymbol{b}', \tag{1.41}$$

where $\boldsymbol{\varepsilon} = \langle \boldsymbol{u}' \times \boldsymbol{b}' \rangle$ is the EMF generated by fluctuating fields, and $\boldsymbol{G} = \boldsymbol{u}' \times \boldsymbol{b}' - \langle \boldsymbol{u}' \times \boldsymbol{b}' \rangle$ (Moffatt, 1978).

In an accretion disk setting, it is more natural to consider the magnetic field as the sum of poloidal and toroidal components, making use of the system's natural symmetries. With this in mind, the axisymmetric mean field ($\overline{\boldsymbol{B}}$) can be expressed as a linear sum of toroidal (T) and poloidal (P) components:

$$\overline{\boldsymbol{B}} = \boldsymbol{B}_T + \boldsymbol{B}_P, \tag{1.42}$$

with

$$\begin{aligned} \boldsymbol{B}_P &= \nabla \times (A(r, \theta, t) \hat{\mathbf{e}}_\phi) \\ \boldsymbol{B}_T &= B(r, \theta, t) \hat{\mathbf{e}}_\phi, \end{aligned} \tag{1.43}$$

for toroidal vector potential $A(r, \theta, t)\hat{\mathbf{e}}_\phi$. The same method can be used for $\overline{\boldsymbol{u}}$.

If we consider a fully axisymmetric field as in Equation (1.40), then, applying the rules from Equation (1.42) yields:

$$\frac{\partial A_\phi}{\partial t} + r^{-1} \boldsymbol{u}_P \cdot \boldsymbol{\nabla}(r A_\phi) = \eta \Delta A_\phi + \overline{\varepsilon}_\phi \tag{1.44}$$

$$\frac{\partial B_\phi}{\partial t} + r \boldsymbol{u}_P \cdot \boldsymbol{\nabla}(r^{-1} B_\phi) = \eta \Delta B_\phi + r \boldsymbol{B}_P \cdot \boldsymbol{\nabla} \Omega + (\boldsymbol{\nabla} \times \overline{\boldsymbol{\varepsilon}})_\phi. \tag{1.45}$$

If there is no active EMF, i.e. $\overline{\boldsymbol{\varepsilon}} = 0$, then by Equation (1.44) $A \to 0$ as $t \to \infty$. For Equation (1.45), the only source term is $r\boldsymbol{B}_P \cdot \boldsymbol{\nabla}\Omega$, which also goes to zero since $\boldsymbol{B}_P = \boldsymbol{\nabla} \times \boldsymbol{A}_T$. Hence, $\boldsymbol{B}_P$ can feed $\boldsymbol{B}_T$, but there is no path to closure on which $\boldsymbol{B}_T$ can feed $\boldsymbol{B}_P$.

Even in the absence of a driving EMF, the poloidal field can feed a toroidal field through $r\boldsymbol{B}_P \cdot \boldsymbol{\nabla}\Omega$; this is known as the Ω-effect after the symbol commonly used for shear. The Ω-effect, however, does not provide a path to closure in which the toroidal field can feed back to the poloidal, and hence, the magnetic field will die away without some additional mechanism. This is *Cowling's anti-dynamo theorem* (Cowling, 1933):

> *No purely axisymmetric magnetic field can be sustained by a self-exciting dynamo.*

For the dynamo to be sustained, we must find a source for $\overline{\boldsymbol{\varepsilon}}$. This can be achieved as long as $\boldsymbol{\nabla} \times (\boldsymbol{u}' \times \overline{\boldsymbol{B}}) \neq 0$; this can act as a source for $\boldsymbol{B}'$ in Equation (1.41) and hence a source for $\overline{\boldsymbol{\varepsilon}}$.

1.3.2 First-Order Smoothing Approximation

Equation (1.41) contains a rather awkward expression in $\boldsymbol{G} = \boldsymbol{u}' \times \boldsymbol{b}' - \langle \boldsymbol{u}' \times \boldsymbol{b}' \rangle$ which it would be convenient to ignore, as its presence makes calculating $\boldsymbol{b}'$ (and hence $\boldsymbol{\varepsilon}$) difficult. The neglect of this term is known as the First-Order Smoothing Approximation (FOSA) and is valid so long as the system satisfies the condition

$$\min(\mathrm{Rm}, S) \ll 1 \tag{1.46}$$

where $\mathrm{Rm} = u_0 l_0/\eta$ is the magnetic Reynolds number and $S = u_0\tau_0/l_0$ is the Strouhal number, with u_0 the rms velocity fluctuation, and τ_0 and l_0 being the characteristic time scale and length scale of the velocity field (Hughes & Tobias, 2010; Krause & Rädler, 1980).

Under FOSA, it is also customary to assume $\overline{\boldsymbol{u}} = 0$, removing the $\overline{\boldsymbol{u}} \times \boldsymbol{b}'$ term from Equation (1.41) and leaving us with,

$$\frac{\partial \boldsymbol{b}'}{\partial t} = \boldsymbol{\nabla} \times (\boldsymbol{u}' \times \overline{\boldsymbol{B}}) + \eta \nabla^2 \boldsymbol{b}'. \tag{1.47}$$

The above argument technically only applies so long as $\overline{\boldsymbol{u}} \sim 0$, however equivalent expressions can be derived in the case of a dominant azimuthal flow $\overline{\boldsymbol{u}} = u_\phi$ (see, e.g. Moffatt, 1978, Chapter 8) using the Braginskiĭ method (Braginskiĭ, 1965), which yields functionally equivalent solutions. We will follow the more straightforward case $\overline{\boldsymbol{u}} \sim 0$ for brevity.

In the case of small Rm, dissipative effects are so strong that $\boldsymbol{b}'$ is effectively stationary and so (Brandenburg & Subramanian, 2005; Krause & Rädler, 1980)

$$\eta \nabla^2 \boldsymbol{b}' = -\boldsymbol{\nabla} \times (\boldsymbol{u}' \times \overline{\boldsymbol{B}}) \tag{1.48}$$

while in the case of Rm $\gg 1$ and/or $S < 1$ (more applicable to astrophysical systems),

$$\frac{\partial \boldsymbol{b}'}{\partial t} = \boldsymbol{\nabla} \times (\boldsymbol{u}' \times \overline{\boldsymbol{B}}). \tag{1.49}$$

In Equation (1.49) above, solutions $\boldsymbol{b}'$ are linearly related to $\overline{\boldsymbol{B}}$ by some relation

$$\boldsymbol{b}' = \boldsymbol{b}'_{ss} + \mathcal{L}(\overline{\boldsymbol{B}}) \tag{1.50}$$

where $\boldsymbol{b}'_{ss}$ is a small-scale fluctuating dynamo field that exists independently of $\overline{\boldsymbol{B}}$ and $\mathcal{L}(\overline{\boldsymbol{B}})$ is a linear operator over the average field $\overline{\boldsymbol{B}}$. We include $\boldsymbol{b}'_{ss}$ for mathematical completeness, however, as it is defined to exist for $\overline{\boldsymbol{B}} = 0$, it should not influence the EMF $\boldsymbol{\varepsilon}$, and hence we may ignore it anyway (Tobias, 2021).

Within the kinematic regime, we assume that $\boldsymbol{u}'$ is independent of $\boldsymbol{B}$, hence $\boldsymbol{\varepsilon} = \langle \boldsymbol{u}' \times \boldsymbol{b}' \rangle$ is also a linear scaling with $\overline{\boldsymbol{B}}$, and can be Taylor expanded as (e.g. Krause & Rädler, 1980; Moffatt, 1978)

$$\varepsilon_i = \alpha_{ij} \overline{\mathbf{B}}_j + \beta_{ijk} \partial_k \overline{\boldsymbol{B}}_j + \dots \tag{1.51}$$

where α_{ij} and β_{ijk} are pseudo-tensors[3] which depend on the statistical properties of the flow, primarily the stratification, angular velocity, and mean magnetic field, although the small-scale magnetic field may also play a role (Brandenburg & Subramanian, 2005). These pseudo-tensors are collectively referred to as the turbulent transport coefficients. The dots $(\dots)$ represent higher-order derivatives in $\overline{\boldsymbol{B}}_j$ that are smaller and hence generally ignored.

The core aim of kinematic mean-field theory under FOSA is, therefore, to solve for the turbulent transport coefficients α_{ij} and β_{ijk} which define the EMF $\boldsymbol{\varepsilon}$ and thus describe the dynamo.

[3]Here, a 'pseudo-tensor' refers to a tensor which is dependent on right-handedness, i.e. those which change sign under reflection

The α_{ij} tensor

Considering first the α_{ij} tensor in isolation, Equation (1.51) above reduces to

$$\varepsilon_i = \alpha_{ij}\overline{\mathbf{B}}_j. \tag{1.52}$$

Under the assumption of homogeneous, isotropic turbulence, and within the regimes covered by FOSA, the α_{ij} coefficient can be reduced to a pseudo-scalar α by

$$\alpha_{ij} = \alpha\delta_{ij}, \tag{1.53}$$

where δ_{ij} is the Kronecker tensor,

$$\delta_{ij} = \begin{cases} 1 \text{ if } i = j \\ 0 \text{ if } i \neq j \end{cases} \tag{1.54}$$

If $S \ll 1$ (i.e. Equation 1.49 applies) then the scalar α takes the form (Charbonneau, 2010; Moffatt, 1978)

$$\alpha = -\frac{\tau_c}{3}\langle \boldsymbol{u}' \cdot \nabla \times \boldsymbol{u}' \rangle, \tag{1.55}$$

where τ_c is the correlation time of the turbulent eddies. Here, α is proportional to the mean flow helicity $\mathcal{H} = \langle \boldsymbol{u}' \cdot \nabla \times \boldsymbol{u}' \rangle$, a point which will become important in later discussion.

If we instead want to consider more complex flows, we may split the α_{ij} tensor into symmetric and antisymmetric parts

$$\alpha_{ij} = \alpha_{ij}^S + \alpha_{ij}^A \tag{1.56}$$

where the symmetric part may be referred to by its principal axis (in spherical polar coordinates) (Hughes & Tobias, 2010)

$$\alpha_{ij}^S = \begin{bmatrix} \alpha_{rr} & & \\ & \alpha_{\theta\theta} & \\ & & \alpha_{\phi\phi} \end{bmatrix} \tag{1.57}$$

and the antisymmetric part given as $\alpha_{ij}^{A} = -\epsilon_{ijk}\gamma_k$, where ϵ_{ijk} is the Levi-Civita skew-symmetric tensor,

$$\epsilon_{ijk} = \begin{cases} 1 & \text{if } (i,j,k) \text{ is } (1,2,3),(3,1,2) \text{ or } (2,3,1) \\ 0 & \text{if } i=j \text{ or } i=k \text{ or } j=k \\ -1 & \text{if } (i,j,k) \text{ is } (2,1,3),(3,2,1) \text{ or } (1,3,2) \end{cases} \tag{1.58}$$

When expressed in this way, the EMF equation (1.52) becomes

$$\varepsilon_i = \alpha_{ij}^{S}\overline{\mathbf{B}}_j + [\boldsymbol{\gamma} \times \overline{\mathbf{B}}]_i, \tag{1.59}$$

and the antisymmetric component of the α_{ij} tensor contributes an additional mean velocity component to Equation (1.40). γ_k and, to a lesser degree, the off-diagonal components of α_{ij}^{S} constitute *turbulent pumping*, which is responsible for expelling mean magnetic field from regions of strong turbulence and is the main transport of magnetic fields out of bulk motions (Brandenburg & Subramanian, 2005). The non-diagonal components of α_{ij}^{S} are generally ignored in accretion studies as they are considered less important in the dynamo process (Bardou et al., 2001; Brandenburg & Donner, 1997).

The diagonal elements of the symmetric component α_{ij}^{S} (or, in the case of isotropic turbulence, α) are known as the α-effect (not to be confused with the Shakura & Sunyaev viscosity term α_{visc}). Of the three components $(\alpha_{rr}, \alpha_{\theta\theta}, \alpha_{\phi\phi})$, the latter is of most interest in accretion disks as it is responsible for replenishing the poloidal field from the toroidal.

The relative strength of the α- and Ω-effects is measured by the ratio of their contributions to (1.45),

$$\frac{r\boldsymbol{B}_P \cdot \boldsymbol{\nabla}\Omega}{(\boldsymbol{\nabla} \times \overline{\boldsymbol{\varepsilon}})_\phi} = \frac{r\boldsymbol{B}_P \cdot \boldsymbol{\nabla}\Omega}{\hat{\mathbf{e}}_\phi \cdot \nabla \times (\alpha \boldsymbol{B}_P)}. \tag{1.60}$$

Where the α contribution dominates over Ω, the α-effect can be considered the primary source of the toroidal field for (1.45). As α is also the primary source for the poloidal field (1.44), this is termed the α^2-dynamo. Conversely, if Ω dominates over α for toroidal field generation, it is known as the $\alpha\Omega$-dynamo.

Both dynamos can act in accretion disks; however, the presence of significant shear through differential rotation means that the $\alpha\Omega$-dynamo model is typically used as the basis for study.

The β_{ijk} tensor

Looking instead briefly at the second tensor in Equation (1.51), we have

$$\varepsilon_i = \beta_{ijk}\partial_k \overline{\boldsymbol{B}}_j. \tag{1.61}$$

Similarly to before, under homogeneous, isotropic turbulence, the tensor β_{ijk} can be expressed as a pure scalar via

$$\beta_{ijk} = \beta\epsilon_{ijk}, \tag{1.62}$$

which, under FOSA, can be calculated as

$$\beta = \frac{\tau_c}{3}\langle(\boldsymbol{u}')^2\rangle. \tag{1.63}$$

When Equation (1.62) is resubstituted into the EMF, it yields

$$\varepsilon_i = -\beta\boldsymbol{\nabla} \times \overline{\boldsymbol{B}}_j \tag{1.64}$$

and hence $\boldsymbol{\nabla} \times \overline{\varepsilon}$ in Equation (1.40) becomes $\beta\nabla^2\overline{\boldsymbol{B}}$. Like terms can then be collected,

$$\frac{\partial\overline{\boldsymbol{B}}}{\partial t} = \boldsymbol{\nabla} \times (\overline{\boldsymbol{u}} \times \overline{\boldsymbol{B}}) + (\eta + \beta)\nabla^2\overline{\boldsymbol{B}} \tag{1.65}$$

showing that the scalar β acts as a turbulent diffusivity (often re-termed η_T) enhancing the molecular diffusivity η.

Quenching α

Early work by Vainshtein & Cattaneo (1992) and Gruzinov & Diamond (1994) suggested that α_{ij} would be catastrophically quenched (and hence the dynamo effectively shut down) once the large-scale field generated by the EMF was sufficiently strong that the Lorentz force back-reaction suppressed the EMF itself. It was initially assumed that the Lorentz force would become significant once the mean magnetic energy $\overline{\boldsymbol{B}}^2 = \boldsymbol{B} \cdot \boldsymbol{B}$ was in equipartition with the flow kinetic energy, leading the α_{ij} tensor to have a $\overline{\boldsymbol{B}}$ dependence which was expressed *ad hoc* as (Hughes & Tobias, 2010; Tobias, 2021)

$$\alpha(\overline{\boldsymbol{B}}) = \frac{\alpha_K}{1 + \overline{\boldsymbol{B}}^2/B_{\text{eq}}^2} \tag{1.66}$$

where α_K is the kinematic scalar determined by Equation (1.55) and $B_{\rm eq}$ is the equipartition value of the magnetic field strength.

However, at high Rm, the energy contained within the small-scale field $\boldsymbol{b}'$ could potentially be more significant than that of the large-scale field, and so the equations may be adjusted assuming $\overline{b}^2 = \boldsymbol{b}' \cdot \boldsymbol{b}' \sim \mathrm{Rm}^q B_0^2$ with the exponent $0 \leq q \leq 2$ (often taken as $q = 1$) being flow and geometry dependent (Tobias, 2021):

$$\alpha(\overline{\boldsymbol{B}}) = \frac{\alpha_K}{1 + \mathrm{Rm}^q \, \overline{\boldsymbol{B}}^2 / B_{\rm eq}^2} \tag{1.67}$$

In most astrophysical systems, $\mathrm{Rm} \gg 1$, and the resulting saturation $|\langle \boldsymbol{B} \rangle|$ suggested by Equation (1.67) is far too low (Charbonneau, 2010). A solution to this issue was found in *dynamical quenching*, whereby the cascade of turbulence to smaller scales leads to a growth in current helicity, adding a new term to Equation (1.55) (Pouquet et al., 1976),

$$\alpha = -\frac{\tau_c}{3}(\langle \boldsymbol{u}' \cdot \nabla \times \boldsymbol{u}' \rangle - \rho^{-1} \langle \boldsymbol{j}' \cdot \boldsymbol{b}' \rangle) \equiv \alpha_K + \alpha_M \tag{1.68}$$

Here, we define the first term in the brackets as the original kinematic contribution α_K. The second term is the addition due to current helicity (closely related to the magnetic helicity) α_M. Under the assumption that $\alpha_K = \text{const}$ and $q = 1$, Equation (1.67) is modified to (Brandenburg & Subramanian, 2005; Gruzinov & Diamond, 1994)

$$\alpha(\overline{\boldsymbol{B}}) = \frac{\alpha_K + \beta \, \mathrm{Rm} \, \langle \boldsymbol{J} \cdot \boldsymbol{B} \rangle / B_{\rm eq}^2}{1 + \mathrm{Rm} \, \overline{\boldsymbol{B}}^2 / B_{\rm eq}^2}. \tag{1.69}$$

We note here the inclusion of our turbulent diffusivity scalar β, although its role is uncertain. It may be assumed to be constant, with a value set by Equation (1.62), or assumed to vary proportionally to $\alpha(\overline{\boldsymbol{B}})$ and hence take on a similar relationship to Equation (1.67).

The impact of α-quenching is uncertain in accretion disks, and in $\alpha\Omega$-type dynamos generally, as in these cases, a toroidal field can be generated by shear without generating much $\langle \boldsymbol{J} \cdot \boldsymbol{B} \rangle$.

In context of the MRI

The MRI complicates the image of the kinematic dynamo presented above. Since the MRI can amplify and sustain the magnetic energy of the disk over time (e.g. Brandenburg et al., 1995; Hawley et al., 1995, 1996; Stone et al., 1996), it is considered to act as a dynamo, but not in a kinematic sense.

The core of the issue is the back-reaction of $\boldsymbol{B}$ onto $\boldsymbol{u}$, particularly from $\boldsymbol{b}'$ (e.g. Kulsrud & Anderson, 1992). In the classical kinematic dynamo theory of Moffatt (1978), $\boldsymbol{u}$ can act upon $\boldsymbol{B}$, but the latter is assumed to be too weak to impact the former significantly, and so Lorentz forces can be ignored. In the context of the MRI, the Lorentz forces are never insignificant; hence the relationship $\boldsymbol{u} \rightarrow \boldsymbol{B}$ appears to be wrong, and there is instead $\boldsymbol{u} \leftrightarrow \boldsymbol{B}$ (Balbus & Hawley, 1998).

Fast and slow, large and small

If the growth rate of the dynamo remains above zero as Rm $\rightarrow \infty$, then the dynamo is said to be *fast*; otherwise, it is *slow*. Dynamos are additionally classified by the length scale over which they act, with *small-scale* dynamos arising from similarly small-scale dynamics (e.g. turbulence, random stretching of field lines) and *large-scale* dynamos being those concerned with larger scale features like density stratification and boundary conditions, often with a dependence on helicity or vorticity (Brandenburg et al., 2012), and which generate fields on scales larger than turbulent eddies.

Herzenberg (Herzenberg, 1958) and Roberts flow (Roberts, 1970, 1972) are two examples of slow dynamos; however, in accretion disks, the microphysical resistivity η tends to be small (Rm large), favouring fast dynamo models. For this reason, slow dynamos will not be covered here, and the reader is directed to Brandenburg & Subramanian (2005) for a review.

An example of a fast dynamo is the Vainshtein & Zel'dovich (1972) folding loop dynamo. In the high-Rm limit (Rm $\gg 1$), field lines are frozen into the fluid flow, and the stretching of these lines leads to exponential growth in field magnitude. The authors suggested that these loops could be formed via convective motions in stars and then rotated and folded by the action of the Coriolis force. A similar process might proceed in accretion disks via shear and turbulent motions.

To see this in more detail, we consider a closed magnetic flux tube with flux

$$\Phi_{B,0} = \int_S \boldsymbol{B} \cdot \boldsymbol{n} dS \tag{1.70}$$

and cross-sectional area S (Fig 1.6). If the loop is stretched to be twice its length, then the total loop magnetic flux doubles, and the cross-sectional area halves as, for an incompressible flow, the volume of the tube is conserved. The stretched loop is then twisted into a figure-8 and folded back onto itself. This folded structure coalesces to the same volume as the initial loop with minimal diffusive effects, but the magnetic flux remains double its initial value. Differential rotation then shears the resultant loop, and the process starts

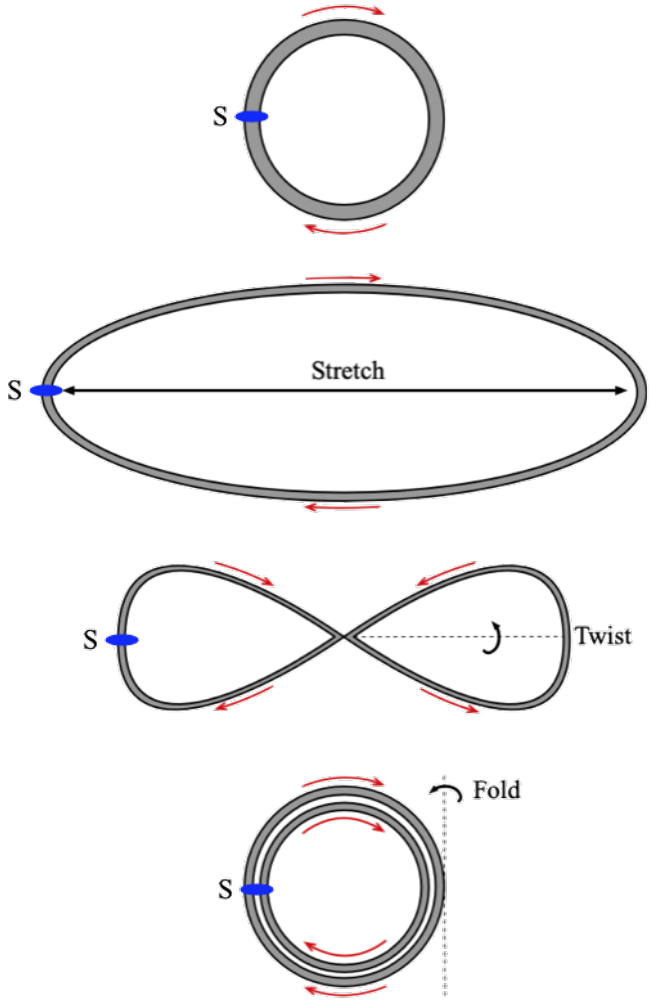


Fig. 1.6 A schematic representation of the Stretch-Twist-Fold (STF) method of field amplification, with appropriate cross-section given by the blue oval. Adapted from Roberts & Soward (1992) and Vainshtein & Zel'dovich (1972).

over. After n repetitions, the flux has grown to $\Phi_{B,n} = 2^n \Phi_{B,0}$ with a growth rate of $\Lambda = n^{-1} \ln (\Phi_{B,n}/\Phi_{B,0}) = \ln 2$ (Roberts & Soward, 1992). This is known as the *Stretch-Twist-Fold* (STF) method.

The STF method does have its limitations. Most notably, it describes a possible pathway to amplify the field but does not propose a cause for the STF motion. Differential rotation is sufficient to perform the 'stretch' action through shear, but it cannot perform the 'twist' and 'fold' actions, which require movement in the vertical direction. For these steps, Tout &

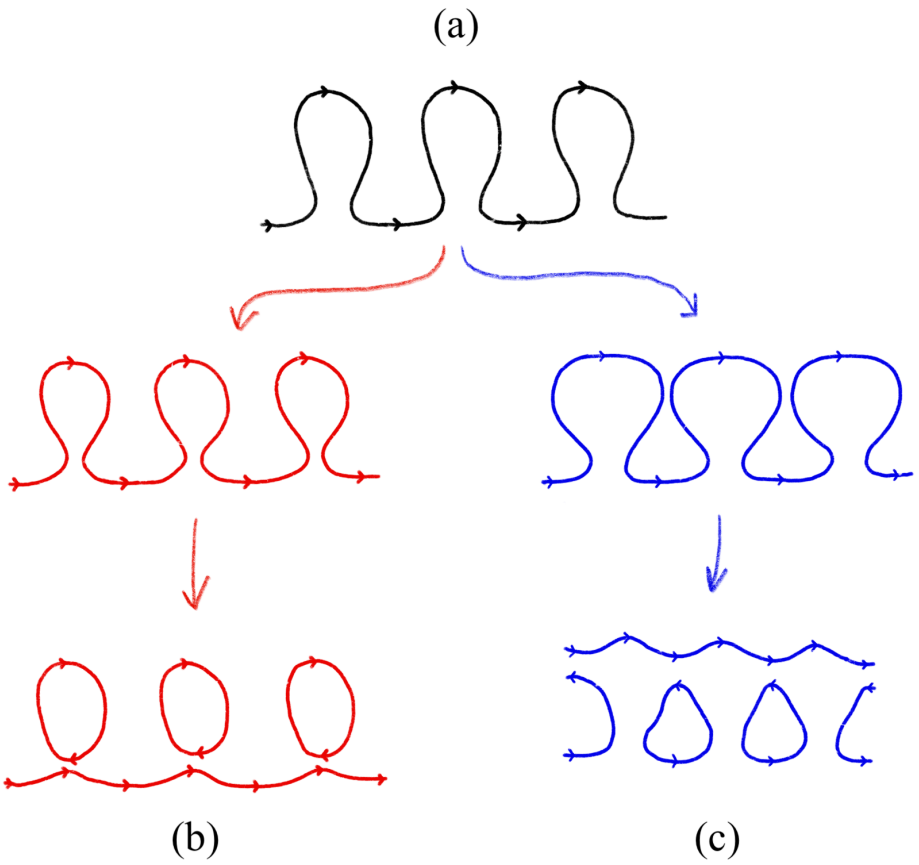


Fig. 1.7 Schematic of the two pathways of the Parker instability (a), which can either connect at their base (b) or through their top (c). Adapted from Parker (1992).

Pringle (1992) suggested that the Parker instability (PI; Parker, 1966, 1992) would play a key role in converting radial and azimuthal field to vertical.

In this instability, shown in Fig. 1.7, magnetic field lines buoyantly rise into arcs, reconnecting and separating from the flow. The vertical fields produced in this process can feed the MRI whilst the PI generates an additional large-scale radial field through loop reconnections, as shown in Keplerian simulations by Johansen & Levin (2008). However, the PI grows at a maximum rate of order $(v_A/c_s)\Omega$, while the MRI grows at $\sim \Omega$, meaning the MRI dominates unless the magnetic field is initially very strong ($v_A \sim c_s$).

The STF method also introduces helicity to the flow. *Writhe* helicity $\mathcal{H}_{\rm wr}$ is related to the number of crossings in the tube (curve C) and is assigned either $W = +1$ or -1 depending on the direction of the underlying flow (Gilbert, 2002).

There is a compensatory *twist* helicity $\mathcal{H}_{\rm tw}$ associated with the field lines inside the tube which takes the opposite sign to the writhe contribution, such that an individual STF cycle has $\mathcal{H} = \mathcal{H}_{\rm wr} + \mathcal{H}_{\rm tw} = 0$. However, over many cycles, the STF method accumulates strong helicity on the largest scales with one dominant sign and the opposite signed helicity on the smallest.

Repeated STF cycles produce twist helicity at ever smaller scales, which acts to resist further bending. The ultimate result is that helicity collects at the smallest scales, severely limiting the growth of the large-scale field and potentially leading to magnetic field decay (e.g. Shukurov et al., 2006; Sur et al., 2007). The large-scale dynamo can be protected if the small-scale magnetic field and its associated helicity are sapped away through some outflow, but if the outflow is too strong, it will also subtract from the large-scale field.

The presence of flow and magnetic helicity in large-scale dynamos generally, and in the construction of α specifically (e.g. Equations 1.55,1.68), points to the necessity of helicity in generating large-scale fields. In a closed system, helicity evolves as (Charbonneau, 2010)

$$\frac{d}{dt}\int \boldsymbol{A}\cdot\boldsymbol{B}dV = -2\eta\int \boldsymbol{J}\cdot\boldsymbol{B}dV \tag{1.71}$$

with magnetic vector potential $\boldsymbol{B} = \boldsymbol{\nabla}\times\boldsymbol{A}$, which in the ideal limit ($\eta = 0$) says that helicity must be conserved.

Under an STF-like picture, helicity of opposing signs collects at different length scales, which we can express by again separating the flow into mean and fluctuating parts (Charbonneau, 2010),

$$\frac{d}{dt}\int \overline{\boldsymbol{A}}\cdot\overline{\boldsymbol{B}}\,dV = 2\int \boldsymbol{\varepsilon}\cdot\overline{\boldsymbol{B}}\,dV - 2\eta\int \overline{\boldsymbol{J}}\cdot\overline{\boldsymbol{B}}\,dV \tag{1.72}$$

$$\frac{d}{dt}\int \overline{\boldsymbol{a'}\cdot\boldsymbol{b'}}\,dV = -2\int \boldsymbol{\varepsilon}\cdot\overline{\boldsymbol{B}}\,dV + 2\eta\int \overline{\boldsymbol{j'}\cdot\boldsymbol{b'}}\,dV \tag{1.73}$$

Hence, the total helicity (the sum of Equations 1.72 and 1.73) is conserved, and non-zero helicity can be collected at opposite scales.

Unfortunately, whilst the codes commonly in use for astrophysical applications (e.g. *Athena++*, *PLUTO*) are constructed to conserve specific quantities and satisfy certain constraints (including $\boldsymbol{\nabla}\cdot\boldsymbol{B} = 0$), they very rarely conserve helicity. Thus, helicity can numerically diffuse or be dissipated. Some work is now being done to correct the issue (e.g. Zenati & Vishniac, 2023), although no helicity-cleaning or helicity-conserving algorithms have yet been implemented in accretion disk studies.

1.3.3 Summary

To summarise, magnetic fields are maintained and grown throughout the Universe via a so-called *dynamo* process. Whilst the general concept of a dynamo is broad and only vaguely defined, we may characterise them by their speed of action and the physical scales they primarily operate over. To understand the fast, large-scale dynamo operating within accretion disks, we turn to two major simplifying theories: kinematic dynamo theory, in which Lorentz forces are ignored, and the mean-field dynamo theory, where we split field and velocity into mean and fluctuating parts that inhabit different length scales. Through the use of these simplifying models, we come to understand the fast, large-scale disk dynamo through tensors α_{ij} and β_{ijk}. Of primary interest is α_{ij}, whose symmetric diagonal components drive the large-scale dynamo. α_{ij} is closely linked to flow and magnetic (or current) helicity within the disk, although the relationship is complicated. Understanding the role of helicity in dynamo behaviour is difficult as many simulation codes do not conserve helicity, although increasing attention is being paid in this area.

1.4 Black hole jets

Approximately 10–20 per cent of AGN are radio-loud (RL), defined as a 5GHz (radio) to B-band (optical) luminosity ratio $\gtrsim 10$ (Kellermann et al., 1989), although the statistics increase with increasing optical and X-ray luminosity (Urry & Padovani, 1995). In some radio-loud galaxies we observe large, collimated outflows in the form of relativistic jets, with simulations suggesting they may be ubiquitous amongst quasars (including those that are radio-quiet, RQ; Macfarlane et al., 2021), merely unresolved. When resolved, jets may appear to extend far beyond the bounds of their host galaxy and as such are thought play an essential role in the form of kinetic feedback onto the galaxy (Figure 1.8).

These structures are thought to be formed close to the central black hole, and two proposed mechanisms exist to explain their formation: the Blandford-Znajek (BZ; Blandford & Znajek, 1977) and Blandford-Payne (BP; Blandford & Payne, 1982) mechanisms. Both concern large-scale poloidal magnetic fields near the black hole and are capable of launching jets with kinetic power $\sim 10^{45-47}\ \mathrm{erg\ s^{-1}}$ for reasonable BH mass, spin, and accretion rate, although each method sources the necessary field (and the energy used to launch the jets) from a different location.

The BP mechanism sources its magnetic field from within a rotating accretion disk, whilst the BZ mechanism sources its field from the rotating BH's ergosphere and horizon and the

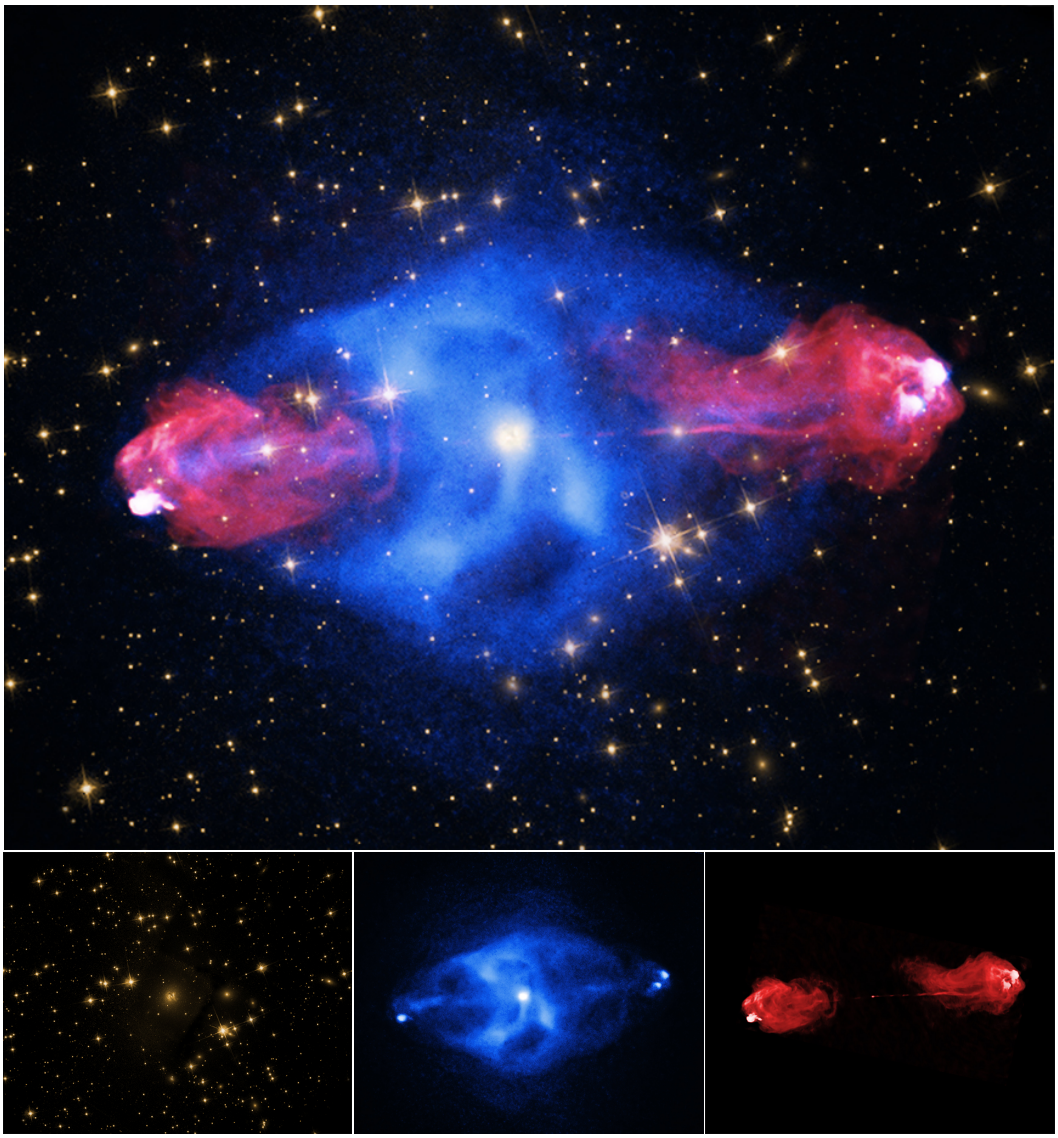


Fig. 1.8 A composite multi-wavelength view of Cygnus A, including optical (yellow), X-ray (blue), and radio (red). Individual components are shown in the lower panels. Image Credit: X-ray: NASA/CXC/SAO; Optical: NASA/STScI; Radio: NSF/NRAO/AUI/VLA.

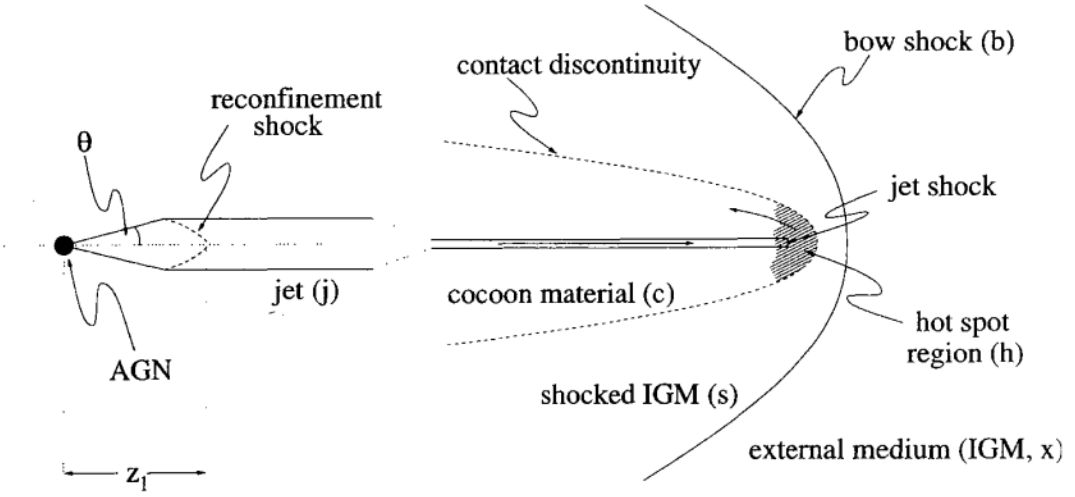


Fig. 1.9 The basic jet components near the BH (*left*) and at termination (for FR-II sources, *right*). Imaged reproduced with permission from Kaiser & Alexander (1997).

BH spin, giving rise to a jet that is dominated by Poynting flux (instead of a magnetically driven wind). Of the two proposed mechanisms, the BZ is preferred as it can theoretically generate more power for an SMBH with maximal spin and retrograde disk motion, however, it is also possible that both processes co-occur (e.g. Xie et al., 2012).

Once launched, the jets are mass-loaded with charged particles of electron/position or electron/proton composition. These electrons (and positrons) range in energy from $\lesssim 100\,\mathrm{MeV}$ to $\gtrsim 100\,\mathrm{TeV}$ (Blandford et al., 2019) and produce radiation with major components contributed from synchrotron radiation in the radio and inverse Compton scattering in the X-ray.

Under the BZ mechanism, the initial velocity of the jet, which we may characterise using the Lorentz factor Γ,

$$\Gamma = \frac{1}{\sqrt{1 - \frac{v^2}{c^2}}}, \tag{1.74}$$

depends on the nature of the mass-loading. Low mass loading can lead to very high velocities $\Gamma \gtrsim 2$–10 with the exact maximum Lorentz factor dependent on the detailed high-energy physics of the loading particles (see e.g. McKinney, 2005a,b), although observations suggest the tail of the distribution may extend as far as $\Gamma \sim 50$ (Lister et al., 2009).

Following launch, jets are expected to take a parabolic shape up to the Bondi radius r_B before being recollimated via a change in the external pressure gradient and transitioning to a conical shape (e.g. Asada & Nakamura, 2012). In the case of more powerful outflows, the

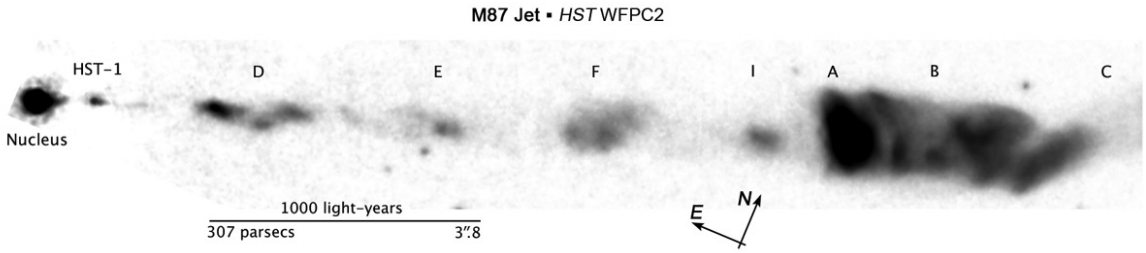


Fig. 1.10 Hubble Space Telescope (HST) observations of the main jet of M87, with knots labelled. Image credit: NASA, ESA, and Z. Levay (STScI/AURA).

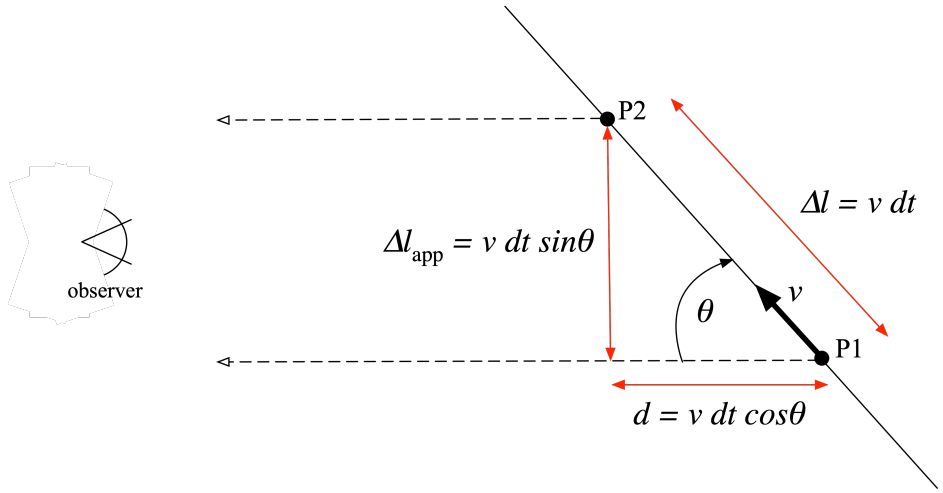


Fig. 1.11 Diagram of apparent superluminal motion of an emitting source moving from point P1 to P2 with velocity v at angle θ to the line of sight.

jets will continue through the ICM/IGM, inflating a bubble of shocked ICM/IGM material, whilst a backflow of jet material fills a smaller cocoon around the jet (Figure 1.9).

Many jets feature 'knots' or blobs of emitting material along their length (Figure 1.10). These knots move along the jet, and by measuring the distance travelled by a knot between observations, an estimate can be made of the jet velocity. However, when such calculations are performed, it is often found that the knots move with apparent velocity $v > c$; this is known as apparent superluminal motion and occurs when a source of relativistic photons (in this case, the knot) is itself moving at a significant fraction of the speed of light close to the line of sight, causing the source to nearly 'catch up' with the photons it emits (Rees, 1966). The effect can be seen diagrammatically in Figure 1.11.

An emitting source travels from point P1 to P2 with a velocity v at an angle θ to the line of sight. In a time dt, the source has moved a real distance $\Delta l = v\ dt$, but in the plane of the sky, this is reduced to $\Delta l_{\rm app} = v\ dt\ \sin\theta$. In moving this length, the distance a signal

must travel to the observer has been decreased by $d = v\ dt\ \cos\theta$, and hence the observer observes signals separated by a time $\Delta t_{\rm app} = dt - d/c = dt[1 - (v/c)\cos\theta]$. Combined, these apparent distances and times give an apparent velocity (Rybicki & Lightman, 1979),

$$v_{\rm app} = \frac{\Delta l_{\rm app}}{\Delta t_{\rm app}} = \frac{v \sin\theta}{1 - \frac{v}{c}\cos\theta}, \tag{1.75}$$

hence if the angle θ is small, and the source velocity v/c large, we may observe $v_{\rm app} > c$.

Relativistic effects also impact the emitted photons themselves. Photons emitted isotropically in the comoving frame are focussed in the direction of movement in the rest (observer) frame. For a relativistic velocity with a Lorentz factor $\Gamma \gg 1$, radiation is beamed into a cone with half opening angle $\theta \sim 1/\Gamma$, and so jets viewed closer to the line of sight receive a luminosity boost whilst those pointing away from the line of sight appear much fainter (or may not be visible at all). For this reason, most observed jets (e.g. Figure 1.14 right) appear one-sided on the parsec scale even though hotspots and lobes are visible on both sides.

1.4.1 Jet morphology

As jets expand beyond the immediate vicinity of the host and into the ICM, they may take on one of two forms depending on the external density/pressure profile, the opening angle of the jet (Krause et al., 2012; Yates-Jones et al., 2021) and jet power (Figure 1.12; Fanaroff & Riley, 1974; Ghisellini & Celotti, 2001), which leads to differences in how energy is deposited. In the first class, characterised by a wide half opening angle $\theta_j \geq 25°$ and low jet power ($\lesssim 10^{25}$ W Hz^{-1} at 1.4 GHz; Fanaroff & Riley, 1974), the jets are quickly decelerated to moderately relativistic speeds at a recollimation-like shock near the BH radius of influence $R_{\rm inf} \sim GM/\sigma^3$, while the electromagnetic energy of the constituent particles is carried out to much higher radii. This early loss of kinetic energy produces jets with a 'plume-like' form, and are termed FR-I or Fanaroff-Riley Type I (Figure 1.13; Fanaroff & Riley (1974)). In the second class, with $\theta_j < 25°$ and high jet power ($\gtrsim 10^{25}$ W Hz^{-1} at 1.4 GHz), the electromagnetic energy is deposited early in the jet expansion, while kinetic energy (i.e., jet velocity) remains highly relativistic until the jet terminates. This second group appear as tear-drop-shaped lobes, often with a terminating hotspot at their edge, and are termed FR-II or Fanaroff-Riley Type II. This classification system is often made quantitative by measuring the location of the brightest part of the jet relative to the jet length through the Fanaroff-Riley

Index (e.g Krause et al., 2012),

$$\mathrm{I_{FR}} = \frac{2x_{\mathrm{bright}}}{x_{\mathrm{length}}} + \frac{1}{2}, \tag{1.76}$$

giving an FR-I classification for $0.5 < \mathrm{I_{FR}} < 1.5$, and FR-II for $1.5 < \mathrm{I_{FR}} < 2.5$. This method works well for the automated classification of straight sources but may struggle with bent sources and otherwise abnormal morphologies or where projection effects are uncertain. In light of this classification scheme, the two classes—FR-I and FR-II—may alternatively be described as edge-darkened or edge-brightened.

As the process of recollimation is thought to depend on external pressure gradients, the pressure profile of the immediate environment is also very likely to play some role in determining jet morphology. The typical luminosity separating line of Fanaroff & Riley (1974) was initially determined based on a severely flux-limited sample. Newer studies by Mingo et al. (2019, 2022) on Low-Frequency Array (LOFAR) data now find a large population of FR-IIs with fluxes extending up to 3 magnitudes below the traditional break, with these ultra-low-luminosity FR-IIs preferentially lying within correspondingly fainter galaxies. This result suggests that it may be the combination of jet power and external pressure profile that determines the morphology.

Within a single source, it is possible, although rare, for the opposing jets to be of different morphology, so-called 'hybrid morphology' radio galaxies (Gopal-Krishna & Wiita, 2000; Harwood et al., 2020). As both jets should be launched with similar initial conditions (e.g. BH spin, magnetic flux), the existence of hybrid sources points further to an environmental cause to jet morphology. There may also exist a subset of sources that exhibit morphologies outside of the two-class FR system, for example, FR-0 sources that show similar core properties to FR-Is but which lack the extended diffuse radio emission associated with jets (e.g. Baldi et al., 2019; Ghisellini, 2011). Additionally, while jet power and opening angle appear as crucial players in determining morphology, it has also been suggested that mass entrainment across the jet boundary (e.g. Bicknell, 1994), stellar mass loading (e.g. Bowman et al., 1996; Laing & Bridle, 2002), direct stellar impact on the jet (Perucho, 2020), centrifugal instabilities (Gourgouliatos & Komissarov, 2018), Kelvin-Helmholtz instabilities(e.g. Perucho et al., 2010), and more, could play a role.

As the jets expand into the ICM/IGM and decelerate, they perform work on the surrounding medium. In the case of FR-IIs, which quite effectively escape the inner galaxy region, much of this work is done in the ICM or IGM through shocks and direct displacement of matter, inflating large cavities or 'bubbles' visible at X-ray wavelengths. Sanders & Fabian

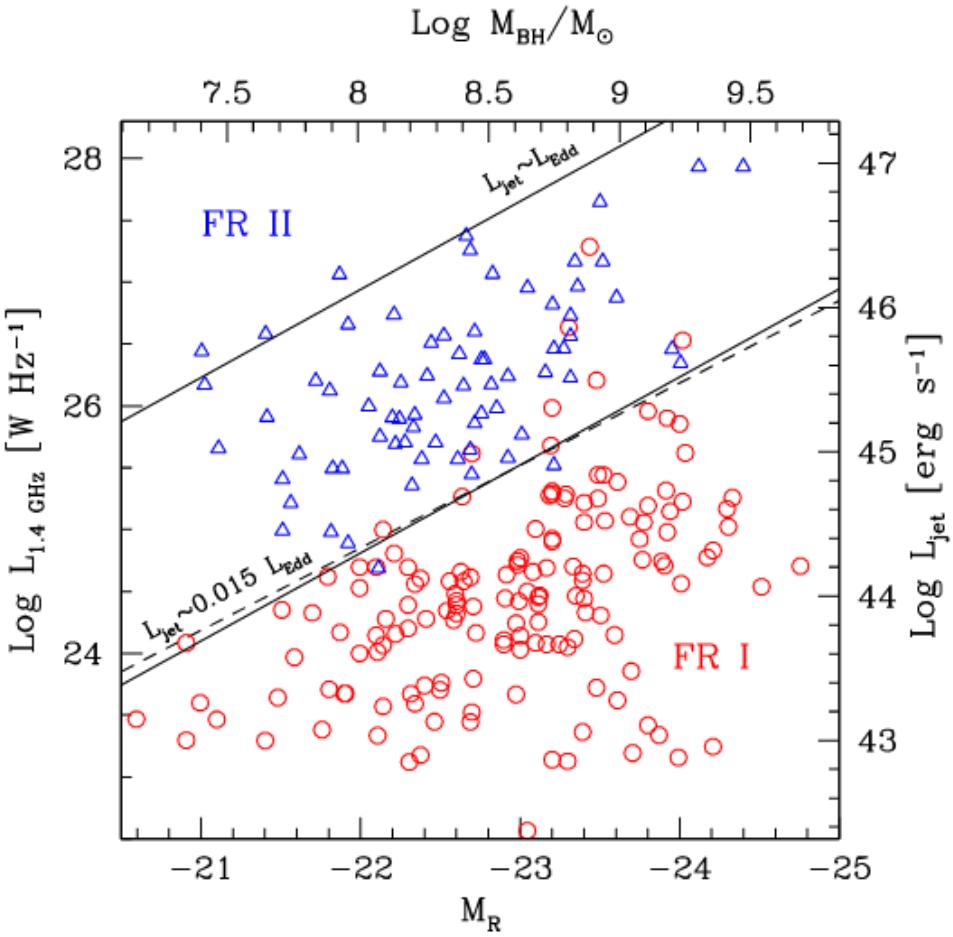


Fig. 1.12 Kinetic jet power (here as $L_{\rm jet}$, right axis) vs host galaxy optical magnitude M_R (bottom axis), with FR-I sources as red circles and FR-II sources as blue triangles. The data can alternatively be expressed in terms of the radio luminosity at 1.4 GHz $L_{1.4\ \rm GHz}$ (left axis) and central black hole mass $M_{\rm BH}$ (top axis), with scaling factors given in Ghisellini & Celotti (2001). The dashed line from Ledlow & Owen (1996) separates the two populations and lies near $L_{\rm jet} = 0.015 L_{\rm Edd}$. Reproduced with permission from Ghisellini & Celotti (2001).

(2007) suggest that 20–40 per cent of the available cavity power may be converted to sound waves which carry energy out into and heat the ICM, although the value may be even higher for the highest power FR-IIs (Perucho et al., 2014).

In addition, the surrounding medium performs work on the jets, too. Correlations between FR-II jet length/luminosity and ambient medium density have been measured both observationally (e.g. Rodman et al., 2019) and in simulations (e.g. Hardcastle & Krause, 2014; Yates-Jones et al., 2021). FR-I sources are generally more symmetric, but environmental impacts cannot be ruled out since FR-I morphology is produced as the jet decelerates

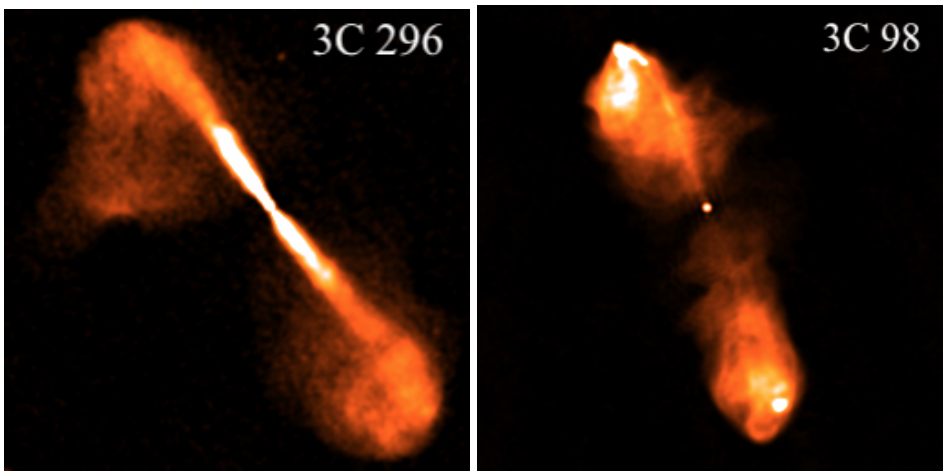


Fig. 1.13 Examples of FR-I (*left*) and FR-II (*right*) type radio jets. Radio images are taken from the online 'Atlas of DRAGNS' at http://www.jb.man.ac.uk/atlas/.

and entrains material near $R_{\rm inf}$. Additionally, FR-I sources have a diffuse structure with luminosity decreasing along the jet length, making extended features hard to observe and highly dependent on telescope sensitivity.

Head-Tail Radio Galaxies

In some cases, sources appear in the sky as having a jet 'tail' (in a 'U', 'V', or 'C' shape) instead of diametric plumes/lobes and are historically termed 'head-tail' (HT) radio galaxies. Head-tail sources have jet powers that straddle the canonical power line between FR-I and FR-II (e.g. Mingo et al., 2019), and the origin of their peculiar shape is widely accepted to be bending by relative movement, i.e. a strong ICM wind or movement of a field galaxy through the cluster. Within this exclusive group, there are two sub-groupings: wide-angle tails (WATs) and narrow-angle tails (NATs), which are defined by the angle between jets (examples in Figure 1.14). The distinction, however, can be murky. Jets with a wide angle may appear narrower when observed at a too-low resolution, and projection effects may also narrow the observed angle between jets artificially.

WATs are are typically associated with dominant, bright D or cD type galaxies within clusters, however, they may also present in groups and filaments and some field WATs are known (O'Dea & Baum, 2023). Many WATs present with thin, well-defined jets in the inner region, which then are disrupted and transition to a more FR-I, plume-like structure, which Sakelliou & Merrifield (1999) suggest could be caused when the jets cross the transition between ISM and ICM. Relative movement of the galaxy to the ICM then provides ram pressure to which the plumes are sensitive, producing the bending observed.

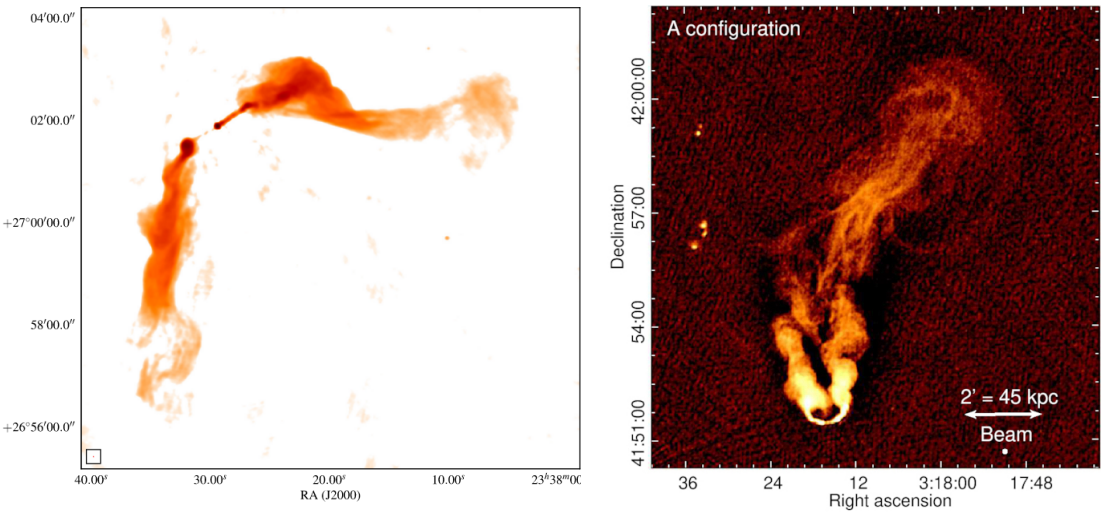


Fig. 1.14 Examples of Wide Angle Tail (WAT) and Narrow Angle Tail (NAT) jets. (*Left*) Archetypal WAT source 3C 465, adapted from Gendron-Marsolais et al. (2020) with permission. (*Right*) Archetypal NAT source NGC 1265, reproduced from Bempong-Manful et al. (2020) with permission.

NATs are also preferentially found in clusters, but their host galaxies are typically not the brightest in the cluster. As they do not reside at the bottom of the potential well, NATs tend to have much higher relative velocities to the cluster centre (Gan et al., 2017), which may account for their more extreme bending.

Where bending is caused by ram pressure from a crosswind, it produces jets bent through 90 degrees at a characteristic bending scale (Jones et al., 2017; O'Neill et al., 2019a),

$$l_b \sim \frac{r_j T_j v_j}{\rho_w v_w^2}, \tag{1.77}$$

for jet radius r_j, momentum density T_j, and velocity v_j and wind density ρ_w and perpendicular velocity v_w.

In the non-relativistic limit, $T_j = \rho_j v_j$, and the above simplifies to

$$l_b \sim r_j \frac{\rho_j v_j^2}{\rho_w v_w^2}, \tag{1.78}$$

and in the further case of a pressure-matched ($p_j = p_w$) supersonic jet, can be expressed in terms of the internal jet and wind Mach numbers (O'Neill et al., 2019a)

$$l_b \sim r_j \left(\frac{M_{j,\mathrm{int}}}{M_w} \right)^2. \tag{1.79}$$

Equation (1.79) also applies for pressure-matched relativistic jets as long as the proper relativistic Mach number is used for the jet.

For a relativistic mass-loaded jet, $T_j = w_j \Gamma_j^2 v_j / c^2$ for jet enthalpy density w_j and Lorentz factor Γ_j, while a Poynting-flux dominated jet would have $T_j = S_j / c^2$ for Poynting energy flux S_j (O'Neill et al., 2019a). When l_b is significantly larger than the jet width, this will give rise to a WAT, while smaller l_b / r_j ratios will produce a source visually similar to a NAT.

1.4.2 Jet magnetic field structure

The two most often proposed launching mechanisms for jets (Blandford-Payne and Blandford-Znajek) both centre around strong, ordered magnetic fields, making this an essential feature in jet formation. Indeed, magnetic fields are not only crucial in the initial launching of the jet but also appear to play an important role in the jet's expansion, confinement, and collimation. As magnetic fields themselves are not luminous, to determine the magnetic structure in and around jets we must instead focus on the interactions between light and magnetic fields through studies of polarisation and Rotation Measure.

The synchrotron radiation produced by a single electron accelerating around a magnetic field line is linearly polarised when viewed edge-on and circularly polarised when viewed face-on; however, due to relativistic beaming, the radiation is focussed into a narrow cone of emission with a half-opening angle $\sim 1/\Gamma$ and so in practice is seen neither fully edge-on nor fully face-on, but is instead elliptically polarised. Radiation from the inner side of the emission cone spirals in a right-handed sense (to the observer), and that from the outer side spirals in a left-handed sense. When an ensemble average is taken over many electrons, the opposite-handed elliptical components cancel out, leaving highly linearly polarised light behind (Rybicki & Lightman, 1979).

We can describe polarisation through the Stokes parameters Q, U, V, and I (Stokes, 1851). Q and U measure the linear polarisation at 90 and 45 degrees, V measures circular polarisation, and I is the total intensity. While I is always positive, Q, U, and V are signed depending on the direction or handedness of the polarisation. For coherent radiation over a narrow frequency range, $I^2 = Q^2 + U^2 + V^2$, otherwise $I^2 < Q^2 + U^2 + V^2$ and we additionally define the polarised intensity $I_p^2 = Q^2 + U^2 + V^2$ and the fraction of polarisation, $p = I_p / I$.

Many lobes contain features which show linear polarisation fractions approaching ∼40–60 per cent (Bridle et al., 1994; Bridle & Perley, 1984), close to the theoretical maximum (75 per cent) for well-ordered fields and a random electron pitch angle. In contrast, the rest

of the lobe sits closer to 5–15 per cent, suggesting significant disordered components to the field as well (Gabuzda, 2021).

The emitted photons are additionally susceptible to changes in their polarisation from the intervening medium they pass through on their way to the observer. Magnetised plasmas are birefringent, and so the dielectric constant ϵ varies for right and left-handed polarisations:

$$\epsilon = 1 - \frac{\omega_p}{\omega_r(\omega_r \pm \omega_B)} \tag{1.80}$$

where ω_p is the plasma frequency, ω_r is the radiation frequency, and ω_B is the cyclotron frequency, $\omega_B = e|\boldsymbol{B}|/m_e c$ for electron charge e and mass m_e, and magnetic field strength $|\boldsymbol{B}|$. The sign in Equation (1.80) above corresponds to right circular polarised light for positive and left circular for negative. Although synchrotron radiation is highly linear (or extremely elliptical), it can be thought of as the combination of a left-circular and a right-circular wave. In the magnetised ICM/IGM, one of these components will travel faster than the other, leading to a rotation of the plane of polarisation χ (Burn, 1966),

$$\chi = \chi_0 + \Delta\chi = \chi_0 + \mathrm{RM}\lambda^2 \tag{1.81}$$

where λ is the wavelength and RM is the *Rotation Measure*, defined in physical units as

$$\mathrm{RM} = \frac{e^3}{8\pi^2\epsilon_0 m_e^2 c^3} \int n_e \boldsymbol{B} \cdot d\boldsymbol{l}, \tag{1.82}$$

for electron charge e, mass m_e, and number density n_e, permittivity of free space ϵ_0, and magnetic field $\boldsymbol{B}$ along the line of sight $\boldsymbol{l}$. The above is often alternatively expressed in cgs units with

$$\mathrm{RM} = 812 \int \left[\frac{n_e}{\mathrm{cm}^{-3}}\right] \left[\frac{\boldsymbol{B}}{\mu\mathrm{G}}\right] \cdot \left[\frac{d\boldsymbol{l}}{\mathrm{kpc}}\right] \ \mathrm{rad\ m}^{-2}. \tag{1.83}$$

Whilst the initial polarisation angle χ_0 cannot be known, the rate of change as a function of wavelength—the RM—can be. Typical RM values for AGN jets vary from $\sim$ 100–10^4, whilst their cores are largely unpolarised (e.g. Zavala & Taylor, 2004).

Early models for the jet magnetic field included the spine and sheath structure of Laing (1980), which features a spine of poloidal field sheathed by an outer toroidal field. However, radio observations by Asada et al. (2002) found transverse gradients in RM across the jet of 3C 273, which the authors interpreted as evidence for a helical magnetic field. Similar gradients have since been found in more than 50 AGN jets at parsec and kiloparsec scales (Gabuzda et al., 2018; Hovatta et al., 2012; Wardle, 2018). Such a consistent, fine-structured

RM signal is unlikely to occur due to chance orientations of the ICM/IGM/ISM, suggesting that the primary source of RM is a thin screen of toroidal field encapsulating the jet material. Within Equation (1.82) only $\boldsymbol{B}$ can change sign and cause a change in sign of RM, a signature that is observed in approximately 50 per cent of transverse RM gradients. However, Carilli & Taylor (2002) argue that extreme RMs in both FR-I and FR-II sources indicate an external Faraday screen associated with large-scale fields in the cluster environment. Additionally, large RMs are observed for background sources viewed *through* an intervening cluster, further suggesting that the high RM values may come from the cluster, not the source.

The direction of the helical/toroidal field supplied by the launching mechanism should be smooth and continuous, otherwise the jets would be disrupted if the launched field exhibited sudden reversals. However, we instead observe apparent reversals in the field direction—both temporally and spatially—in some sources (e.g. Asada et al., 2008; Gabuzda et al., 2018; Mahmud et al., 2009). These reversals can be explained by nesting helical fields of opposite direction. When combined, their effects on the RM partially cancel out; however, as the jet moves and the disturbed, magnetised plasma surrounding the jet changes, a different helical 'shell' can dominate the RM signature, leading to a change in the gradient direction that varies with space and time (Gabuzda, 2021).

Calculating the strength of $\boldsymbol{B}$ and its spatial distribution from the RM is difficult, as Equation (1.82) gives only the cumulative effect, and we must account for additional contributions from the intervening ICM, IGM, and ISM.

Within radio lobes themselves, equipartition arguments give field strength estimates of 20–100 μG (Brunetti et al., 1997) with 100–600 μG in hotspots (Hardcastle, 2001; Hardcastle et al., 2002, 2001), although these values may be underestimates as they do not account for a contribution from relativistic protons (Beck & Krause, 2005).

Outside of the jet itself, if an estimate of the electron density is known, for example, through Hα observations in the ultraviolet (UV) and far-UV (e.g. McClure-Griffiths et al., 2010), then an average $\boldsymbol{B}$ along the line-of-sight can be determined simply by inverting the equation,

$$\langle \boldsymbol{B} \rangle = \frac{8\pi^2 \epsilon_0 m_e^2 c^3}{e^3} \frac{\langle \mathrm{RM} \rangle}{\int n_e d\boldsymbol{l}}. \tag{1.84}$$

However, the intervening magnetic field is turbulent and highly disordered, featuring reversals and large- and small-scale structure, which a simple line-of-sight average fails to capture. The next level of complexity is to split the magnetic field into cells along the line of sight, with the field having the same strength in all cells but a random orientation. Such a distribution will lead to an RM ~ 0 but a dispersion σ_{RM} that goes as the square root of the number of

cells along the line of sight. Visual analysis of the variation in RM across a lobe suggests the magnetic field in clusters may have coherence lengths on the order of 5–10 kpc, providing cluster field strength estimates within a typical range 5–20 μG (Carilli & Taylor, 2002).

Simulations of jets within a magnetised medium generally rely on this kind of description for the external field (e.g. Huarte-Espinosa et al., 2011a,b; Jerrim et al., 2024; Murgia et al., 2004). However, by randomly selecting amplitudes and directions, the resulting field will lack the kind of coherent, large-scale structure that may arise from natural turbulence. If the modelled intervening field is not appropriately modelled, its contribution to the RM may not be accounted for properly, leading to further confusion in measuring the intrinsic jet field.

1.4.3 Radio filaments and threads

Filamentary or 'thread-like' structures are considered to be a natural consequence of MHD turbulence and both HD and MHD fluid instabilities (e.g. Das & Gronke, 2024; Lazarian & Vishniac, 1999; Porter et al., 2015), and have been long been observed within radio lobes (e.g. Dreher & Feigelson, 1984; Fomalont et al., 1989; Owen et al., 2000; Perley et al., 1984; Swain et al., 1998). With the advent of increasingly sensitive low-frequency radio arrays—in particular, the *LOw-Frequency ARray* (LOFAR; van Haarlem et al., 2013) and *MeerKAT* (Jonas & MeerKAT Team, 2016)–such filamentary structure is now being observed more widely, including in the galactic centre (Heywood et al., 2019, 2022) and in radio phoenixes previously thought to be diffuse (e.g. Knowles et al., 2022; Raja et al., 2023). Additionally, a new population of filaments is being unearthed which connect to and extend beyond radio jets/lobes, including Abell 194 (Figure 1.15; Govoni et al., 2017; Knowles et al., 2022; Rudnick et al., 2022; Sakelliou et al., 2008), ESO 137-006 (Ramatsoku et al., 2020), Abel 1314 (van Weeren et al., 2021; Wilber et al., 2019), IC 4296 (Figure 1.15; Condon et al., 2021), 3C 34 and 3C 320 (Mahatma et al., 2023), and Abell 3528S and Abell 3532 (Di Gennaro et al., 2024). Previous observations of some sources at lower resolution showed diffuse 'bumps' in emission where well-defined filaments are now seen; as a result, there are additionally several candidates which would benefit greatly from further imaging, including Abell 2034 (Shimwell et al., 2016), Abell 2255 (Pizzo et al., 2011), NGC 6086 (Candini et al., 2023), Abell 3562 (Giacintucci et al., 2022), and NGC 1272 (Gendron-Marsolais et al., 2021). In the discussion that follows, we will refer primarily to filaments connecting to and extending beyond jets, although some aspects of the discussion may be relevant to other filament populations, particularly those in the galactic centre.

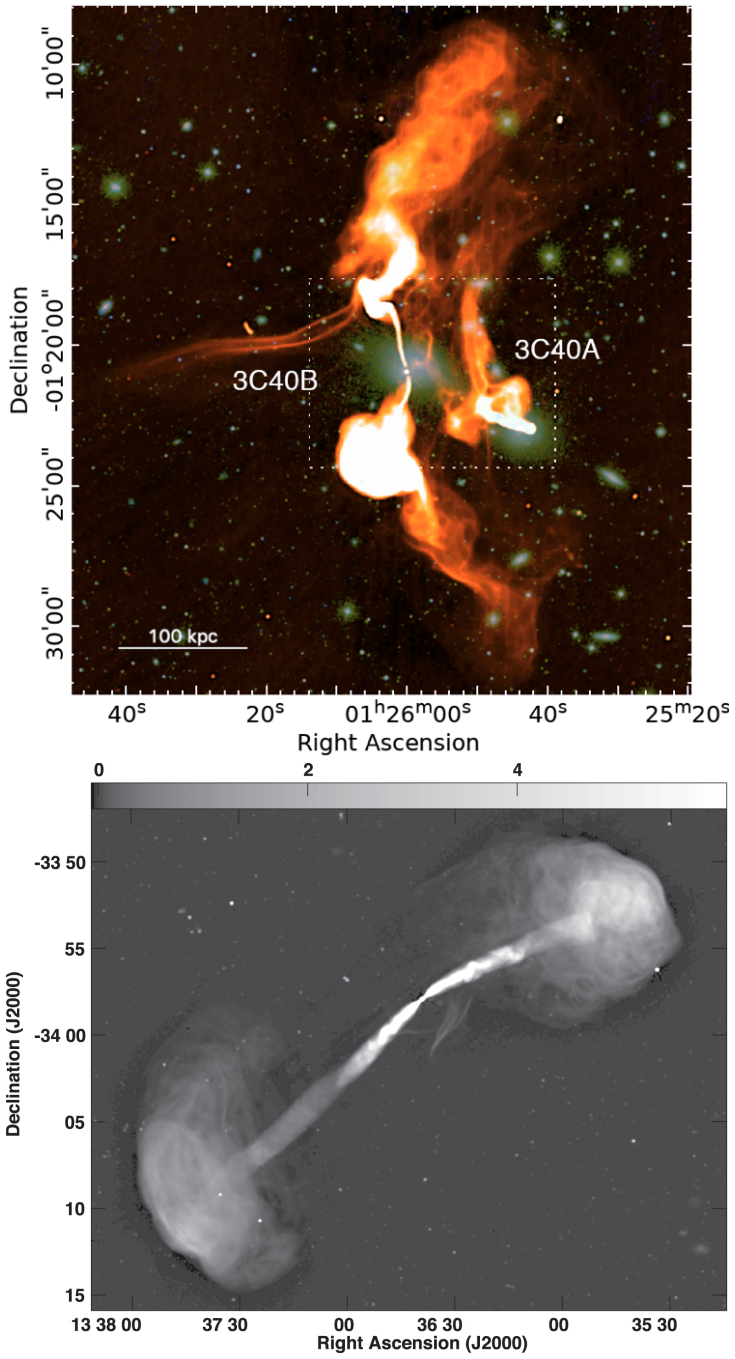


Fig. 1.15 Examples of jet filaments. (*Left*) 3C 40B in Abell 194, featuring a pair of clearly distinct filaments connecting to the upper lobe and extending to the left. Further filamentary structure can be seen within the outer regions of the radio lobes. Reproduced from Knowles et al. (2022) with permission. (*Right*) IC 4296, with a faint bundle of wispy filaments seen to extend from the right jet into the region below the central source, reproduced from Condon et al. (2021) with permission.

Despite their increasing frequency, remarkably little is known of these filaments. As radio emission is synchrotron, we may be reasonably sure that filaments reflect highly energetic electrons gyrating about magnetic field lines, although the source for either component is still hotly debated. Still, as the number of sources with radio filaments increases, several common features can be identified.

The filaments are thin, with typical widths on the order of 1–2 kpc (e.g. Condon et al., 2021; Ramatsoku et al., 2020; Rudnick et al., 2022), and preferentially appear in multiples, bundles, or with significant sub-filamentation. The magnetic field within the filaments is highly linearly polarised, 50–70 per cent, suggesting a highly ordered field, and has a typical strength of a few μG (Yusef-Zadeh et al., 2022). The magnetic energy density is less than or comparable to that of the surrounding ICM (e.g. Condon et al., 2021), and the magnetic pressure may be sufficient to expel X-ray emitting gas from around the filaments in some cases (e.g. Rudnick et al., 2022).

The radio spectral index of jet filaments is typically steep, between 1–2, although higher indices have also been recorded. The spectral index in the filament matches that of the jet or lobe at the apparent intersection point, strongly suggesting that it is a true intersection and not a chance line-of-sight alignment. For those filaments which appear disconnected on one end (e.g. 3C 40B; Rudnick et al., 2022), the spectral index tends to increase with distance from the lobe, reflecting an older and less energetic electron population, although some filaments have been observed with relatively flat spectral index (Condon et al., 2021). The current 'endpoint' of such disconnected filaments is almost certainly sensitivity-limited, with more diffuse LOFAR observations suggesting that filaments continue for some distance further (Rudnick et al., 2022).

Concerning the radio jets/lobes themselves, many sources appear disrupted, bent, or otherwise of complex morphology and are primarily hosted in BCGs, although selection effects will play a role. As a result, the presence of filaments may suggest recent merger activity, the relative motion of a galaxy through a cluster, or other sources of turbulent ICM 'weather' (Brunetti & Jones, 2014). It is unknown if these are preconditions for the formation of all filaments or merely *highly luminous* filaments.

The source of the filament electrons and magnetic field is an open question. Electrons may be sourced directly from the jet, or they may be 'fossil' ICM electrons that have been reaccelerated through interaction with the passing jet. Diffuse radio relics and radio halos are thought to involve the reacceleration of ICM electrons through turbulence or by passing shocks (e.g. ZuHone et al., 2021a,b), although it is not clear that such processes can produce filamentation as finely structured as now observed. The source of the magnetic

field is additionally unknown, although the high degree of coherence has been interpreted as indicative of a large-scale ICM field (e.g. Condon et al., 2021). The expected typical field strength in the ICM is much lower than the inferred field strength in the filaments (Yusef-Zadeh et al., 2022), however, it may be amplified through a turbulent dynamo, a condition likely met in many of the recorded cases. If the field is sourced from the ICM and the electrons from the jet, then there must be magnetic reconnection between the jet magnetic field and the ICM magnetic field, potentially providing another pathway through which relativistic jets can interact with the nearby environment.

1.4.4 Summary

To summarise, jets are a striking feature of many radio-loud AGN which can extend far beyond the boundaries of the AGN host galaxy and likely play an essential role in AGN feedback. Launched by twisted magnetic fields and loaded with charged particles, the jet quickly collimates and generally takes on one of two main shapes—the wispy, meandering FR-I or the powerful FR-II—although exceptions exist. When exposed to a strong external wind, these jets may bend to an extreme degree and appear as head-tail sources, further defined by the angle between the jets as WAT or NAT. The magnetic field within the jet is challenging to study, although some is known due to observations of polarised light and its depolarisation to the intervening magnetised plasma, measured through the Rotation Measure, or RM. From these studies, we may be reasonably sure that the jet contains a significant ordered toroidal or helical component with equipartition field strengths of order 20–100μG. Estimating the RM contribution from the foreground medium is more complicated and relies on substantial simplifying assumptions for the magnetic field distribution, potentially leading to greater confusion in interpreting jet RM values and variance. With the newest generation of radio observatories, *LOFAR* and *MeerKAT*, a new population of radio filaments has been discovered linking to radio lobes, characterised by steep radio spectral indices. While little is known for certain of how these filaments form, or from where they source their relativistic electrons and magnetic fields, they promise to play an important role in furthering our understanding ICM fields and of ICM-jet interactions.

1.5 Numerical Challenges

Just as observational facilities must deal with issues of sensitivity and resolution of their telescopes, simulations must also contend with a number of inherent challenges.

The use of any grid-based method first requires discretisation of the physics equations, in this case the Navier-Stokes equations. This discretisation occurs in two spaces—physical space and time—and each has consequences.

In physical space, an appropriate finite domain must be defined and then subdivided into smaller units—the cell. The cell size dictates the finest detail and structure that can be resolved by the simulation, and as such the ideal resolution is problem-dependent. Too-low resolution often has resounding effects on the subsequent evolution of the system being studied. In the case of accretion disks in particular, a too-low resolution risks under-resolving the MRI modes, with Sano et al. (2004) finding that MHD turbulence would decay under grid diffusion if the fastest growing MRI mode were resolved with less than six cells, leading the magnetic field to saturate far below prediction. Energy saturation levels in driven turbulence also vary with resolution, with higher resolutions saturating at higher energy, most likely due to a complex interplay between shock dynamics and the dissipation of small-scale turbulence at the grid scale (e.g. Stone et al., 1998)

Discretizing in time also introduces a new minimal requirement–the Courant–Friedrichs–Lewy (CFL) condition. In essence, the CFL condition requires that the timestep at step n, $\delta t = t^{n+1} - t^n$, of computation be small enough such that the fastest wave in the system crosses less than one cell distance in that time, characterised under the CFL number $C_0 < 1$. In 1D, this requirement is given as (Stone et al., 2008),

$$\delta t = \frac{C_0 \delta x}{\max(|v_{x,i}^{n+1}| + C_{fx,i}^{n+1})} \tag{1.85}$$

where δx is the cell size in the x-direction, $v_{x,i}^{n+1}$ is the x-direction, cell-centred velocity at cell location i and timestep $n+1$, and $C_{fx,i}^{n+1}$ is the fast magnetosonic velocity in the same direction, location, and time. In the accretion system and non-relativistic MHD, the fastest wave is the fast magnetosonic wave. For relativistic jets and special relativistic MHD, it is the speed of light $c = 1$ and so the requirement simplifies to $\delta t = \delta x/2$ (Qamar & Warnecke, 2005). Different combinations of time-stepping and reconstruction algorithms will, however, place more stringent requirements on the CFL number, most commonly $C_0 \lesssim 1/N_{\rm dim}$.

Further problems are introduced as information is passed from one cell to another. Codes like *Athena++* and *PLUTO* employ finite volume methods, in which the quantities under consideration are defined at the centre of each grid cell and are updated according to the flux through the boundaries of the cell. As a result, the method used to reconstruct the fluxes is of critical importance, and, much like using the Trapezoidal rule for integration, the fewer steps used in the reconstruction of the flux, the less accurate the solution. More

accurate methods can involve more steps or take advantage of characteristics in the MHD equations (so-called 'upwind' methods). In practice, 'lower order' schemes have the effect of smearing the simulation data (i.e. are inaccurate), while 'higher order' schemes tend to produce oscillations, particularly near discontinuities and shocks (i.e. are unstable). In this setting, we introduce the concept of a *limiter*, which allows one to switch between a high-accuracy, low-stability method and a low-accuracy, high-stability method where needed. This is a useful tool in simulating relativistic jets, where shocks are expected, with test cases by Beckwith & Stone (2011) finding that highly magnetised ($\beta \lesssim 10^{-2}$), fast ($\Gamma_{\text{Alfvén}} \gtrsim 5$) flows suffer from grid-scale artefacts under the most accurate solvers.

Within a standard cubical simulation domain, for any astrophysical object there will be regions in the domain of high activity (for example, high density) and regions of low activity (low density), and hence the minimum resolution requirement may differ greatly over the domain. If a uniform grid is used in these situations, much computational time and energy are wasted on unimportant regions. To address this issue, we turn to static or adaptive mesh refinement (SMR and AMR, respectively). In static mesh refinement, levels of increased resolution are defined ahead of time and the grid remains static for the duration of the simulation. This method is useful for systems with well-defined and consistent geometry, for instance, a standard accretion disk. It also finds value in grids with problematic geometries, such as spherical polar coordinates. In Chapter 2, we use SMR to de-resolve the grid in the polar region, as in spherical polar coordinates the cells in this region would otherwise become very small and the corresponding CFL condition prohibitive. In adaptive mesh refinement, on the other hand, one sets a per-cell refinement condition, which, when triggered, increases the resolution locally. While considerably more expensive than SMR, AMR is useful in situations where the geometry of a system is changing and not easily predicted, for example, in the case of a tilted disk. Regardless of the method used, care must be taken in refining the grid as a step-change in resolution can introduce artefacts to the fluid flow and so generate spurious signals—an issue we will see demonstrated in Chapter 2.

Grid refinement as a tool exists to address one issue that arises due to the vast dynamical range of many astrophysical systems, namely the varying resolution requirements. There are, however, other ways in which dynamical range can impact simulations. In particular, discretization leads to rounding errors in arithmetic. In extreme cases, where densities or energies are very low or velocities very high ($\sim c$), these rounding errors can lead to unphysical values, usually in the form of negative pressures or energies, or velocities $|v|^2 > c$. Several techniques exist to address these issues. For negative pressures, the pressure can instead be calculated from the entropy, guaranteeing a positive value although breaking

energy conservation in the process and potentially leading to unphysical velocities (Noble et al., 2009). Negative densities and pressures can also be solved by setting floor values $\rho_{\rm floor}$ and $P_{\rm floor}$ such that any $\rho < \rho_{\rm floor}$ ($P < P_{\rm floor}$) is replaced by the floor value (e.g. Beckwith & Stone, 2011), however, this process also violates conservation laws by effectively creating new matter at the grid scale. The appropriate method/s used to handle such errors depends on the flexibility of the code, the frequency of errors, and the amount of matter or energy which is added to the grid; as such, it must be determined on a case-by-case basis and by trial.

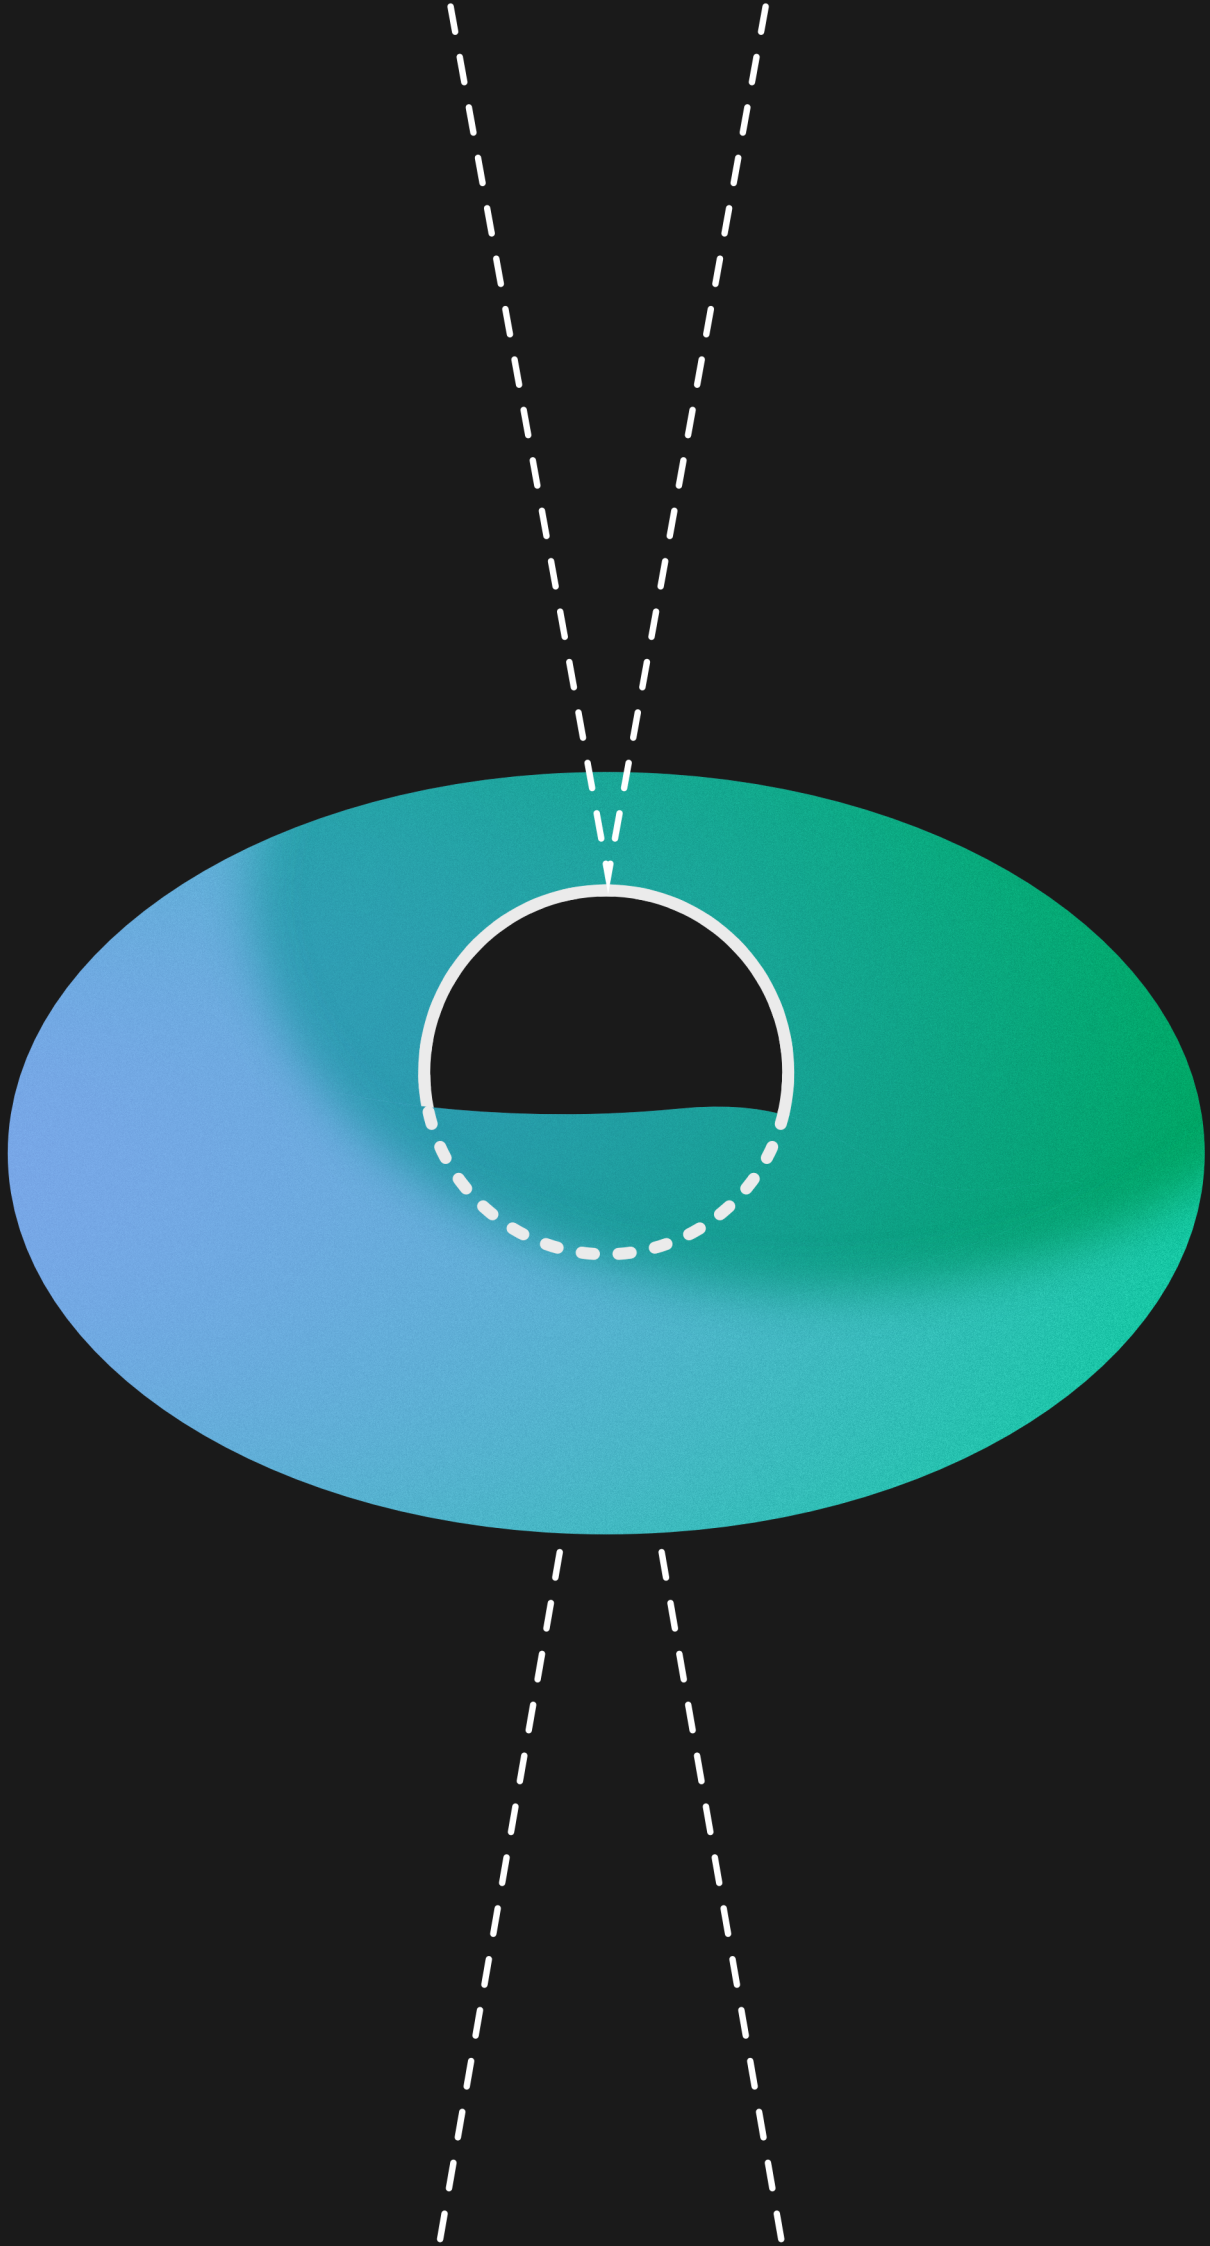
On a practical level, issues also arise in data management and processing. To study natural phenomena, multiple variables (e.g. density, pressure, magnetic field vectors, velocity vectors, and more) must be saved to file at appropriately regular intervals to capture the physics of interest. In the case of a simulation that is run in batches, as opposed to being run continuously, then outputs must include all evolving variables in order to necessitate a restart. As the resolution increases, though, and additional physics beyond ideal hydrodynamics is considered, the size of the resulting files can become quite large. For long-time-scale simulations with many outputs like those presented in Chapter 2, this quickly leads to an enormous amount of data (in our case, $\approx$80 TB) that must be stored somewhere at significant cost. Processing such large amounts of data also introduces its own problems: as memory is limited, analysis must be parallelised and the overhead associated with opening and reading files becomes limiting.

In combination, the abovementioned issues lead to a complex situation in which the ideal physics setup for a given problem is often at odds with the computational resources available, and concessions must be made, either in the resolution or in the amount of non-ideal physics included. The true difficulty in simulation is the art of knowing when and where simplifying assumptions can be made and their resulting impact on the physics. However, while individual algorithms (e.g. solvers and limiters) have been well-tested on a suite of simple test problems, their performance in more complex problems and with more complex initial conditions is more challenging to measure. Regardless, the impact of simulation methods on the final simulation result cannot be overstated, a fact that becomes clear in our studies of accretion disks (Chapters 2 and 3) and jets (Chapter 4).

1.6 Thesis outline

The rest of the thesis is structured as follows. Chapter 2 is a reproduction of a published work describing magnetic field evolution within simulated thick accretion disks, specifically looking at large-scale poloidal magnetic field generation from an initially toroidal morphology.

A small extension of this work is presented in Chapter 3, analysing the dynamo coefficient α_{ij} terms of our simulated disks. We reorient in Chapter 4, turning to look at simulations of bent jets and the complex interactions between jet and ICM magnetic fields. Finally, we draw these results together for our conclusion in Chapter 5.



Chapter 2

Evolution of the magnetic field in high- and low-β disks with initially toroidal fields

As appeared in *Evolution of the Magnetic Field in High- and Low-β Disks with Initially Toroidal Fields*, **Rodman, P. E.**, Reynolds, C. S., 2024, ApJ, 960, 97

2.1 Introduction

Accretion disks are highly turbulent and complex astrophysical systems formed by infalling gas circularising around a massive central body such as a young star, a white dwarf, a neutron star, or a black hole. In the case of accretion onto a supermassive black hole, the disk physics governs both the growth of the black hole and the form of the resulting energy output and thus has wide-reaching implications not just for the immediate environment but on galaxy scales (e.g. the M–σ relation; King, 2003). The strong gravitational forces, relativistic velocities, and magnetic fields present make these accreting black hole systems an invaluable laboratory for studying a wide range of extreme physical processes, including the properties of strong gravitational fields, particle acceleration, and the interaction of magnetic fields and fluid turbulence, the last of which forms the primary focus of this paper.

For an accretion disk to actually accrete it must find a way to remove angular momentum from orbiting fluid elements. This could, in theory, occur through local exchanges between gas parcels within the disk and/or through a wholesale loss of angular momentum from the system via, for example, magnetic braking from a disk wind. Given the extremely

high Reynolds numbers of most astrophysical disks, intrinsic microscopic viscosity can be immediately ruled out as an angular momentum transport mechanism as it would drive accretion over timescales that are far too long to explain the observed luminosity of disks (see Balbus, 2003, for derivation). The seminal work of Shakura & Sunyaev (1973) acknowledged that a form of anomalous viscosity could exist as a result of turbulence, potentially related to the presence of magnetic fields, but the mechanism for generating this turbulence was then unable to be specified. Simply invoking a very high Reynolds number as a driver of turbulence is not sufficient as Keplerian disks are linearly and non-linearly stable to shearing instabilities even at very high Reynolds numbers (Balbus et al., 1996; Hawley et al., 1995, 1996). Later work by Balbus & Hawley (1991, 1998) identified a powerful magnetohydrodynamic (MHD) instability which can drive the necessary turbulence. This magneto-rotational instability (MRI) has since become the paradigm for accretion disk theory.

Magnetohydrodynamic (MHD) and (more recently) General Relativistic Magnetohydrodynamic (GRMHD) simulations are now common tools for studying the internal dynamics of the disk due to the highly non-linear nature of the MRI-driven turbulence (see e.g. Balbus, 2003; Davis & Tchekhovskoy, 2020, for review). In spite of much attention in recent years, however, many questions still remain about the evolution of magnetic fields within accretion disks and the nature of the turbulence that is formed. Equally importantly, uncertainties remain about the influence of computational realities (finite spatial resolution, memory of initial conditions in a limited temporal span, and non-helicity preserving MHD algorithms) on the outcome of accretion disk simulations.

Choices in the initial vertical magnetic field were shown to have important implications for turbulence and overall disk evolution even in early simulations (Hawley et al., 1995), with turbulent stresses increasing with increasing net vertical field. Later studies of both shearing box (e.g. Bai & Stone, 2013; Salvesen et al., 2016) and global (e.g. Mishra et al., 2020) simulations showed that the presence of an initially vertical field affects the structure of the simulated accretion flow. If the vertical field in the inner regions of the disk is sufficiently strong, it may cause the disk to enter a Magnetically Arrested Disk (MAD; Narayan et al., 2003) state in which mass accretion is suppressed or even temporarily halted due to magnetic pressure and, if the central black hole is spinning, jets are launched via the Blandford-Znajek mechanism (Blandford & Znajek, 1977; Tchekhovskoy et al., 2011). MADs have been invoked to explain the low-luminosity and powerful jet in M87 and GRMHD models of MADs can quantitatively match the horizon-scale mm-band image obtained by the Event Horizon Telescope (Event Horizon Telescope Collaboration et al., 2019). The counterpart to MAD is the Standard And Normal Evolution (SANE) state, which is effectively characterised

by the absence of a MAD state. Recent work by Begelman et al. (2022), however, suggests that while the presence of a *net* vertical field is important, it plays a secondary role to that of the toroidal field, with the toroidal component being critically important in the saturation of magnetic flux. To add further complexity to this discussion, the net vertical magnetic flux may be zero on a global scale but *non-zero* locally (Sorathia et al., 2010). A zero-net-vertical-flux (ZNVF) global disk may well be composed of many net-vertical-flux (NVF) local 'shearing box' patches, magnetically connected at high altitudes through the low-density corona of the disk.

In addition to questions on the role of vertical fields and fluxes, the final state of a simulation appears to depend significantly on the particular initial field morphology used, even where the initial net vertical field is kept constant. Early work by Machida et al. (2001) suggested that the initial magnetic field configuration could affect the mass outflow rate in Radiatively-Inefficient Accretion Flows (RIAFs), disks in which the radiative cooling is negligible and local cooling driven by advection—such disks are generally geometrically-thick and have low-accretion rate. Such flows lack a large separation of scales between their inflow/viscous times, thermal times, and dynamical times and so seem to retain memory of their initial conditions during the inflow (White et al., 2020). Further, since they lack radiative cooling, the sign of the Bernoulli parameter (and hence the degree to which regions of the disk are unbound) is sensitive to both the feeding prescription and the detailed internal energetics associated with the initial field. Later simulations by Narayan et al. (2012) found that a MAD state could be produced by varying the number and direction of poloidal loops in the initial magnetic field, with a single-loop configuration being most conducive to building up vertical field necessary for the MAD state. White et al. (2020) then built upon this work, showing that within SANE models variations in initial magnetic field configuration have strong effects on the energetics and structure of RIAFs, with smaller initial poloidal loops generating a less coherent radial field (leading to less horizon-penetrating magnetic flux) and producing polar inflows while larger poloidal loops showed the opposite. Similar work by Jiang et al. (2019) on thin, radiation-dominated disks also found factors of three differences in accretion rate and vertical disk structure depending on whether the initial poloidal field was dipolar or quadrupolar, while Mishra et al. (2022) found that thin RIAFs initialised with dipolar poloidal fields were thermally unstable and those with quadrupolar oscillated between stable and unstable states.

These combined results clearly demonstrate that the properties of simulated accretion disks are dependent on the artificial and somewhat arbitrary initial magnetic conditions used

in their construction, and such a dependence raises serious questions as to the validity of simulations as emulations of real astrophysical disks.

Magnetic fields with net vertical flux or with significant large-scale vertical/polar components play an important role not only for the emergence of MAD states, but also for the production of astrophysical jets through the Blandford-Znajek (Blandford & Znajek, 1977) and Blandford-Payne (Blandford & Payne, 1982) processes. An open question remains whether, and under what conditions, a disk can eventually forget the details of the initial field configuration and grow a self-consistent large-scale field via dynamo action given sufficient time. Fragile & Sądowski (2017) found that initially strongly magnetised disks (where the ratio of thermal-to-magnetic pressure is only $\beta = P_{\rm gas}/P_{\rm mag} \sim 10$) with purely toroidal fields were not locally self-sustaining, albeit the simulations were run over a short timescale ($700\ GM/c^3$) that may not have allowed enough time for poloidal growth. Global simulations by Liska et al. (2020) run for much longer times ($1.3 \times 10^5\ GM/c^3$) successfully generated a strong enough poloidal field *in situ* to enter the MAD regime, but only from an initially strong field ($\beta = 5$), which is likely not reflective of initial conditions in nature.

Given the long times needed for a disk with an initially weak toroidal magnetic field to reach equilibrium out to reasonable radii, most simulations with initially toroidal fields are initialised with a low magnetic plasma β (i.e. a 'strong' magnetic field), meaning little is known of the amplification process for $\beta > 100$ and few simulations cover the intersection of weak and toroidal initial field configurations.

In this work, we study the ability of simulated disks initialised with a purely toroidal field to grow a large-scale poloidal field. We agree with and extend the previous findings of Liska et al. (2020), finding large-scale field growth for disks initialised with either a strong ($\beta = 5$) or a weak ($\beta = 200$) magnetic field. The core of this work is based on two long-duration simulations with initially toroidal fields, a strong-field ($\beta = 5$) and a weak-field ($\beta = 200$) case. Our longest simulation (the weak-field case) is run out to $t \sim 4 \times 10^5\ GM/c^3$ in order to capture the long-term evolution of the magnetic field. Given our focus on understanding the influence of numerics on results, we also run two additional simulations that have identical setups except for higher resolution. We see evidence of poloidal flux generation in all cases, with the weaker initial field case saturating with weaker and more disorganised poloidal field. All of our model disks display dynamo-cycles in the toroidal field even though the disks are geometrically thick. We also identify the subtle ability of internal mesh-refinement boundaries to drive vorticity-induced density asymmetries.

The basic equations of the problem and computational setup are given in Section 2.2. The analysis of our four simulations, with special emphasis on the growth of large-scale

poloidal field, are presented in Section 2.3. We discuss and summarise our findings with our conclusions in Section 2.4.

We make use of the scale-free nature of accretion systems, with quantities scaled by the gravitational radius $r_g = GM/c^2$ and the speed of light c. As a result, physical units can be derived and a physical distance set by defining a specific black hole mass M and the units for c.

2.2 Numerical model

2.2.1 Foundational equations

Our accretion disk models are evolved using the MHD code *Athena++ v21.0* (Stone et al., 2020) on a spherical polar grid (r, θ, ϕ), with the piecewise linear method (PLM) for spatial reconstruction, second-order (VL2) time-stepping, a Harten-Lax-van Leer discontinuities (HLLD) Riemann solver, and a Courant, Friedrichs, & Lewy (CFL) number of 0.3. The code is parallelized over 1008 Cascade Lake cores on the Cambridge Service for Data Driven Discovery (CSD3) system *Peta4*, achieving a full-node performance of $\sim$ 21–22 MZone-cycles/core-second.

Athena++ solves the equations of ideal MHD, given in conservative form as (Stone et al., 2008)

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}) = 0, \tag{2.1}$$

$$\frac{\partial \rho \boldsymbol{v}}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}\boldsymbol{v} - \boldsymbol{B}\boldsymbol{B} + \boldsymbol{P}^*) = -\rho \nabla \Phi_{\rm PW}, \tag{2.2}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot [(E + P^*)\,\boldsymbol{v} - \boldsymbol{B}\,(\boldsymbol{B} \cdot \boldsymbol{v})] = 0, \tag{2.3}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} - \nabla \times (\boldsymbol{v} \times \boldsymbol{B}) = 0, \tag{2.4}$$

where $\boldsymbol{P}^* = (P + B^2/2)\boldsymbol{I}$, $\boldsymbol{I}$ is the diagonal unit tensor, $B^2 = \boldsymbol{B} \cdot \boldsymbol{B}$, and all other variables having their usual meaning. The total energy density is given by

$$E = \frac{P}{\gamma - 1} + \frac{1}{2}\rho v^2 + \frac{B^2}{2} + \rho \Phi_{\rm PW}, \tag{2.5}$$

where we have imposed a γ-law equation of state relating the internal energy density $\mathcal{E}$ to the pressure, $P = (\gamma - 1)\mathcal{E}$. We use the pseudo-Newtonian Paczyński–Wiita gravitational potential, $\Phi_{\rm PW}$,

$$\Phi_{\rm PW}(r) = -\frac{GM}{r - 2r_g}, \qquad r_g \equiv GM/c^2, \tag{2.6}$$

which features an innermost stable circular orbit (ISCO) at $r_{\rm ISCO} = 6r_g$ and is an adequate approximation (within Newtonian physics) of the gravitational field around a non-rotating black hole. The solenoidal condition on the magnetic field,

$$\nabla \cdot \boldsymbol{B} = 0, \tag{2.7}$$

is enforced by upwind constrained transport (CT; Evans & Hawley, 1988).

The system of equations 2.1–2.7 is closed and may be solved numerically. During the numerical solution, we enforce density and pressure floors at $\rho_{\rm floor} = 3 \times 10^{-5}$ and $P_{\rm floor} = 1 \times 10^{-10}$. Due to the high Alfvén speeds that arise in the low-density regions of our `b5_hi` run, we increased the density floor in this simulation to $\rho_{\rm floor} = 3 \times 10^{-4}$ at $t \sim 2 \times 10^4\, GM/c^3$, representing a 0.16% increase in the total mass at that moment which does not appreciably affect our results.

2.2.2 Mesh, boundary, and initial conditions

Mesh and boundary construction

Our computational domain spans $r \in [5r_g, 290r_g]$, $\theta \in [0, \pi]$, and $\phi \in [0, 2\pi)$. Our fiducial pair of simulations have a base resolution of $(N_r, N_\theta, N_\phi) = (112 \times 32 \times 32)$ and 3 additional levels of static mesh refinement (SMR), with each level double the resolution of the one below (Figure 2.1). This gives us a maximum effective resolution of $(N_r, N_\theta, N_\phi) = (896 \times 256 \times 256)$ for $r < 100r_g$ and height $\lesssim 2H$, resulting in 20 cells per scale height H. Using SMR in this way, our resulting mesh (Figure 2.1) is de-resolved near the poles, increasing cell sizes in this region and avoiding untenably small timesteps that would arise from the CFL condition. As a result of this mesh construction we are able to simulate the full spherical domain including the poles. Our two high-resolution runs use a base resolution of $(N_r, N_\theta, N_\phi) = (112 \times 64 \times 64)$, i.e. double the resolution in θ and ϕ, but are otherwise identical. This gives a maximum effective resolution of $(N_r, N_\theta, N_\phi) = (896 \times 512 \times 512)$ and 40 cells per scale height H.

The base grid is logarithmically spaced in r and linearly spaced in θ and ϕ. We use 'polar' boundary conditions (BCs) at the poles (see Stone et al., 2020, for details), periodic BCs in the ϕ-direction, and outflow BCs on the inner and outer r boundaries. Our polar boundary condition allows magnetic field lines to correctly pass through the coordinate singularity and allows a polar flux bundle to stay in place without (unphysically) being destroyed by the boundary.

Initial conditions

To contrast with previous work and further explore robustness to initial conditions, we initialise our disks with an approximate disk-like configuration following that of Hogg & Reynolds (2018a), rather than the commonly-used Fishbone-Moncrief torus (Fishbone & Moncrief, 1976).

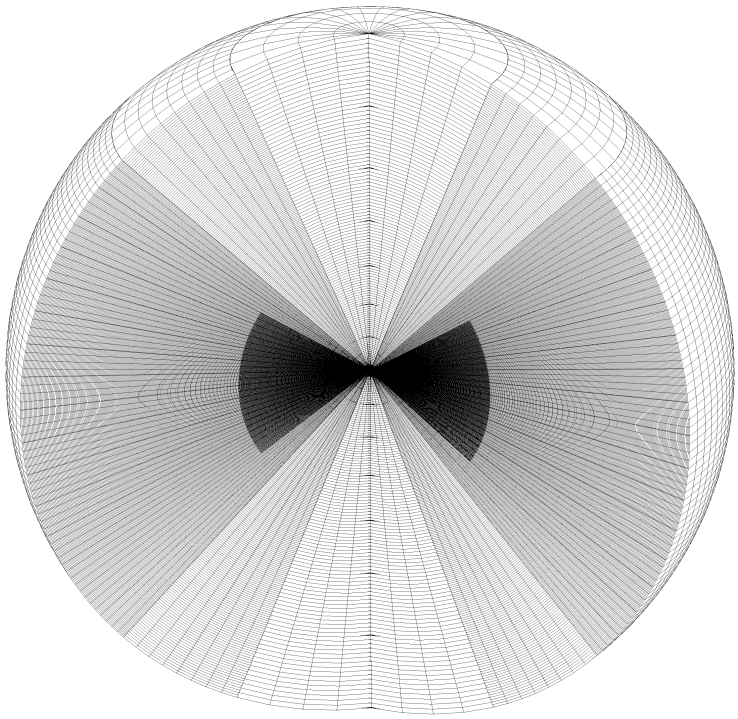


Fig. 2.1 Visualisation of the fiducial-resolution simulation mesh in 3D, with a cutout to show internal zones.

We initialise our disks with a power-law density profile

$$\rho(R) = \rho_0 R^{-3/2} e^{-(z^2)/(2c_s^2 R^3)} \, F_{\rm s}(R), \tag{2.8}$$

where R is the cylindrical radius, $F_s(R)$ is a function motivated by the zero-torque inner boundary condition of Shakura & Sunyaev (1973),

$$F_s(R) = 1 - \left(\frac{r_{\rm ISCO}}{R}\right)^{1/2}. \tag{2.9}$$

$\rho_0 = 100$ is a scaling factor set to give $\max\{\rho(r)\} \sim 1$, and c_s is the isothermal sound speed in a thin disk. The resulting density scale height of the initial disk H_i is given by,

$$c_s = v_{\rm K} H_i \tag{2.10}$$

for Keplerian rotational velocity v_K.

To this density profile, we add a non-axisymmetric perturbation of the form $\sin(10\phi)$ at the 1% level. Breaking the formal axisymmetry of the initial conditions lessens the impact of (transitory) channel modes and hastens the development of saturated turbulence.

The initial velocity profile is that of a test particle on a circular orbit (i.e. 'Keplerian') in the Pseudo-Newtonian potential,

$$v_\phi = \frac{\sqrt{R}}{R-2}, \tag{2.11}$$

with the other components of velocity set initially to zero ($v_r = v_\theta = 0$). This initialises the disk close to, but not exactly matching, its natural equilibrium state, reducing the time needed for the disk to settle into a quasi-steady state and the simulation data to become useful.

The initial magnetic field is a purely toroidal field, set by

$$B_\phi = \sqrt{\frac{2P}{\beta_0}}, \tag{2.12}$$

with all other components equal to zero, i.e. the magnetic pressure is set to be a constant fraction $1/\beta_0$ of the gas pressure P. We calculate P from the density before the introduction of the non-axisymmetric perturbations. Thus the corresponding magnetic field is axisymmetric and, being toroidal, is hence divergence free everywhere. We explore two cases of field strength: a stronger field scenario ($\beta_0 = 5$) and a weaker field scenario ($\beta_0 = 200$). Combined with the two resolutions that we explore, this defines the four simulations presented in this paper and summarised in Table 2.1.

Simulation ID	Max. Resolution ($r \times \theta \times \phi$)	Duration (GM/c^3)	ISCO orbits	β_0	$\langle\langle\beta\rangle\rangle$	$\langle\langle Q_\theta\rangle\rangle$	$\langle\langle Q_\phi\rangle\rangle$	$\langle\langle\theta_B\rangle\rangle$
`b5`	$896 \times 256 \times 256$	1.67×10^5	2704	5	9.39	20.56	63.22	11.85
`b200`	$896 \times 256 \times 256$	4.02×10^5	6535	200	23.22	9.22	27.38	10.62
`b5_hi`	$896 \times 512 \times 512$	3.15×10^4	511	5	7.87	39.89	106.56	11.14
`b200_hi`	$896 \times 512 \times 512$	1.54×10^5	2502	200	20.90	10.75	47.82	10.75

Table 2.1 Simulation parameters. NOTE: Averages denoted by $\langle\langle X\rangle\rangle$ are taken over $r \in [5, 100]$, ϕ, and three scale heights in θ, at late times in the disk as defined in Section 2.2.3.

2.2.3 Diagnostics and convergence

MHD accretion disks are highly turbulent systems and thus, even when they have reached a 'quasi steady state', will be characterised by strong spatial and temporal fluctuations in all quantities. Useful diagnostics of disk structure necessarily require some averaging process.

In this paper, spatial averages are denoted by $\langle X \rangle$; by default, this means averaging over $r < 100 r_g$, all ϕ, and across $\theta = \text{constant}$ shells within ± 3 scale heights of the mid-plane. Quantities denoted by $\langle\langle X \rangle\rangle$ are additionally averaged over time; by default, we average our weak field simulations (`b200` and `b200_hi`) for $t \gtrsim 2000 T_{\rm ISCO}$ ($t > 1.2 \times 10^5\ GM/c^3$) and our strong-field simulations (`b5` and `b5_hi`) for $t \gtrsim 500 T_{\rm ISCO}$ ($t > 3.0 \times 10^4\ GM/c^3$). Any deviations from these defaults are indicated by $\langle X \rangle^\star$ or $\langle\langle X \rangle\rangle^\star$ and details specified within the text.

In addition to looking at the evolution of quantities already described (β, ρ, v_r), we calculate and track several additional standard quantities. The mass accretion rate across some closed surface S is

$$\dot{M} = \iint_S \rho \boldsymbol{v} \cdot d\boldsymbol{S}, \tag{2.13}$$

which we calculate across the sphere $r = 6$, namely the innermost circular stable orbit (ISCO) of this potential.

The geometrical scale height $h(r)$ is calculated as in Hogg & Reynolds (2018b), with the scale height at each radius given by

$$\frac{h(r)}{r} = \left\langle \frac{\int [\theta(r) - \bar{\theta}(r)]^2 \rho d\Omega}{\int \rho d\Omega} \right\rangle \tag{2.14}$$

for solid angle $d\Omega = \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi$ and

$$\bar{\theta}(r) = \frac{\int \theta(r) \rho d\Omega}{\int \rho d\Omega}. \tag{2.15}$$

The radial bins from equation 2.14 are then averaged by weighting them according to the radial width at each shell (as cells are logarithmically spaced in r), giving a disk-averaged scale height H. To prevent confusion with other parameters within the text, we use H to refer to the disk-averaged geometrical scale height rather than $\langle \frac{h(r)}{r} \rangle$.

With the focus of this paper on the magnetic evolution of the disk, a parameter of key interest is the horizon penetrating magnetic flux,

$$\Phi_{\rm BH} = \frac{\sqrt{4\pi}}{2} \iint_{S_H} \boldsymbol{B} \cdot d\boldsymbol{S}, \tag{2.16}$$

where S_H is the upper half-sphere near the inner boundary of our domain ($r = 6, \theta > 0$). This is then re-normalised by the late-time-averaged accretion rate to $\phi = \Phi_{\rm BH} / (\langle\langle \dot{M} \rangle\rangle r_g c^2)^{-1/2}$ (Tchekhovskoy et al., 2011). For a disk scale height $H \sim 0.3$ and BH spin $a = 0$, the

transition from SANE to MAD states is expected to occur for $\phi \sim 50$ (Tchekhovskoy et al., 2012)[1].

We also calculate the Shakura-Sunyaev alpha (Shakura & Sunyaev, 1973),

$$\alpha_{\rm visc} = \frac{T_{\hat{r}\hat{\phi}}}{P} = \frac{T_{{\rm Max},\hat{r}\hat{\phi}} + T_{{\rm Rey},\hat{r}\hat{\phi}}}{P} \tag{2.17}$$

where $T_{\hat{r}\hat{\phi}}$ is the fluid frame stress, and $T_{{\rm Max},\hat{r}\hat{\phi}} = -B_r B_\phi$ and $T_{{\rm Rey},\hat{r}\hat{\phi}} = \rho u_r \delta u_\phi$ are the Maxwell and Reynolds components of that stress, where $\delta u_\phi \equiv u_\phi - r\Omega_{\rm Kep}$ is the fluctuating component of the ϕ-velocity. For thin disks, $\alpha_{\rm visc} \sim 0.1$, however, this value is known to vary with disk thickness, initial magnetic field morphology/strength, and radius within the disk (e.g. Penna et al., 2013).

A key question is the degree to which a numerical model of an accretion disk is converged. There are two distinct senses in which we must consider convergence. Firstly, since our disks are initialised in a laminar state that is close to hydrodynamic equilibrium but strongly unstable to the MRI, they must first undergo evolution before settling into a (turbulent) statistical steady state. Only then can the simulation data be considered physically meaningful (out to the radius where we achieve inflow equilibrium). Diagnosing this state informs the times on which we focus our analysis. Secondly, as we are interested in studying the behaviour of the magnetic field we need to be sure that the underlying physical mechanisms that govern its evolution are spatially well-resolved within the disk. To assess these criteria, we turn to two measures of convergence: the quality factors Q_θ and Q_ϕ, and the magnetic tilt angle θ_B. Additionally, we are able to do a very limited resolution study by comparing our low and high resolution simulations.

The quality factors, Q, defined by Noble et al. (2010) as

$$Q_\theta = \frac{\lambda_{{\rm MRI},\theta}}{r\Delta\theta} \tag{2.18}$$

$$Q_\phi = \frac{\lambda_{{\rm MRI},\phi}}{R\Delta\phi}, \tag{2.19}$$

are a direct measure of the resolvability of the linear MRI modes via a comparison of the characteristic wavelength of the MRI modes $\lambda_{\rm MRI} = 2\pi v_A/\Omega$, where v_A is the Alfvén velocity, with the size of the basic simulation voxel. The Q's are simply a measure of the number of resolution elements that span one wavelength of the fastest growing MRI mode.

[1] We note that $\phi \sim 50$ is the saturation value in code or Gaussian units. When using Lorentz-Heaviside units the factor of $\sqrt{4\pi}$ in equation 2.16 disappears, giving $\phi \sim 15$ instead.

Values as low as $Q_\theta = 6$–8 are sufficient to capture the linear MRI (Flock et al., 2010) in low-spatial order methods, but must rise to $Q_\theta > 10$ with $Q_\phi > 20$ to capture the non-linear interactions between MRI modes (Hawley et al., 2011, 2013; Sorathia et al., 2012), although a higher Q_ϕ or Q_z (when in cylindrical coordinates) may offset a slightly lower Q_θ.

Our second diagnostic tool is the magnetic tilt angle, θ_B. First used by Guan et al. (2009), θ_B is a measure of the correlations between the radial and toroidal fields,

$$\theta_B = -\arctan\left\langle \frac{B_r}{B_\phi} \right\rangle. \tag{2.20}$$

Guan et al. (2009) originally found this to have a characteristic value of $\theta_B \sim 15^\circ$ for a well-resolved saturated MRI-driven turbulence with later works on local (Hawley et al., 2011) and global (Hawley et al., 2013; Hogg & Reynolds, 2016, 2018a; Sorathia et al., 2012) simulations revising this number down to 11–13°. The physics dictating the saturation of MRI-driven turbulence and hence the value of the tilt angle is still not clear, although Pessah (2010) describe such behaviour with a model in which saturation corresponds to disruption of linear-MRI modes by the parasitic instabilities of Goodman & Xu (1994). Our primary focus with θ_B is not the final value itself, as this is not a direct measure of resolution and so may be affected by other factors in the disk, but we instead look for a stabilisation in the value of this parameter.

The time-evolution of these convergence metrics for all of our runs is shown in Figure 2.2, and the late-time average is reported in Table 2.1. Through these metrics we find that our strong-field simulation `b5` is well-resolved with $Q_\theta \sim 16$, $Q_\phi \sim 38$, and a tilt angle that stabilises at $\theta_B \approx 12^\circ$. The numbers are less comfortable for `b200` (Figure 2.2). The lower values of $Q_\theta \sim 9$ indicates that `b200` only marginally resolves the dominant linear MRI modes, although θ_B does eventually stabilise for $t \gtrsim 2 \times 10^5\ GM/c^3$, albeit at a notably lower value ($\theta_B \sim 10$) than the stronger field case.

A more direct way to assess the effects of resolution is to compare our standard resolution runs (`b200` and `b5`) with their high-resolution counterparts (`b200_hi` and `b5_hi`). As expected, these high-resolution runs achieve much higher Q_θ and Q_ϕ values, well above the fiducial levels needed. The magnetic tilt angle θ_B takes longer to stabilise than in the lower-resolution counterparts but eventually stabilises at comparable levels to the low-resolution runs.

We conclude that all simulations have achieved a converged state by the end of their respective runs, although `b200` may be only marginally resolving both the linear and non-linear evolution of the MRI.

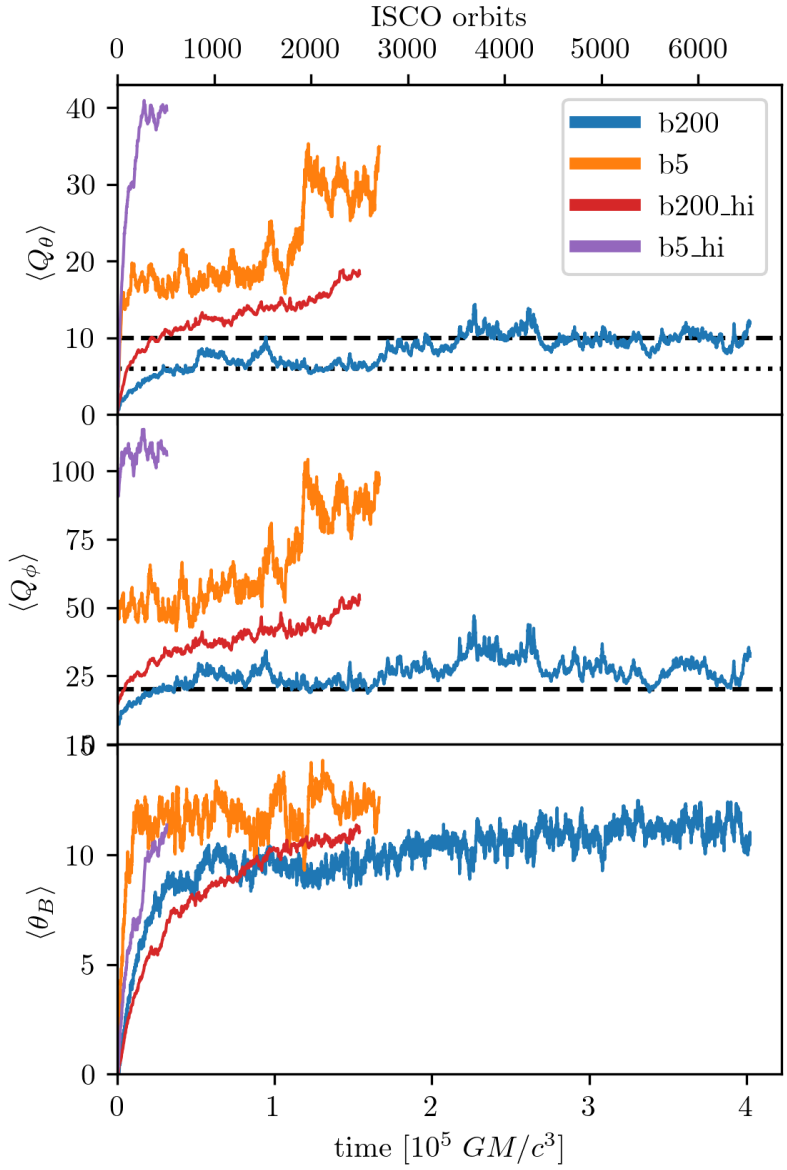


Fig. 2.2 Convergence metrics: (*top* and *middle*) quality factors in the θ- and ϕ-directions, and (*bottom*) magnetic tilt angle, θ_B. The dotted line indicates the minimum value required to resolve the linear MRI, while the dashed lines represent the requirements for resolving the non-linear MRI. The `b5` simulation along with the high resolution `b5_hi` and `b200_hi` resolve both the linear and non-linear MRI modes, while `b200` only marginally resolves the non-linear MRI in Q_θ. Both `b5` and `b200` show a plateauing in θ_B, indicating convergence. The high resolution runs `b5_hi` and `b200_hi` largely follow the trends of their fiducial counterparts, although late-time evolution cannot be ruled out.

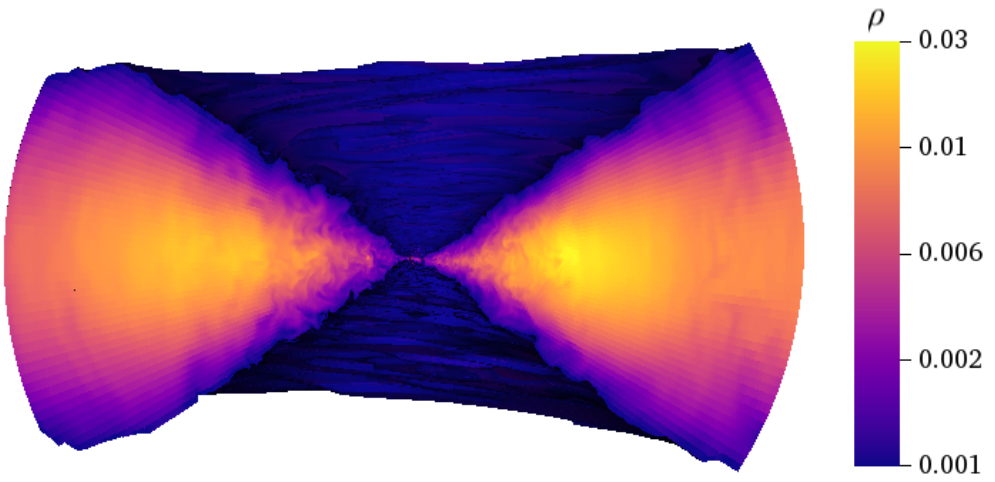


Fig. 2.3 Three-dimensional side-on volume rendering of the simulation domain for one of our fiducial runs, b200, at a late time in the disk. Colour denotes relative number density, ρ.

2.3 Results

2.3.1 Disk structure and evolution

Starting from the initial state described in Section 2.2.2, the disks undergo rapid evolution driven principally by the growth and eventual saturation of MRI-driven turbulence. The outer parts of the simulated disks never achieve a state of inflow equilibrium; we focus our analysis on the inner (and most highly resolved) regions of the disks within $r = 100r_g$. Figure 2.3 shows a late-time three-dimensional volume rendering of run b200, and Figure 2.4 shows the evolution of mass accretion rate across the ISCO as well as the disk-averaged plasma-β, geometric scale height H, and Shakura-Sunyaev $\alpha_{\rm visc}$. Both b5 and b200 show a characteristic early peak and slow drop-off in mass accretion rate $\dot{M}$, as expected for disks with finite mass (Figure 2.4). b200 stabilises to a higher mass accretion rate of $\dot{M} \sim 0.25$, possibly related to the presence of a density lump that inhibits early-time accretion. In all models, we find that $\langle\beta\rangle$ reaches an approximately steady value relatively quickly. The memory of the strength of the initial magnetic field is retained, with run b200 stabilising at $\langle\beta\rangle \sim 20$ and run b5 stabilising at $\langle\beta\rangle \sim 10$. The initial magnetisation also makes a difference to the disk thickness, with stabilisation at $\langle H\rangle \sim 0.27$ and $\langle H\rangle \sim 0.32$ for b200 and b5 respectively. The dominant pressure dictating the disk thickness in all cases is gas pressure, with a modest amount of additional magnetic support thickening the disk in the stronger initial field case b5. The spatially-averaged $\langle\alpha_{\rm visc}\rangle$ (lower panel) remains relatively

stable around 10^{-2}–10^{-1}, as is expected for accretion disks, with a small jump in `b5` at late times which coincides with a thinning of the disk (i.e. lower H).

For the stronger field case, doubling the θ- and ϕ-resolution (going from `b5` to `b5_hi`) makes very little difference to the overall evolution of the disk, with $\dot{M}$, $\langle\beta\rangle$ and $\langle H\rangle$ following very similar time-traces. In the case of the weaker field model, the early-to-intermediate timescale evolution is very similar for the high-resolution models, whereas the late-time magnetisation is a little higher ($\langle\beta\rangle$ marginally smaller), leading to a slightly thicker disk. This further suggests that the fiducial resolution weaker field run `b200` only marginally resolves the non-linear evolution of the MRI.

The radial profiles of density $\langle\langle\rho\rangle\rangle^*$, radial velocity $\langle\langle v_r\rangle\rangle^*$, temperature $\langle\langle P/\rho c^2\rangle\rangle^*$, and effective Shakura-Sunyaev $\langle\langle\alpha_{\rm visc}\rangle\rangle^*$ are shown for all four runs in Figure 2.5. These are time-averaged from the point where the run is deemed to have stabilised (as laid out in Section 2.2.3) out to the end of the simulation. Given that the different runs have quite different end-times, some caution must be exercised in making direct comparisons between profiles from different runs. To aid in this comparison, we also provide the profiles for the low-resolution runs (`b200` and `b5`) averaged over the same time range as the high-resolution counterparts shown as dashed lines. For example, the large difference in the overall normalisation of density between `b5` and `b5_hi` is almost all due to the fact that the lower-resolution run was evolved much longer and hence drained much more of its mass; when the average is calculated over the same time range, this difference largely disappears. However, the comparison of `b200_hi` and `b5` is particularly interesting. Both runs well-resolve the MRI and have evolved for a similar time, with the principal difference being the initial magnetic field strength. We find that both runs have very similar radial velocity profiles during these late-time averages, although `b5` experienced high-accretion rate transients at early times that suppresses its density. This early-time behaviour is clearly seen in the profiles for `b5_hi` which clearly shows significantly higher radial inflow within the central $30r_g$.

The density profile of `b200` shows a broad hump around $r = 100r_g$. This signposts a particularly interesting feature of this disk that, we argue, is a subtle and pernicious imprint of the SMR de-refinement boundaries. The overdensity, or $m = 1$ mode, develops at a relatively early time of $t \sim 1 \times 10^5\ GM/c^3$ and is then sustained throughout the subsequent evolution. This structure is most apparent in the mid-plane density map shown in Figure 2.6 (left panel). The time-dependence of this lump can be seen by forming space-time plots of the density along a midplane line that cuts through the centre, ($\theta = \pi/2,\ \phi = 0, \pi$); see Figure 2.7 (top panel). The sloshing motion of this lump as it orbits the central body is seen as the long-lived 'see-saw' feature, causing an excavation of material from the central region and creating

the density maximum at $r \sim 100r_g$ (Figure 2.5). An examination of the mid-plane velocity field reveals this lump to be a large-scale vortex. It has been argued that such features might arise from a Rossby Wave Instability (RWI; Lovelace et al., 1999) or generalised form of the Papaloizou-Pringle Instability (PPI; Papaloizou & Pringle, 1984). The presence of the RWI within disks has already been argued in the case of Tidal Disruption Events (TDEs) and the spirals it produces may explain the observed flaring behaviours in Sgr A* (Falanga et al., 2007; Mignon-Risse et al., 2021). A similar feature was also reported in an MHD simulation by Mishra et al. (2020), who observed spiral-like inhomogeneities in the density of disks with initially weak poloidal fields, however, these simulations were only run for 23 orbits (versus our ~ 6500 for `b200`) and so the late-time evolution of the effect was not noted.

As seen in Mishra et al. (2020), the lump creates a kind of anomalous accretion where material preferentially accretes from the atmosphere or ceiling of the disk rather than in the midplane. This is clearly seen in Figure 2.8 which shows the poloidal velocity field on a vertical cross-sectional slice for late-time snapshots from `b200` and `b5`. Run `b200`'s velocity streamlines (Figure 2.8, left panel) show strong inflows in the jet regions of the `b200` simulation and eddies along the disk-jet boundary which take material out of the disk and deposit it in the jet streams, while the radial velocity within the disk is reduced (Figure 2.5).

The `b5` disk, by comparison, does not develop such an overdensity and so retains a significant amount of its mass below its $\rho_{\max}$ at $20r_g$ (Figures 2.5 and 2.6, right panel). This smoother disk is also accompanied by a greater inflow velocity v_r and a higher temperature $P/\rho c^2$. Thus stronger magnetic fields suppress the formation of the vortex. More importantly, an examination of the corresponding space-time diagram (Figure 2.7, bottom panel) shows that the lump is not observed in the high-resolution simulation `b200_hi`. This fact, together with the location of the `b200` lump, strongly suggests that the lump is an artefact caused by vorticity generation at the SMR refinement boundary, which affects only the weaker-field case. This is a cautionary note for any long-duration accretion disk models that employ SMR.

2.3.2 Magnetic field strength, morphology, and evolution

A primary motivation for our study is to understand the evolution of the magnetic field, particularly the growth and morphology of the poloidal field, in these disks that start with purely toroidal configurations. Furthermore, we wish to understand whether the strength of the initial toroidal field has any long-lasting imprint on the magnetic evolution of the disks.

Our initially weak-field simulation, `b200`, shows rapid growth of the magnetic field at early times, settling to $\langle\beta\rangle \sim 20$ (Figure 2.4), compared to $\langle\beta\rangle \sim 10$ for the initially strong-field case `b5`. Interestingly, we see the development of large-scale poloidal field in both cases. Figure 2.9 shows the horizon-threading (hence necessarily poloidal) flux as defined in equation 2.16. While there is clearly non-zero poloidal flux in all cases, `b200` shows a consistently low flux $\phi \sim 5$–15 with little sustained growth, whereas `b5` exceeds $\phi = 40$ and nears the transitional value for a MAD state of $\phi = 50$ (Liska et al., 2020; Tchekhovskoy et al., 2012) for a non-spinning BH, although fails to cross. The high-resolution simulations `b5_hi` and `b200_hi` again match their lower-resolution analogues remarkably well.

It is not merely the net horizon-crossing flux that matters, though, when considering features like jet launching. The degree of order of this field as it crosses the inner boundary is also important to questions such as whether the jet is 'striped' (Giannios & Uzdensky, 2019). Most disk models that start with initially poloidal field configurations end up with a well-organized (unidirectional) horizon-threading flux bundle. However, that is not necessarily the case for our models, which begin with an initially toroidal field. To examine this, we calculate the ratio of signed to unsigned fluxes threading the inner boundary of our simulation,

$$\phi_{\parallel} = \frac{\iint_S \boldsymbol{B} \cdot d\boldsymbol{S}}{\iint_S |\boldsymbol{B}| \cdot d\boldsymbol{S}}, \tag{2.21}$$

where S is a hemisphere surface at $r = 6r_g$ and the final value given is the average of the two hemispheres. A value of $\phi_{\parallel} = 1$ indicates agreement in the direction of the field, while a value $\phi_{\parallel} = 0$ indicates that (for example) for every positive B_r there is an equal negative, cancelling the flux everywhere on the surface integrated.

With this new measure, we see that `b5` has a high degree of field order with $\phi_{\parallel} \sim 1$ at late times, while `b200` barely reaches $\phi_{\parallel} \sim 0.8$ at $t \sim 2.5 \times 10^5$ before decreasing again and spending much of its time with $\phi_{\parallel} < 0.6$. Again, we see a high degree of agreement between the high- and low-resolution runs.

The nature of these differences is seen more clearly in renderings of the poloidal magnetic field morphology at late times in Figure 2.10. In `b200`, close to the black hole, we see a well-defined field direction dominating the polar regions and then a strong region of oppositely directed field along a lower-latitude boundary funnel. The mid-plane field is weak and disorganised, existing within small magnetic loops distributed seemingly at random. Recalling that this region is flowing *inwards*, some of this magnetic topology may be due to large-scale circulations between the disk atmosphere and the polar regions. By contrast, `b5` displays noticeably stronger fields in the midplane of the disk, and where there are magnetic

loops these are larger than in `b200`. The polar region contains an organised and sustained flux bundle along which there is outflowing material. It must be recognised, though, that while $\phi_{\parallel} \sim 1$ for this case, the flux bundle still contains interesting structure (field loops and current sheets) downstream, possibly due to Kelvin-Helmholtz modes excited at the jet-disk boundary. Such large magnetic field loops, particularly those on the right side of the upper jet region and the left side of the lower jet region of Figure 2.10, are reminiscent of the large-scale poloidal field loops identified by Liska et al. (2020), and so leads us to tentatively conclude that simulation `b5` is capable of generating a large-scale poloidal field from an initially-strong toroidal field, however, evidence is much weaker for the weak-field case `b200`.

Our discussion of magnetic evolution so far has focused on the poloidal field, but we note that there are also interesting dynamics in the toroidal field. We recover the well-known butterfly patterns—a 'feathering' pattern in the space-time diagram for the azimuthally-averaged B_ϕ at a particular radius (Figure 2.11). These are interpreted as dynamo-cycles, an emergent phenomenon that arises from the non-linear MRI physics (Gressel & Pessah, 2015; Hogg & Reynolds, 2018b; Rincon et al., 2007). In `b200` (Figure 2.11, top panel), the butterfly pattern is rapidly established and is then sustained for the full duration of the simulation. This plot also explicitly shows that (after azimuthal averaging) a non-zero but weak toroidal flux exists within the polar region. These findings completely track over into the higher-resolution version of this model `b200_hi`. By contrast, `b5` establishes a dominant toroidal field sign in the polar regions at a relatively early time ($t \sim 5 \times 10^4\ GM/c^3$), coincident with the creation of the strong and organised poloidal flux bundle (Figure 2.9). From this time onward the butterfly pattern is significantly weakened, showing that the magnetic configuration in the polar region is affecting the operation of the dynamo cycle.

While we have attributed its presence to the (artificial) SMR refinement boundary, it is still interesting to consider the effect of the $m = 1$ lump found in `b200` and discussed in Section 2.3.1. As the inflow of material from large radius is impeded by the lump, a lower-density cavity forms within the inner regions of the disk (Figure 2.5). This increases the magnetisation of the region (decreases the plasma-β) which may be responsible for weakening the amplitude of the dynamo cycles in `b200` from $t \approx 1.0 \times 10^5\ GM/c^3$ onward.

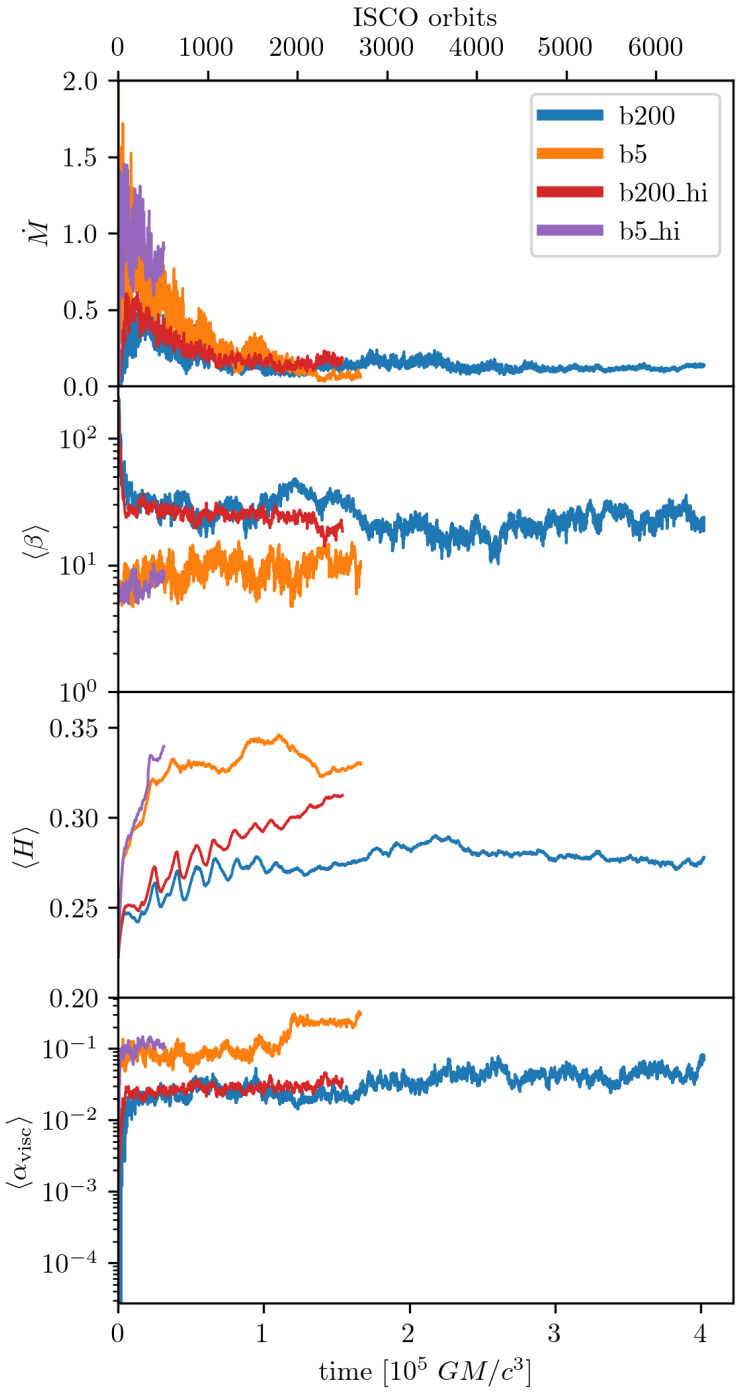


Fig. 2.4 (*Top*) Mass flux $\dot{M}$ through the surface at $r = 6r_g$, (*top middle*) spatially-averaged magnetic plasma $\langle\beta\rangle$, (*bottom middle*) spatially-averaged geometric scale height $\langle H\rangle$, and (*bottom*) spatially-averaged Shakura-Sunyaev alpha $\langle\alpha_{\rm visc}\rangle$. Of note, we find that the plasma β values of `b200` and `b5` do not converge. Although `b200_hi`'s scale height shows a late-time trend towards the `b5` results, a similar trend is not observed in other variables plotted.

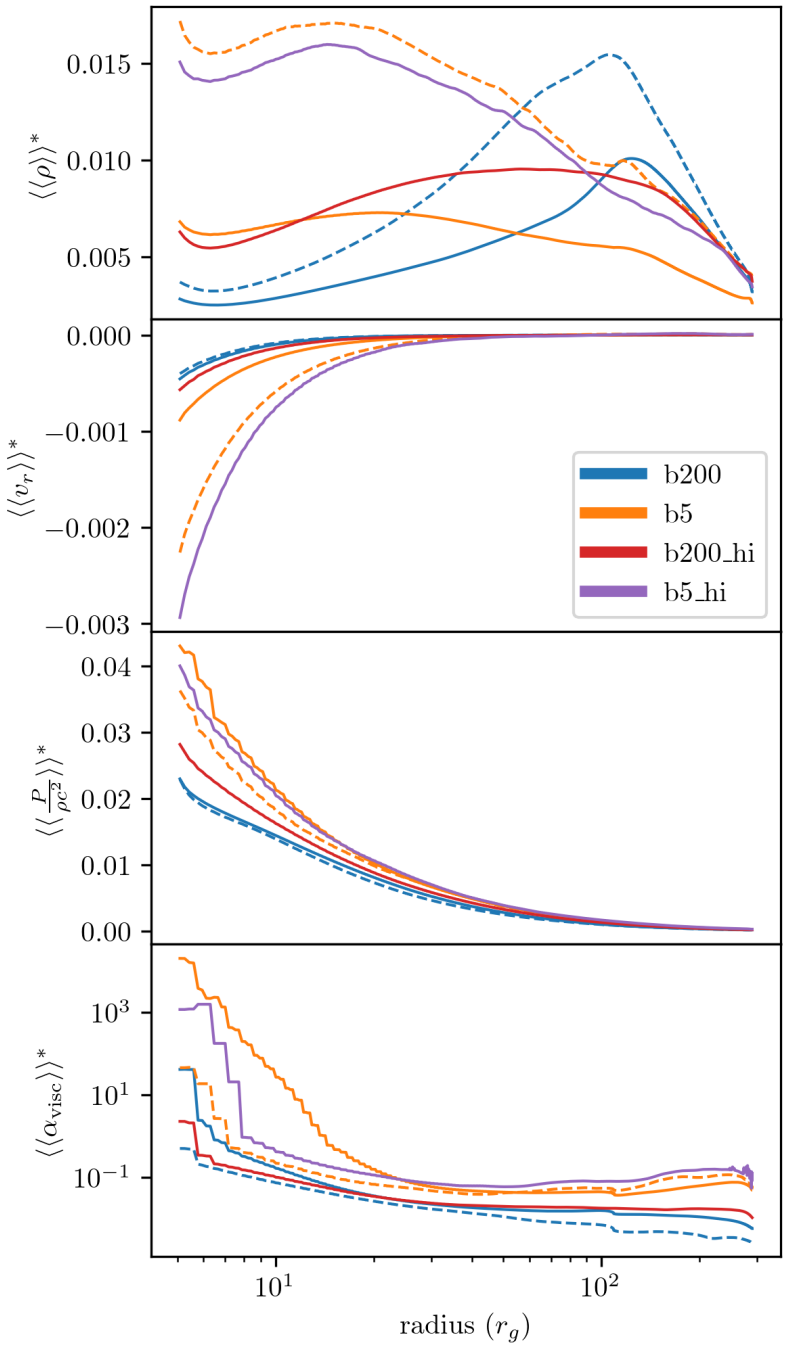


Fig. 2.5 Late-time-averaged profiles for the density (*top*), radial velocity (*top middle*), temperature (*bottom middle*), and Shakura-Sunyaev alpha (*bottom*) as a function of disk radius. Dashed lines for `b200` and `b5` are averaged over the same time windows as `b200_hi` and `b5_hi` respectively. We find the main differences between `b5` and `b5_hi` are resolved when averages are taken over like times (purple vs orange dashed), but differences remain in the `b200` and `b200_hi` cases (red vs blue dashed). Note: Spatial averages are taken azimuthally and within $3H$.

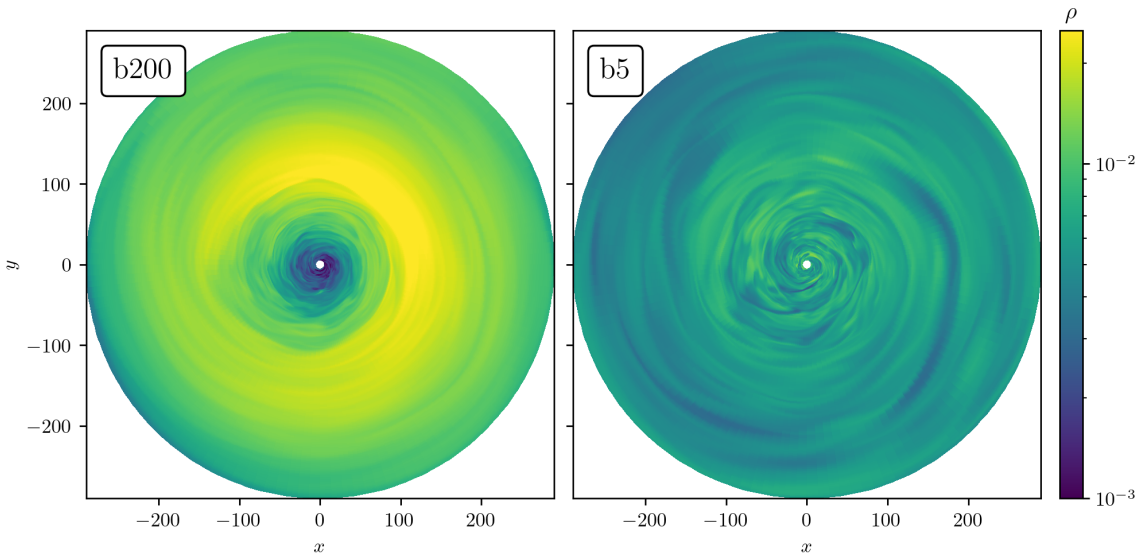


Fig. 2.6 Late-time snapshots of midplane density for `b200` (*left*) and `b5` (*right*), taken at $t = 4.05 \times 10^5\ GM/c^3$ and $t = 1.25 \times 10^5\ GM/c^3$ respectively. Significant differences in the density profile are apparent, most notably the presence of a large co-orbiting 'lump' in `b200`, leading to a relative evacuation of material in the inner part of the disk compared to `b5`.

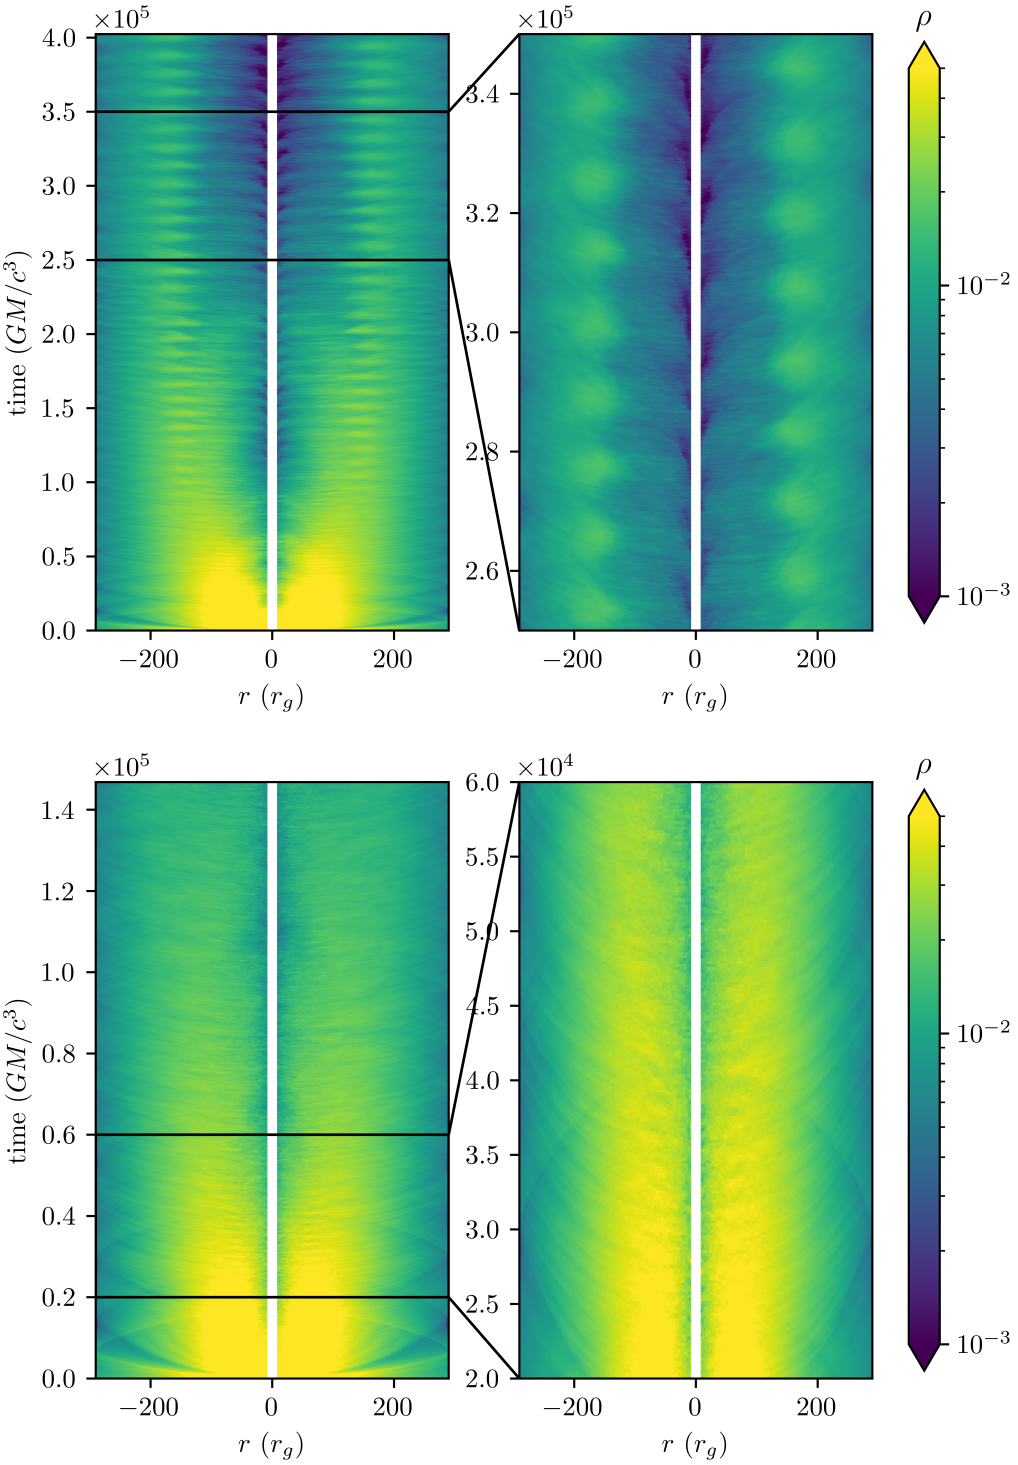


Fig. 2.7 Density along a midplane slice ($\theta = \pi/2$, $\phi = \{0, \pi\}$) over time, for `b200` (*top*) and `b200_hi` (*bottom*). The overdensity, or $m = 1$ mode, is clearly visible in `b200` as striations between radii 100–200r_g. No comparable pattern is observed in `b200_hi`.

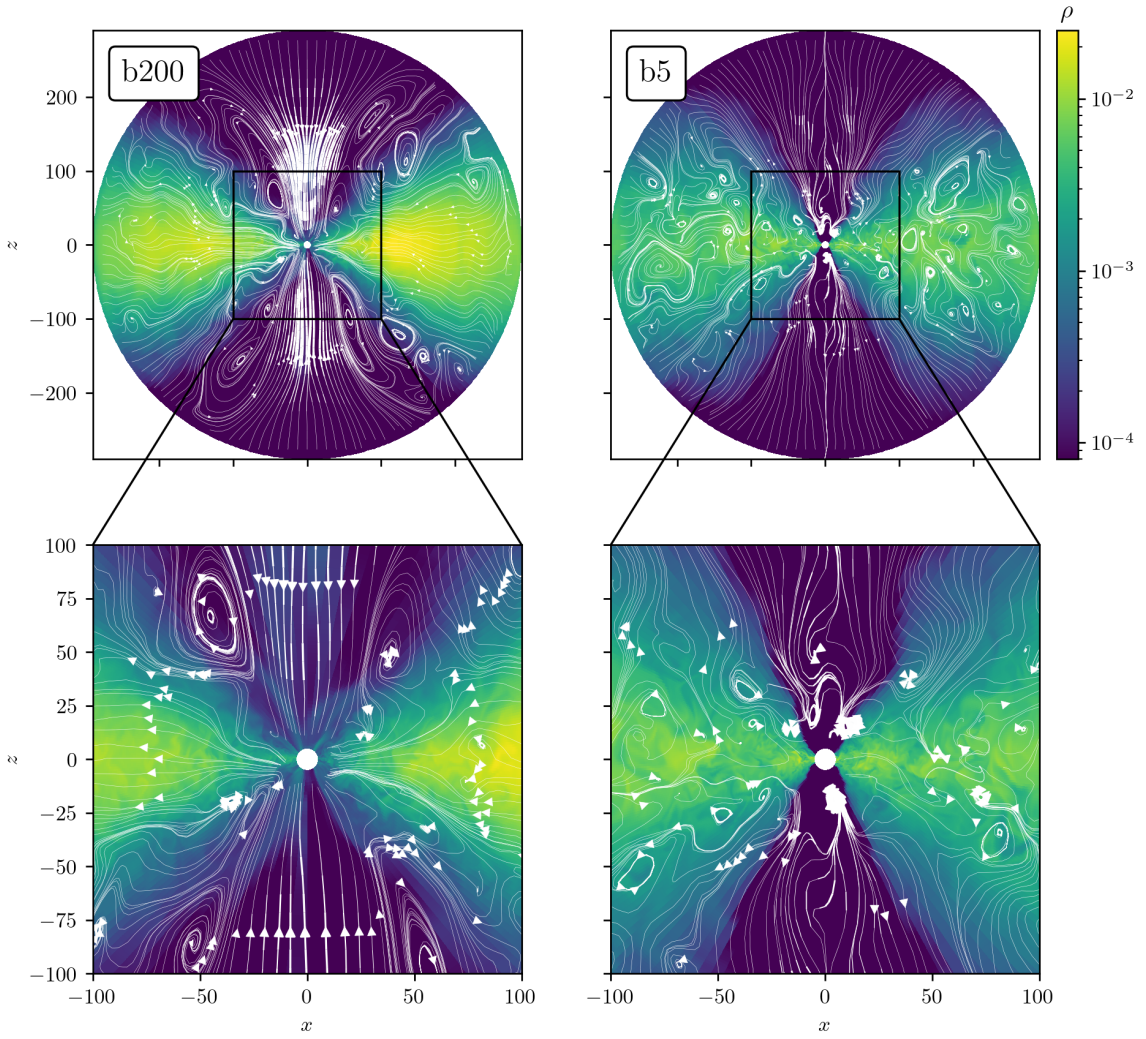


Fig. 2.8 Velocity streamlines for the whole domain (*top*) and zoom-in on the $r < 100r_g$ region (*bottom*) overlaid onto density for `b200` (*left*) and `b5` (*right*). Relative velocity magnitude is indicated by the linewidth of streamlines. Stills are taken from late times: $t = 4.05 \times 10^5\ GM/c^3$ for `b200` and $t = 1.25 \times 10^5\ GM/c^3$ for `b5`. Simulation `b200` features a well-ordered inwards polar velocity field with large circulations along the jet/disk boundary. `b5`, in contrast, displays a less ordered polar velocity field and significantly more small-scale turbulent structure within the disk.

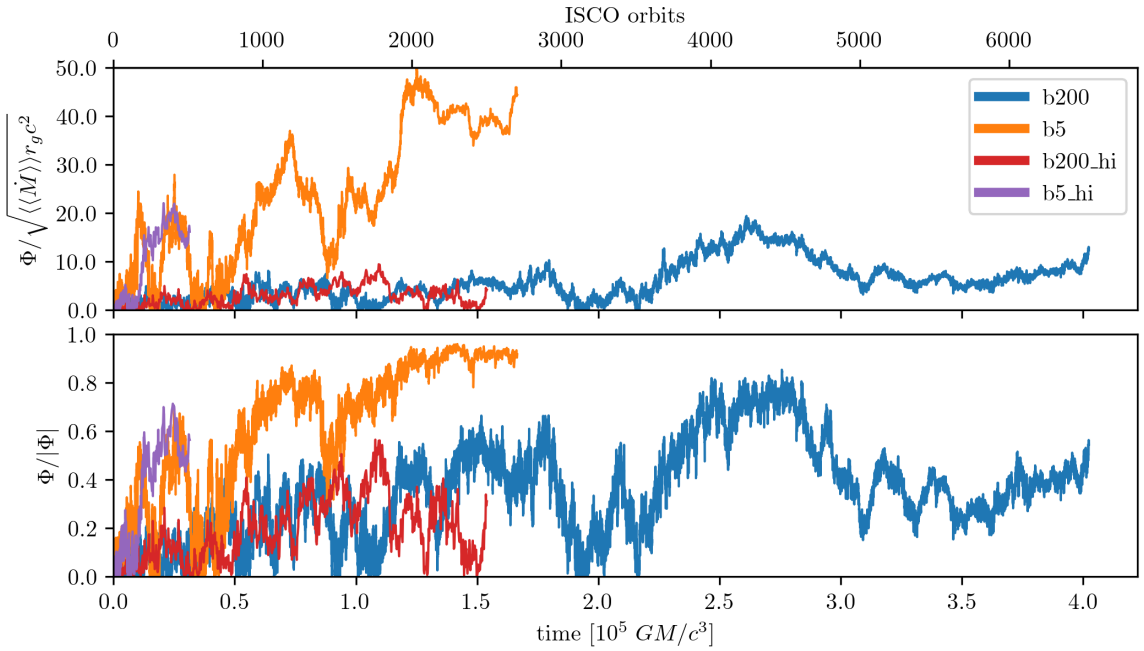


Fig. 2.9 (*Top*) Magnetic flux, $\Phi_B = \int B_r \cdot dS$, through the ISCO ($r = 6r_g$). The signed flux is calculated for each hemisphere individually before the absolute value is taken, and the two hemispheres are averaged to give $|\langle\Phi_B\rangle|$. The values are re-normalised according to Tchekhovskoy et al. (2011) (Section 2.2.3). While `b5` approaches the critical MAD value of $\phi \sim 50$, it fails to cross. (*Bottom*) The ratio of signed and unsigned magnetic fluxes, indicating how much of the flux is aligned in a common direction. The initially stronger field simulations `b5` and `b5_hi` exhibit a significantly higher degree of flux alignment.

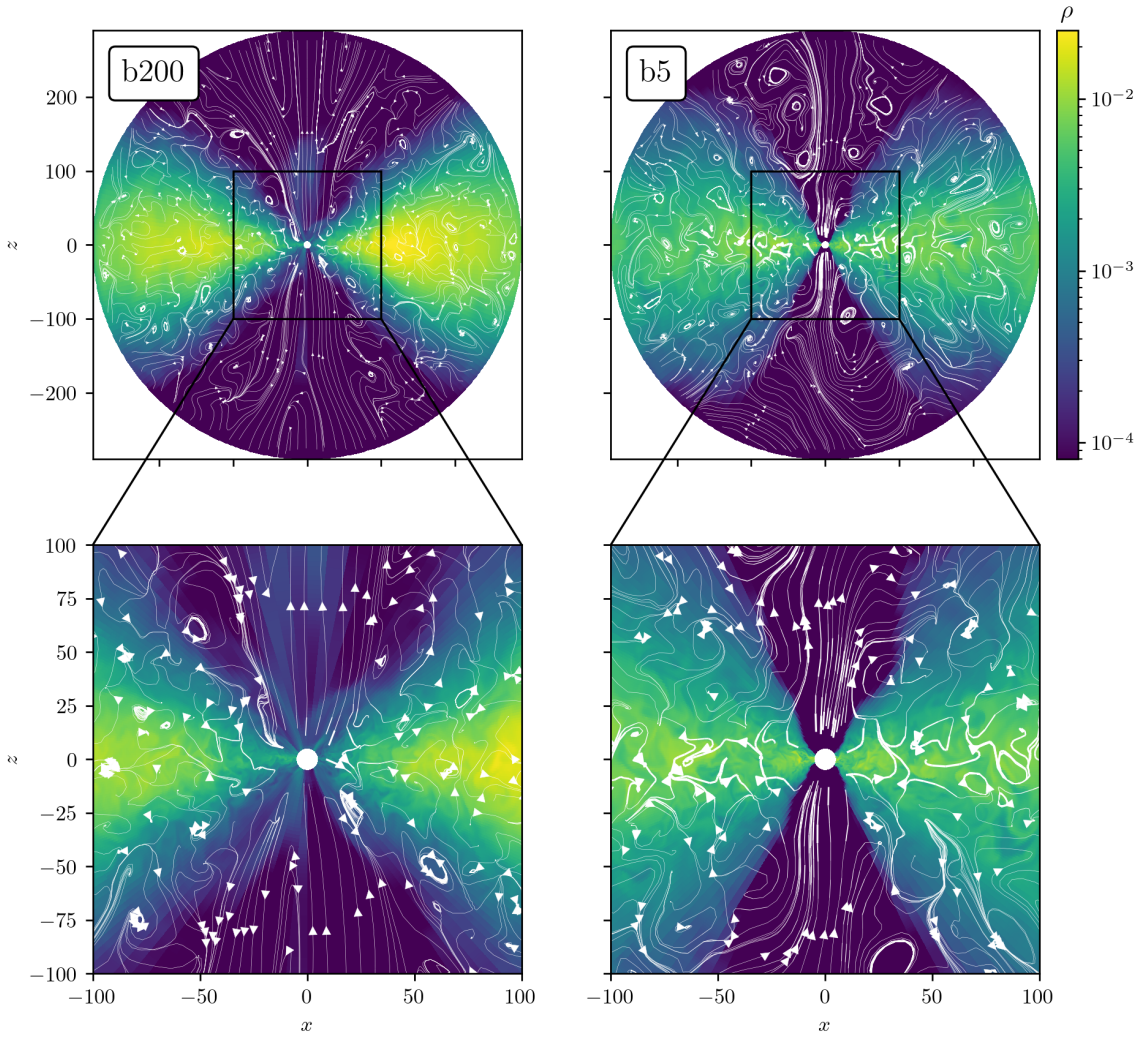


Fig. 2.10 As in Figure 2.8, but with magnetic field streamlines. `b5` displays stronger magnetic field along the disk midplane and within the polar regions, while for `b200` field is amplified in the circulation regions between the disk and the poles.

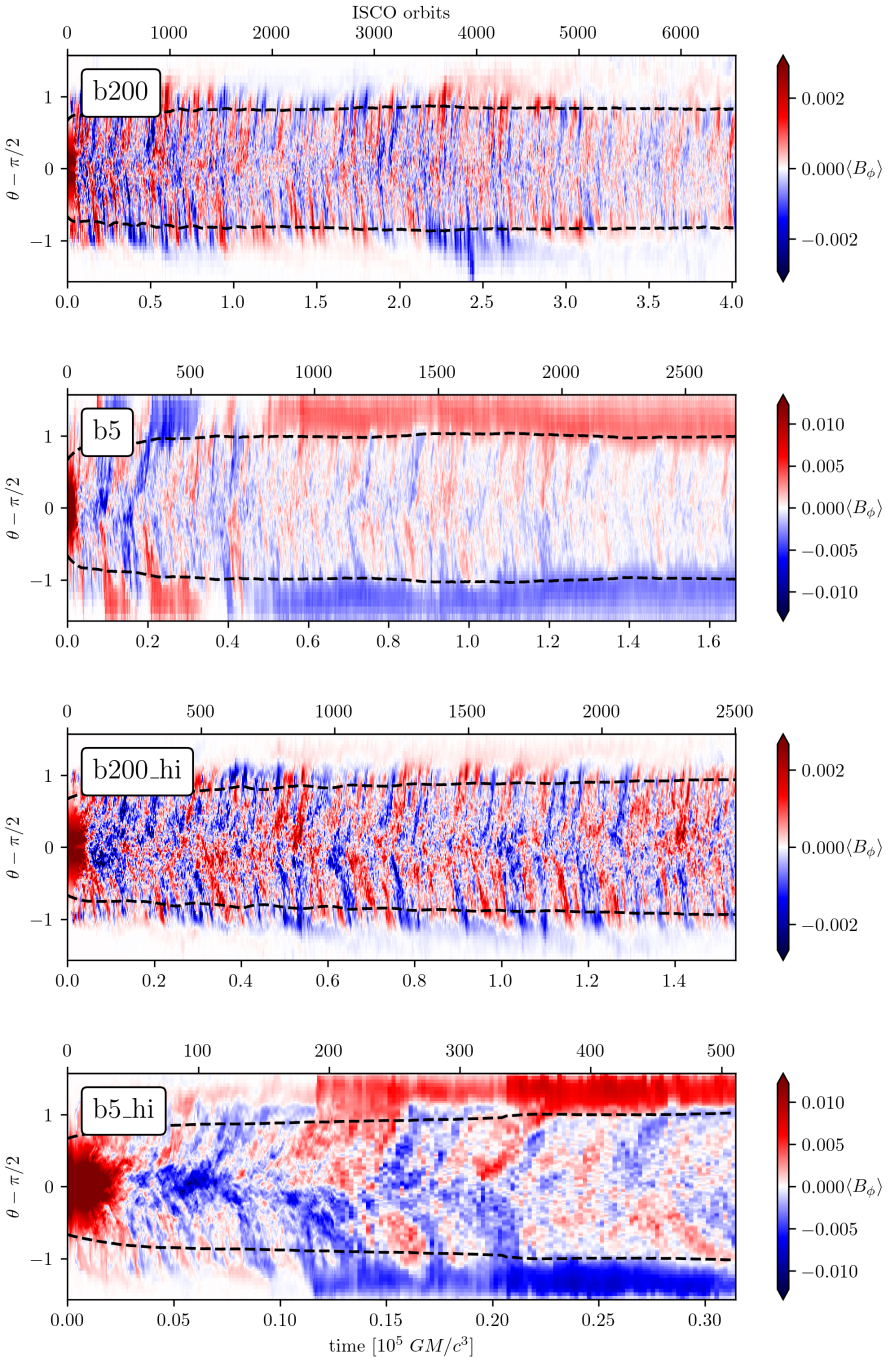


Fig. 2.11 Butterfly diagrams of the azimuthally-averaged B_ϕ at $r = 25r_g$ showing signs of a sporadic butterfly pattern. Dashed lines mark $\pm 3H$. Note there is a change in output cadence for `b5_hi` (*bottom*) for $t > 1.25 \times 10^4\ GM/c^3$. All simulations show signs of disordered field reversal, characterised by patterns of alternating field. The pattern is significantly weaker in `b200` and `b200_hi` perhaps indicating the saturation of the MRI. The strong-field simulations (`b5`, `b5_hi`) establish a dominant polar sign easily, while the process is weaker and more sporadic in the weak-field simulations (`b200`, `b200_hi`).

2.4 Discussion and conclusion

In this work we have presented the results of two 3D MHD accretion disk simulations with initially toroidal fields. These simulations were identical in all respects bar their initial field strength, defined through the initial magnetic plasma beta, $\beta = P_{\rm gas}/P_{\rm mag}$. One simulation, `b200`, was initialised with $\beta = 200$, while the other, `b5`, was initialised with $\beta = 5$. Additionally, two more high-resolution runs were done with double the θ and ϕ resolution for comparison (`b200_hi` and `b5_hi`; for details, see Section 2.2).

We find that our two fiducial simulations diverge early in their evolution and remain on separate paths throughout each run, with notable differences in both hydrodynamic and magnetic quantities. We show that these differences (with one exception) are not sensitive to resolution for the values studied. As the only difference between these two simulations is the initial field strength, we conclude that it is this initial magnetisation that determines the long-term growth (or lack thereof) of the magnetic field. We produce only a weak large-scale poloidal field in `b200` despite evolving for a very long time ($4.0 \times 10^5\ GM/c^3$). We do, qualitatively, produce a large-scale poloidal field in `b5`, and broadly reproduce the results of Liska et al. (2020) whose simulation features a comparable strong toroidal field. This finding shows that the results of Liska et al. (2020) can be reproduced in a pseudo-Newtonian framework and do not rely on the inclusion of General Relativistic physics or the peculiarities of spinning black holes.

All simulations produced clear—although disordered—butterfly patterns in the sign of $\langle B_\phi \rangle$, indicating quasi-periodic cycles of the MRI dynamo. We additionally find that *all* simulations, both weak and strong fields, establish a dominant signed field in the jet region, however, in `b200` and `b200_hi` this field is much weaker than that of the disk, and so may be easily missed if quantities are not plotted with care.

Additionally, although neither simulation produces sufficient horizon-crossing magnetic flux to exceed the MAD transitional value of $\phi = 50$ for a non-spinning BH, our initially-strong-field case `b5` comes close at $\phi \sim 40$, much higher than is typically seen in SANE disks. In disks such as those considered here where the magnetic flux (and magnetic field strength through the magnetic plasma-β) appears saturated but the MRI simultaneously appears active, the question of SANE vs MAD relies heavily on the definition of a MAD state, for which there is still heated debate (e.g. Begelman et al., 2022).

Our fiducial high β simulation `b200` produced an overdensity at early times that persisted throughout the run. No comparable overdensity was observed in the high-resolution counterpart `b200_hi`. The impact of the overdensity on disk evolution is mixed: while

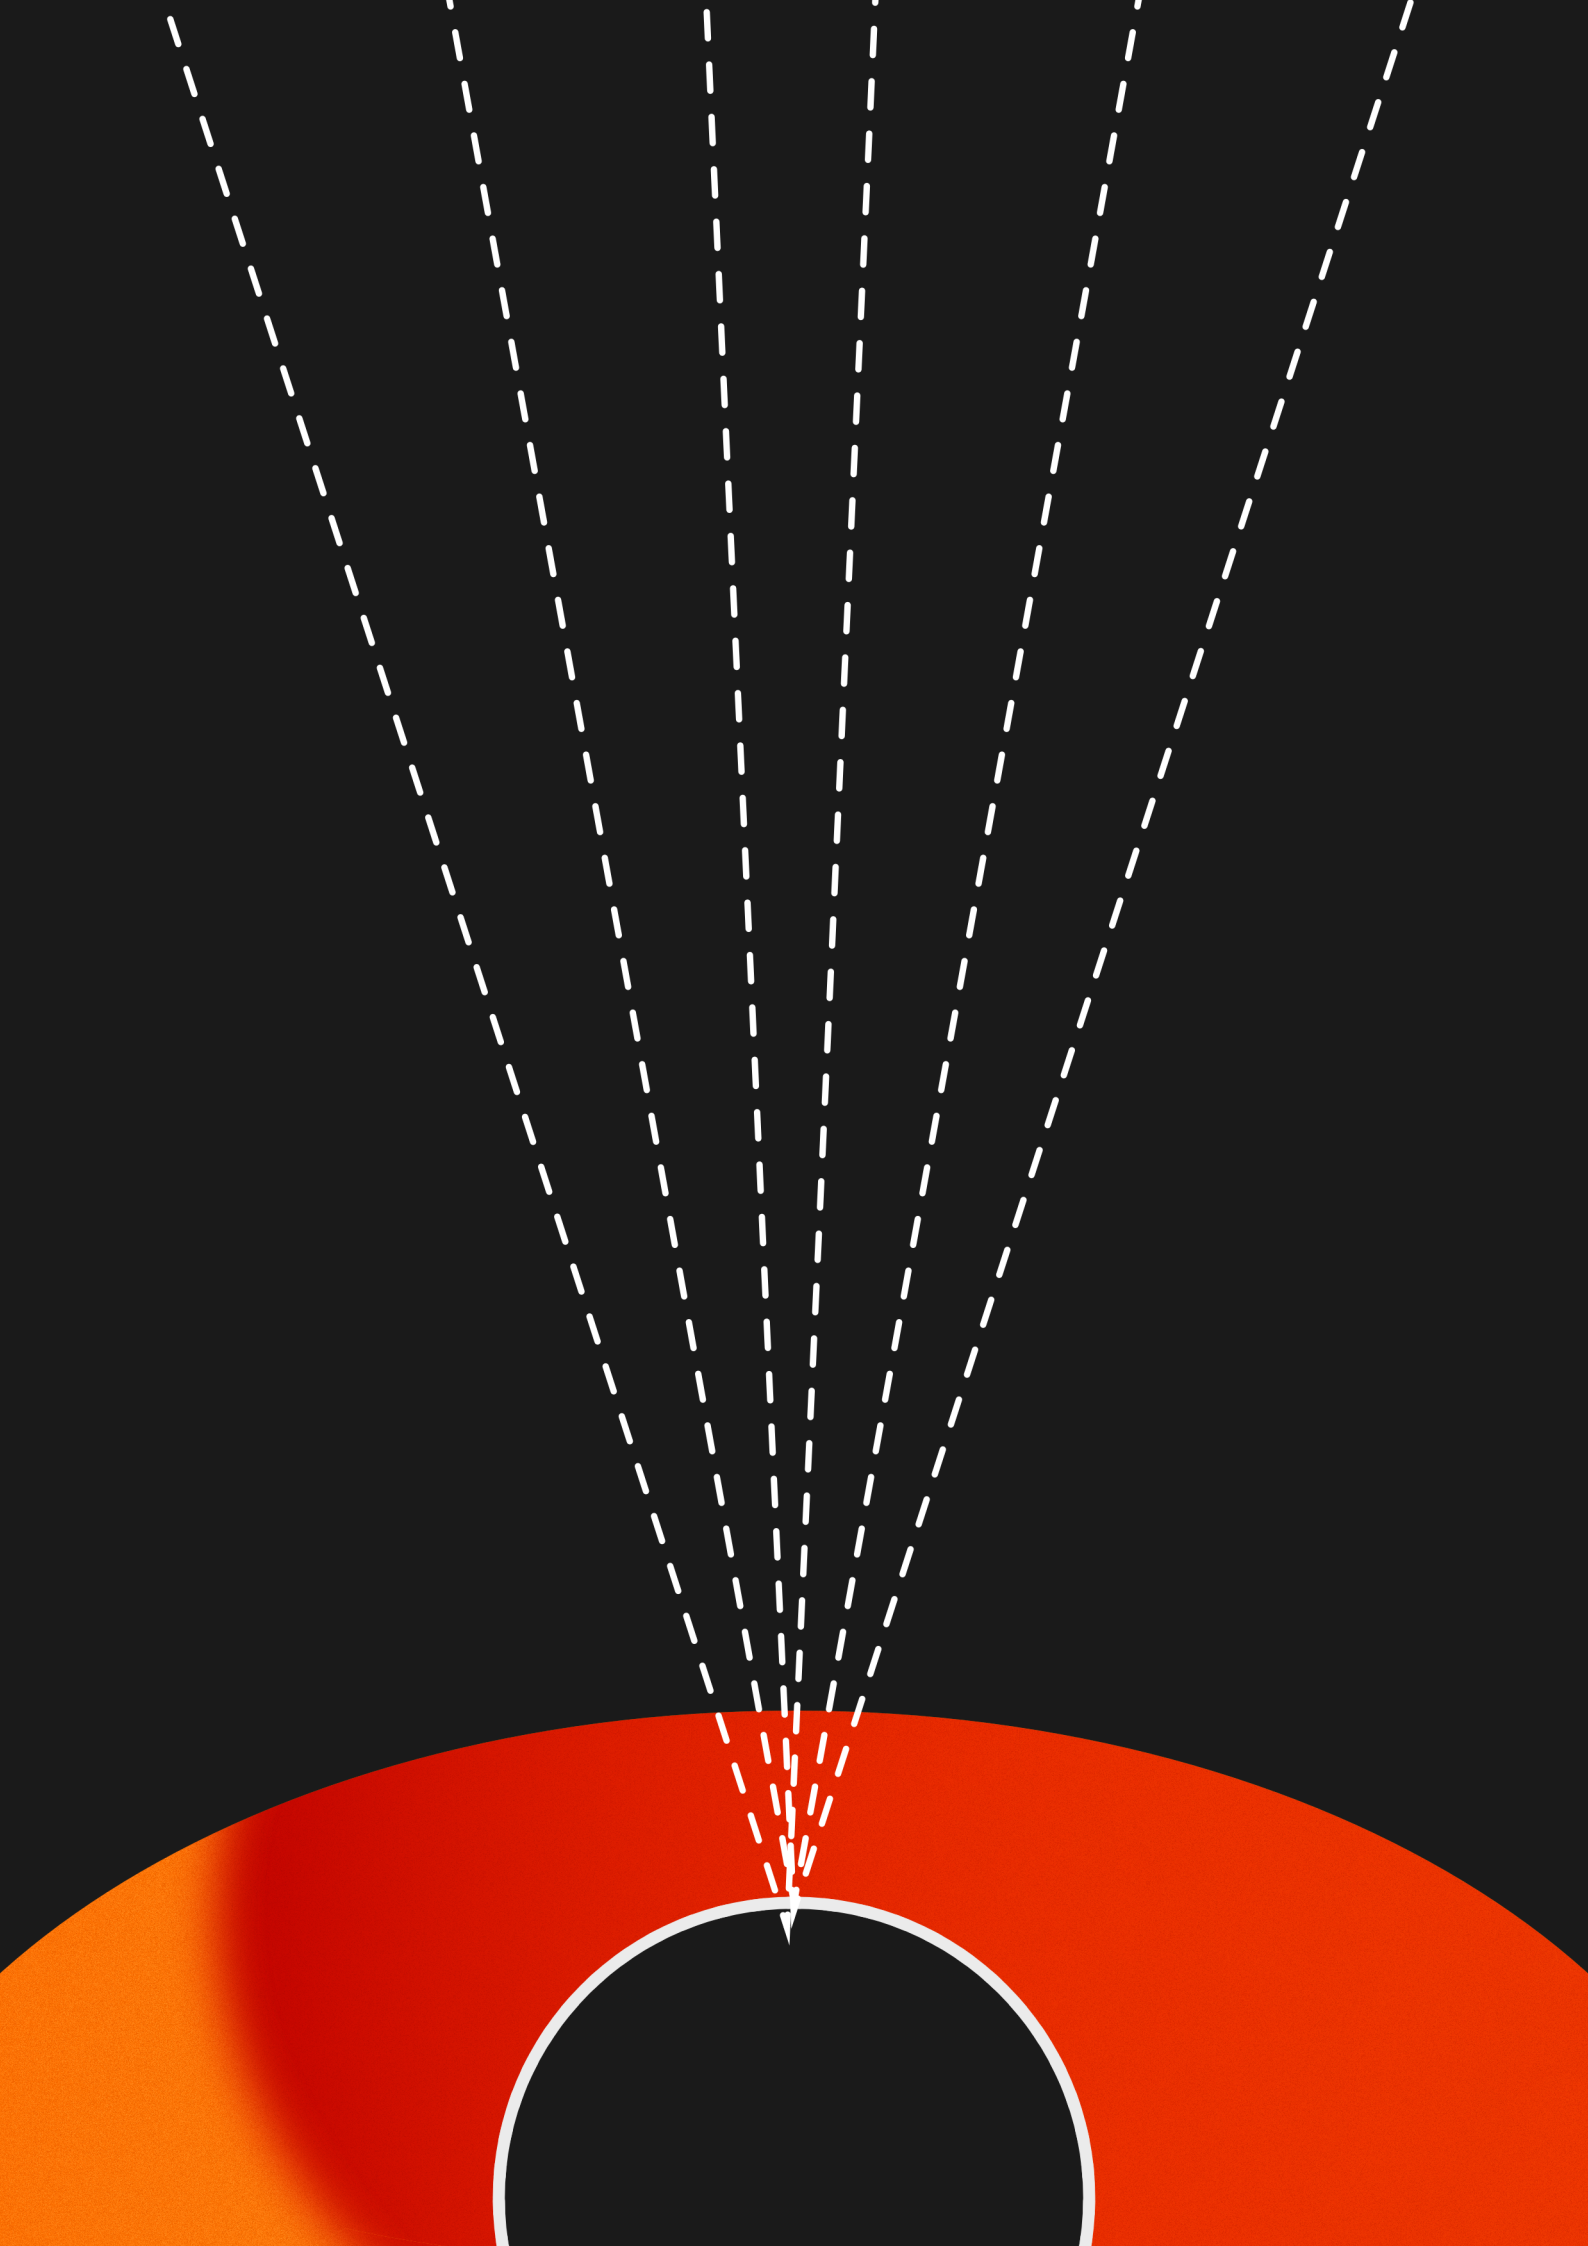
there are clear changes in the disk scale height, there are no clearly observable differences in the magnetic field strength or structure according to the metrics used (magnetic plasma β and magnetic flux Φ). Nonetheless, this finding highlights the need for great care when choosing a simulation mesh, even at high resolutions.

Our results suggest that disk metrics like the magnetic plasma β and magnetic flux Φ are not inherent to our disks but are instead influenced by the initial magnetic conditions—in our case, the initial strength of the toroidal field—supporting the conclusions made by White et al. (2020). Such a finding has implications for both simulated and real accretion disks. It has been previously hoped that, given enough time, a simulated accretion disk eventually forgets its initial conditions and settles to a common stable state. Previous work with NVF and ZNVF poloidal fields has suggested that this is not the case, and our results now also extend this to the purely toroidal case for RIAFs. It is also possible that different modes of material inflow (e.g. feeding via a binary companion vs the ambient medium) could lead to different final states in real accretion disks. Further work utilising more realistic initial conditions and environments may help to quantify whether, and under what conditions, real RIAFs are sensitive to these changes.

Taken at face value, the results of models presented in this work suggest that initial magnetic field conditions can drive highly dissimilar evolution in AGN with otherwise identical physical properties, e.g. BH spin, BH mass, and size of the disk matter reservoir. If this is the case, then significant variations should be observed across AGN with similar non-magnetic properties, with a possible dependence on the method of disk feeding. Additionally, although some poloidal field is naturally generated over the duration of our simulations, neither `b200` nor `b5` or any of the high resolution counterparts achieve a MAD state. If these findings are replicated at higher resolution, it may suggest that disks fed by material with a predominantly toroidal configuration (e.g. in TDEs) are less likely to enter the MAD regime.

Data Availability: Raw data were generated at the Cambridge Service for Data Driven Discovery. Derived data supporting the findings of this study are available from the corresponding author P.E.R. on request. The problem generator and input scripts associated with the simulations are archived in Zenodo (Rodman, 2023a). The software used to analyze data and generate plots is available on GitHub (`athena-analysis` codebase: https://github.com/paytonrodman/athena-analysis) under a 3-Clause BSD License and version 1.0 is archived in Zenodo (Rodman, 2023b).

Software: Athena++ (Stone et al., 2020), NumPy (Harris et al., 2020), pandas (McKinney, 2010), SciPy (Virtanen et al., 2020), VisIt (Childs et al., 2012).



Chapter 3

Mean-field dynamo coefficients in thick accretion disks with initially toroidal fields

3.1 Introduction

The launching of relativistic jets via the Blandford-Znajek (e.g. Blandford & Znajek, 1977) mechanism requires an ergosphere magnetic field with a significant poloidal or vertical component, however, the magnetic field contributed to the accretion disk through accumulation of passive IGM gas does not meet these requirements, and so the disk field must be first amplified and reoriented through a so-called dynamo process. The most well-described dynamo process in the accretion disk is that of the $\alpha\Omega$ dynamo, in which differential rotation shears poloidal fields into toroidal (the Ω-effect) and a secondary action, most likely analogous to magnetic buoyancy, transforms the toroidal back to poloidal (the α-effect). As jet driving is most sensitive to poloidal field, the α-effect is of significant interest and much effort has been made to characterise and quantify this effect, both in shearing box (e.g. Brandenburg & Donner, 1997; Brandenburg et al., 1995; Davis et al., 2010; Gressel, 2010; Gressel & Pessah, 2015; Ziegler & Rüdiger, 2001) and global (e.g. Dhang et al., 2020; Flock et al., 2012; Hogg & Reynolds, 2018a) simulations. However, due to computational constraints, many disk simulations are initialised with fields already containing a significant poloidal component. There is increasing concern that such initial conditions (ICs) are unrealistic and may lead to biased results. Very little is known of the evolution of the α-effect from a purely toroidal IC.

There are many ways to measure the dynamo α, including, as we saw in Chapter 1 Section 1.3, the First Order Smoothing Approximation (FOSA), also sometimes referred to as the Second Order Correlation Approximation (SOCA), under which we make use of Mean-Field Electrodynamics and an assumption of kinematics to arrive at a relationship between the electromotive force (EMF, ε) and the mean magnetic field ($\overline{\mathbf{B}}$):

$$\varepsilon_i = \alpha_{ij}\overline{\mathbf{B}}_j + \beta_{ijk}\partial_k\overline{\boldsymbol{B}}_j + \dots$$

Here, the pseudo-tensor α_{ij} is the embodiment of the α-effect, while β_{ijk} acts as an additional effective viscosity term which is generally considered subdominant to the α-effect.

To make analysis simpler, the α_{ij} term can be split into symmetric and antisymmetric parts

$$\alpha_{ij} = \alpha^S_{ij} + \alpha^A_{ij}.$$

Of these components, the symmetric α^S_{ij}, and in particular its diagonal components diag(α^S_{ij}),

$$\text{diag}(\alpha^S_{ij}) = \begin{bmatrix} \alpha_{rr} & & \\ & \alpha_{\theta\theta} & \\ & & \alpha_{\phi\phi} \end{bmatrix},$$

are the most important for determining large-scale field growth, with the $\alpha_{\phi\phi}$ term being responsible for replenishing the poloidal field from the toroidal field.

There are numerous methods in employment for measuring the dynamo coefficients, and in particular the α_{ij} coefficient, including the minimal τ approximation (MTA; e.g. Brandenburg & Subramanian, 2005), the test-field method (TF; e.g. Brandenburg et al., 2009; Schrinner et al., 2007), and measurement from Direct Numerical Simulation (DNS; e.g. Dhang et al., 2020; Hogg & Reynolds, 2018a). Each method has associated costs and benefits. The minimal τ approximation and test-field method both require one to impose a simple external magnetic field from which a response can then be measured and the α coefficient derived. DNS methods instead employ post-processing of existing disk simulation data, and for that reason are the more flexible approach and the one taken in this work. In the DNS method, one directly measures the turbulent EMF ε' from the fluctuating velocity $\mathbf{u}'$ and fluctuating magnetic field $\mathbf{b}'$ as

$$\varepsilon' = \langle \mathbf{u}' \times \mathbf{b}' \rangle, \tag{3.1}$$

which then allows one to calculate the associated α_{ij} via Equation (3.1).

Shearing box simulations by Brandenburg et al. (1995) were the first to verify such a correlation between $\langle \varepsilon_\phi \rangle$ and $\langle B_\phi \rangle$, on the order of

$$\langle \varepsilon_\phi \rangle = (-2, 4) \times 10^{-3} \langle B_\phi \rangle \langle \boldsymbol{u}^2 \rangle^{1/2} \tag{3.2}$$

where the brackets indicate average values above (the Northern Hemisphere, NH) and below (the Southern Hemisphere, SH) the midplane.

As suggested by the Brandenburg et al. (1995) results, the diag(α_{ij}^S) terms are known to flip sign across hemispheres, however there is disagreement on what the 'correct' polarity should be. Local shearing box studies (e.g. Brandenburg et al., 1995; Davis et al., 2010) found negative $\alpha_{\phi\phi}$ in the upper hemisphere and positive in the lower hemisphere. Global simulations are more mixed, with Flock et al. (2012) finding positive $\alpha_{\phi\phi}$ in the upper hemisphere, Hogg & Reynolds (2018a) finding weakly negative, and the thicker disks of Dhang & Sharma (2019) having no clear dominant sign. The precise reasons for this asymmetry are not yet known, however, some hints may be found in early work on the solar dynamo.

Arguments first applied to explain latitudinal migration of $B_{\rm tor}$ in the solar dynamo (see, e.g. Brandenburg et al., 1995; Stix, 1976; Yoshimura, 1975, and references therein) tell us that the combination $\alpha_{\phi\phi} \partial u_\phi / \partial r > 0$. For Keplerian rotation, $\partial u_\phi / \partial r < 0$, and so there must be $\alpha_{\phi\phi} < 0$ in the NH and positive in the SH in an accretion disk. However, the picture is more complicated than it first appears. The sign of $\alpha_{\phi\phi}$ should additionally be opposite to the sign of the helicity $\langle \boldsymbol{u} \cdot \nabla \times \boldsymbol{u} \rangle$ in a given hemisphere (c.f. Equation (1.55)). Early applications of this theory to accretion disks by Brandenburg et al. (1995) produced the expected sign of $\alpha_{\phi\phi}$ (negative) but the wrong sign for helicity (also negative), perhaps indicating that more parameters set hemisphere sign. In particular, they identified that the current helicity $\langle \boldsymbol{j}' \cdot \boldsymbol{b}' \rangle$ had a positive sign in the NH, so by Equation (1.68) could affect the sign of α if the current helicity had greater magnitude than the flow helicity.

Further evidence comes from another aspect of solar physics–the equatorward migration of sunspots, itself thought to be a sign of latitudinally-travelling dynamo waves (Charbonneau, 2010). A similar analysis of wave solutions can be performed for accretion disks. Per Brandenburg & Subramanian (2005), within the context of the $\alpha\Omega$ dynamo for a disk, in cylindrical coordinates (r, θ, z), with mean velocity $\overline{\boldsymbol{u}} = (0, rS, 0)$ and shear $S = -3/2\Omega$,

we can find wave solutions to the induction equation of the form

$$B_r = B_0 \sin[k_z(z - ct)] \tag{3.3}$$

$$B_\theta = \sqrt{2} B_0 \left| \frac{c}{\alpha} \right| \sin[k_z(z - ct) + \varphi] \tag{3.4}$$

for an unknown amplitude B_0, phase shift $\varphi = -3\pi/4$, and phase speed c,

$$c = \frac{\omega_{\text{cyc}}}{k_z} = \frac{\alpha S}{|2\alpha S k_z|^{1/2}} = \text{sign}(\alpha S)\eta_T k_z. \tag{3.5}$$

For a Keplerian disk $S < 0$ and, hence, if $\alpha < 0$ then $c > 0$, indicating wave propagation away from the midplane. These wave solutions are thought to be the origin of the so-called 'butterfly' patterns seen in the movement of sunspots (e.g. Charbonneau, 2010) and the toroidal field component in accretion disks (Figure 2.11), and hence a negative $\alpha_{\phi\phi}^{\text{NH}}$ is considered indicative of an $\alpha\Omega$ dynamo at work. α_{ij} becomes noisier with increasing disk thickness (e.g. Hogg & Reynolds, 2018b), which may partly explain why butterfly patterns also become irregular in thicker disks.

In this work, we perform a direct numerical calculation of the diagonal, symmetric components of the α_{ij} tensor under mean-field dynamo theory, using data from the fiducial simulations of Chapter 2. These simulations are of thick ($H/R \sim 0.3$–0.4) accretion disks initialised with a purely toroidal magnetic field with plasma beta of $\beta = 5$ or $\beta = 200$ (termed `b5` and `b200` respectively), and as such provide a valuable test for the evolution of α_{ij} in the absence of an initial poloidal field.

Section 3.2 describes the calculation of the dynamo coefficients under mean-field theory and introduces the two spatial averaging methods used to define the 'mean' in 'mean'-field theory. In Section 3.3, we review the calculated α_{ii} and associated offsets C_i, comparing their magnitude, sign, and symmetry to those in the literature, and discuss their implications. We conclude in Section 3.4.

3.2 Calculating the mean-field dynamo coefficients

In Chapter 2, we performed long-timescale MHD simulations of thick disks, which we can now study under mean-field dynamo theory.

As we have seen previously, the fluctuating EMF $\boldsymbol{\varepsilon}'$ can be expressed as a simple linear expansion around $\overline{\mathbf{B}}$ in the form of (3.1),

$$\boldsymbol{\varepsilon}'_{\boldsymbol{i}} = \alpha_{ij}\overline{\boldsymbol{B}}_j + \beta_{ijk}\partial_k\overline{\boldsymbol{B}}_j + \dots$$

where α_{ij} and β_{ijk} are pseudo-tensors. If the scale of $\overline{\boldsymbol{B}}$ is sufficiently large compared to the scale of fluctuations (typically characterised by the eddy scale), then higher order derivatives are expected to be smaller than that of β_{ijk} and so are frequently ignored (Brandenburg & Subramanian, 2005; Moffatt, 1978). Our primary area of interest is the growth of a large-scale poloidal field from an initial toroidal field, so we will focus on calculating the diagonal components of the symmetric part of α_{ij} (Equation 1.57), namely α_{rr}, $\alpha_{\theta\theta}$, and $\alpha_{\phi\phi}$. Hence Equation (3.1) becomes

$$\boldsymbol{\varepsilon}'_{\boldsymbol{i}} = \alpha_{ii}\overline{\boldsymbol{B}}_i \tag{3.6}$$

for components $i = (r, \theta, \phi)$.

By calculating both $\boldsymbol{\varepsilon}'_i$ and $\overline{\boldsymbol{B}}_i$ within our accretion disks, we can, therefore, retrieve an estimate for the associated α_{ii}.

The fluctuating EMF $\boldsymbol{\varepsilon}'_{\boldsymbol{i}}$ of the flow can be determined using the fluctuating components of the velocity $\boldsymbol{u}'$ and magnetic field $\boldsymbol{b}'$ as $\boldsymbol{\varepsilon}' = \langle \boldsymbol{u}' \times \boldsymbol{b}' \rangle$. For our three components, this comes to

$$\begin{aligned}
\boldsymbol{\varepsilon}'_r &= \boldsymbol{u}'_\theta \boldsymbol{b}'_\phi - \boldsymbol{u}'_\phi \boldsymbol{b}'_\theta, \\
\boldsymbol{\varepsilon}'_\theta &= -(\boldsymbol{u}'_r \boldsymbol{b}'_\phi - \boldsymbol{u}'_\phi \boldsymbol{b}'_r), \\
\boldsymbol{\varepsilon}'_\phi &= \boldsymbol{u}'_r \boldsymbol{b}'_\theta - \boldsymbol{u}'_\theta \boldsymbol{b}'_r.
\end{aligned} \tag{3.7}$$

3.2.1 Fluid variable averaging

The fluctuating component X' of some fluid flow variable X can be calculated directly by subtracting the flow average,

$$X' = X - \langle X \rangle, \tag{3.8}$$

where $\langle X \rangle$ is some suitable average. In accretion studies, the average quantity $\langle X \rangle$ is generally taken over one or two coordinate directions. For 3D simulations, the average is typically taken over the azimuthal direction, $\langle X \rangle = 1/2\pi \int_0^{2\pi} X d\phi$. Following Hogg & Reynolds (2018a), we perform azimuthal averages over the full azimuthal range ($\phi \in [0, 2\pi]$ in our case), radially between $r \in [15r_g, 15r_g + 15H]$, and vertically from $\theta \in [H, 2H]$ above and $\theta \in [-H, -2H]$ below the disk midplane.

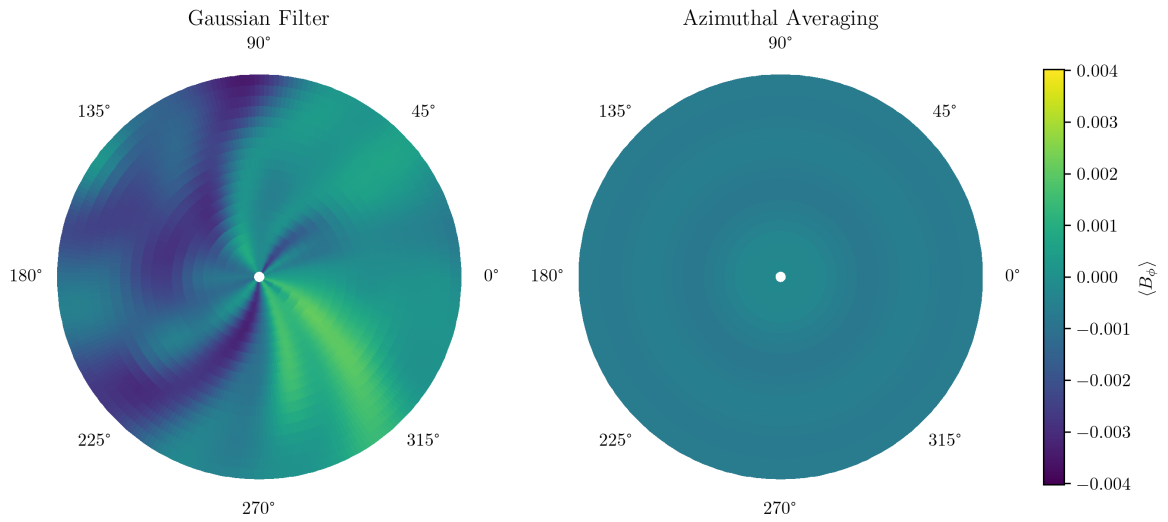


Fig. 3.1 The effect of different azimuthal averaging methods upon the variable B_ϕ in the upper hemisphere of `b200` at late time. Data is taken for $r \in [15r_g, 15r_g + 15H]$, $\theta = 1.5H$, and $\phi \in [0, 2\pi]$ A 1D Gaussian filter with standard deviation $\sigma = 5$ (*left*) preserves more of the local structure than simple azimuthal averaging (*right*).

This method is helpful as it uses the system's natural azimuthal symmetry; however, if the flow has significantly different values within different regions in ϕ, a simple azimuthal average may destroy this information and the resulting derived α_{ii} would be excessively noisy. To account for this possibility, we additionally calculate averages by a 1D Gaussian filter in the ϕ-direction with a standard deviation $\sigma = 5$ cells. Values are wrapped in ϕ, and the filter is truncated at $4\sigma = 20$ cells. The Gaussian filter has the effect of 'blurring' the underlying dataset, turning $\langle X \rangle$ into a local average (Figure 3.1). As such, Gaussian filtering may be helpful in the analysis of `b200`, which features a ϕ variation in the form of a large lump that is expected to cause u_ϕ to deviate significantly from Keplerian in the region, hence affecting $\langle \boldsymbol{u} \rangle$ and $\boldsymbol{u}'$. It bears noting that Gaussian filtering does not obey $\langle X \rangle = \langle \langle X \rangle \rangle$ under Reynolds averaging rules; however, it is known that the tenets of mean-field theory are not strictly held in accretion disks generally—for example, $\boldsymbol{u}'$ and $\boldsymbol{b}'$ are not statistically independent—and are even less applicable in the case of significant azimuthal variations as seen in `b200`. By using an averaging method that acknowledges the deviations from Ω_K, we may recover the underlying correlations in the kinematic, mean-field theory (and hence α_{ii}) better.

With $\langle \boldsymbol{u} \rangle$ and $\langle \boldsymbol{B} \rangle$ determined using either averaging method described above, the turbulent EMF components are then calculated as in Equation (3.7) for every point within $r \in [15r_g, 15r_g + 15H]$, $\theta \in [\pm H, \pm 2H]$, and $\phi \in [0, 2\pi]$. This range returns 360 spatial

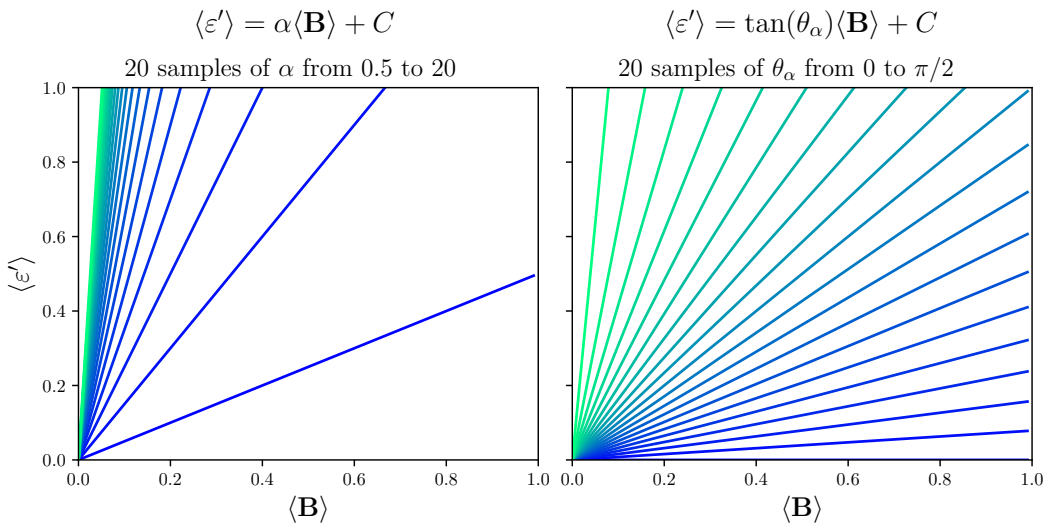


Fig. 3.2 Sampling in α (*left*) and $\tan(\theta_\alpha)$ (*right*) space, with $C_i = 0$. Uniform sampling in α leads to large gaps in ε', B-space as the slope range $m = \alpha$ is unbounded at ∞.

data points per timestep under azimuthal averaging as cells in the ϕ direction are averaged over, but 92,160 under Gaussian filtering. For ease of plotting, a random selection of 500 points is taken from the Gaussian-averaged datasets, although the calculation of α_{ii} and C_i proceeds on the full dataset.

3.2.2 Fitting for α_{ii}

Once the EMF $\boldsymbol{\varepsilon}'_i$ has been calculated for each cell in the range specified for a given timestep, they are similarly averaged (azimuthally, or via Gaussian filtering) and the correlation between $\langle\boldsymbol{\varepsilon}'_i\rangle$ and $\langle\boldsymbol{B}_i\rangle$ is determined for each timestep.

We assess two methods of fitting: (1) simple linear least-squares fitting on the form

$$\langle\boldsymbol{\varepsilon}'_i\rangle = \alpha_{ii}\langle\boldsymbol{B}_i\rangle + C_i, \tag{3.9}$$

where C is some vertical offset, and (2) fitting via a Bayesian Markov Chain Monte Carlo (MCMC) method on the form

$$\langle\boldsymbol{\varepsilon}'_i\rangle = \tan(\theta_\alpha)\langle\boldsymbol{B}_i\rangle + C_i. \tag{3.10}$$

with flat (uniform) priors on θ_α and C_i. By fitting for θ_α rather than α directly, we can sample the space more uniformly (Figure 3.2).

We find that the more expensive MCMC fitting does not return a superior, or even different, fit to that of simple linear least-squares. For this reason, we revert to using linear least-squares to determine α_{ii} (Equation 3.9).

An example distribution for $\langle \mathcal{E}'_\phi \rangle$ and $\langle \boldsymbol{B}_\phi \rangle$ under both azimuthal averaging and Gaussian filter averaging, along with the best linear fit, are shown for a single timestep for `b5` in Figures 3.3.

We note that there is a significant scatter in the distribution of $\langle \mathcal{E}'_i \rangle$ and $\langle \boldsymbol{B}_i \rangle$, with the resulting linear fit often lying close to horizontal. When averages are taken via Gaussian filtering, the centroid of the data lies closer to the origin and with greater variation in $\langle \boldsymbol{B}_i \rangle$. In some cases, the different averaging methods lead to an opposite signed slope.

With these fits, we calculate α_{ii} and C_i for our fiducial `b200` and `b5` simulations, in both hemispheres (upper and lower), for all times $T > 50\ T_{\rm ISCO}$ to allow for disk settling. The resulting α_{ii} and C_i are shown in Figures 3.4 (azimuthal) and 3.6 (Gaussian) for `b200` and Figures 3.5 (azimuthal) and 3.7 (Gaussian) for `b5`. We see that, over time, the distribution for α_{ii} and C_i narrow, with values in the NH often mirroring that of the SH, indicating a change in sign of α_{ii} across hemispheres.

3.3 Results and discussion

The estimated values for α_{ii} and C_i are given in Table 3.1 for azimuthal averaging, and Table 3.2 for Gaussian filtering. Previous values for α_{ii} in shearing box and global accretion disk simulations are given in Table 3.3 for comparison and reference.

We retrieve α_{ii} magnitudes (~ 1–50×10^{-4}) that are in broad agreement with previous works, with offset C_ϕ comparable to that found by Hogg & Reynolds (2018a) who performed a similar fit to $\alpha_{\phi\phi}$ for disk thickness up to $H/R = 0.4$.

Under azimuthal averaging, $\alpha_{\theta\theta}$ has the greatest magnitude of the three components, although its value decreases under Gaussian filtering. Dhang et al. (2020) found a similarly higher (relative to α_{rr} and $\alpha_{\phi\phi}$) absolute magnitude for $\alpha_{\theta\theta}$ under the Singular Value Decomposition (SVD) method, although it is highly disordered and incoherent. $\alpha_{\theta\theta}$ and α_{rr} relate to the generation of toroidal field from poloidal under the α-effect, but this process is expected to be subdominant to the Ω-effect (shear) in accretion disks, and so the interpretation of these two coefficients—both their magnitude and sign—is uncertain. Notably, both `b200` and `b5` return similar α_{ij} magnitudes under Gaussian filter averaging, suggesting that the local strength of the α-effect is largely insensitive to magnetic field strength.

Simulation ID	Component	Value	C_i
		(upper, lower) [$\times 10^{-4}$]	(upper, lower) [$\times 10^{-8}$]
b5	α_{rr}	$(22.9^{\pm 2.49}, -23.0^{\pm 2.59})$	$(105^{\pm 22.7}, -45.6^{\pm 22.8})$
	$\alpha_{\theta\theta}$	$(56.1^{\pm 5.26}, -39.6^{\pm 5.15})$	$(15.2^{\pm 28.4}, -36.8^{\pm 26.9})$
	$\alpha_{\phi\phi}$	$(-3.6^{\pm 0.392}, 2.61^{\pm 0.393})$	$(11.6^{\pm 9.41}, 9.81^{\pm 9.33})$
b200	α_{rr}	$(5.88^{\pm 1.96}, -4.25^{\pm 2.02})$	$(6.58^{\pm 5.57}, -22.6^{\pm 5.74})$
	$\alpha_{\theta\theta}$	$(-37.0^{\pm 4.74}, 41.5^{\pm 4.83})$	$(-24.7^{\pm 8.27}, -4.92^{\pm 8.84})$
	$\alpha_{\phi\phi}$	$(1.3^{\pm 0.23}, -1.39^{\pm 0.24})$	$(-0.135^{\pm 2.06}, 0.484^{\pm 2.29})$

Table 3.1 Mean values of α_{ii} and C_i with 95 per cent CI, for given simulation and dimension, using azimuthal averaging.

Simulation ID	Component	Value	C_i
		(upper, lower) [$\times 10^{-4}$]	(upper, lower) [$\times 10^{-8}$]
b5	α_{rr}	$(18.5^{\pm 0.34}, -19.0^{\pm 0.33})$	$(36.0^{\pm 4.60}, -10.9^{\pm 4.57})$
	$\alpha_{\theta\theta}$	$(1.56^{\pm 0.61}, -4.4^{\pm 0.61})$	$(7.46^{\pm 5.61}, -7.71^{\pm 5.80})$
	$\alpha_{\phi\phi}$	$(-2.54^{\pm 0.044}, 2.55^{\pm 0.045})$	$(-1.19^{\pm 1.66}, 3.38^{\pm 1.70})$
b200	α_{rr}	$(18.0^{\pm 0.28}, -17.1^{\pm 0.28})$	$(-0.57^{\pm 1.4}, -3.80^{\pm 1.4})$
	$\alpha_{\theta\theta}$	$(8.23^{\pm 0.51}, -7.58^{\pm 0.51})$	$(-5.2^{\pm 1.8}, 0.012^{\pm 1.9})$
	$\alpha_{\phi\phi}$	$(-2.03^{\pm 0.03}, 1.95^{\pm 0.03})$	$(0.28^{\pm 0.44}, 1.2^{\pm 0.45})$

Table 3.2 Mean values of α_{ij} and C_i with 95 per cent CI, for given simulation and dimension, using the Gaussian filter method.

We include results for all three components here for completeness, but we will focus primarily on $\alpha_{\phi\phi}$ whose impact, while not at all clear, is clearer still than that of α_{rr} and $\alpha_{\theta\theta}$.

Reference	Study type	Component	Magnitude $[\times 10^{-4}]$	Sign in NH
Brandenburg et al. (1995)	Shearing box	$\alpha^{\dagger}_{yy}$	$20-40$	$-$ve
Brandenburg & Donner (1997)	Shearing box	$\alpha_{\phi\phi}$	$1-10$	$-$ve
Ziegler & Rüdiger (2001)	Shearing box	$\alpha^{\dagger}_{yy}$	$5-40$	$-$ve
Davis et al. (2010)	Shearing box	$\alpha^{\dagger}_{yy}$	$30-50$	$-$ve
Gressel (2010)	Shearing box	$\alpha^{\dagger}_{yy}$	$\sim 100^{\ddagger}$	$+$ve
		$\alpha^{\dagger}_{xx}$	$\sim 100-200^{\ddagger}$	$-$ve
Gressel & Pessah (2015)	Shearing box	$\alpha^{\dagger}_{yy}$	$\sim 50^{\ddagger}$	$+$ve
		$\alpha^{\dagger}_{xx}$	$\sim 50^{\ddagger}$	$+$ve
Hogg & Reynolds (2018a)	Global	$\alpha_{\phi\phi}$	$1-2$	$-$ve
Flock et al. (2012)	Global	$\alpha_{\phi\phi}$	$20-30$	$+$ve
Dhang et al. (2020)	Global	$\alpha_{\phi\phi}$	$\sim 10-100^{\ddagger}$	$-$ve
		α_{rr}	$\sim 10-250^{\ddagger}$	$+$ve
		$\alpha_{\theta\theta}$	$\lesssim 300^{\ddagger}$	$+$ve

Table 3.3 Values of α_{ii} from literature. † Where Cartesian coordinates are used, $xx \equiv rr$ and $yy \equiv \phi\phi$.‡ Visual estimate, value not stated in the text.

We find that all three components are antisymmetric under both averaging methods, although the direction of this antisymmetry changes for `b200` $\alpha_{\theta\theta}$ and $\alpha_{\phi\phi}$ depending on the averaging method. For `b200` under azimuthal averaging, $\alpha^{\rm NH}_{\theta\theta} < 0$ and $\alpha^{\rm NH}_{\phi\phi} > 0$, while for Gaussian filtering $\alpha^{\rm NH}_{\theta\theta} > 0$ and $\alpha^{\rm NH}_{\phi\phi} < 0$. For `b5`, the sign of $\alpha^{\rm NH}_{\theta\theta}$ and $\alpha^{\rm NH}_{\phi\phi}$ is consistent under either averaging method. Due to the known irregularities of the `b200` simulation, these differences support the idea that azimuthal averaging is inappropriate for significant azimuthal variations in the flow.

Under Gaussian filter averaging, we find negative $\alpha_{\phi\phi}$ in the NH and positive in the SH. This is in agreement with most previous shearing box (Brandenburg & Donner, 1997; Brandenburg et al., 1995; Davis et al., 2010; Ziegler & Rüdiger, 2001) and global (Dhang et al., 2020; Hogg & Reynolds, 2018a) simulations, with the exception of Gressel & Pessah (2015) and Gressel (2010) (shearing box) and Flock et al. (2012) (global). The expected sign of $\alpha_{\phi\phi}$ remains a debated topic, and the reason for the discrepancy is unknown.

A change in the sign of $\alpha_{\phi\phi}$, either in space or time, could be caused by a change in the relative importance of kinetic and magnetic α, where in Chapter 1 Equation (1.68) we saw,

$$\alpha = \alpha_K + \alpha_M$$

where

$$\alpha_K = -\frac{\tau_c}{3}\left(\langle \boldsymbol{u}' \cdot \nabla \times \boldsymbol{u}' \rangle\right)$$
$$\alpha_M = \frac{\tau_c}{3}\left(\rho^{-1}\langle \boldsymbol{j}' \cdot \boldsymbol{b}' \rangle\right)$$

are the kinematic and magnetic α. The kinematic α_K ignores the Lorentz back-reaction of the magnetic field on fluid flow and is concerned only with the helicity of the velocity field. The magnetic α_M, on the other hand, relies on the current helicity density. Analysis of these components by Gressel (2010) found that α_M took on a positive value in the NH, with a fit matching closely to that predicted by Rüdiger & Kichatinov (1993) for stratified, rotating turbulence, indicating that a positive NH α-effect may be related to a Parker-type instability. They additionally found that when analysed during the linear growth phase of the MRI, α_M instead took on a negative value in the NH. In both cases $|\alpha_M| > |\alpha_K|$. Hence, the sign of α_M determines the combined α sign. Positive α^{NH} may indicate saturated MRI or α-quenching or represent nonlinear and inhomogeneous effects not accounted for in the predictions for α_K.

The value of offsets C_i is heterogeneous, varying significantly between hemispheres, components, and averaging methods. As we only fit for the components of $\mathrm{diag}(\alpha^S_{ij})$, an increased C_i could be interpreted as a greater contribution from the un-fitted components, namely the off-axis α^S_{ij}, the antisymmetric α^A_{ij}, and the diffusive term β_{ijk}, along with higher order components.

In Figures (3.4–3.7), we plot the 95 per cent point density contours for data within equally-spaced N_{orb} bins. By comparing these contours, we can gain some insight into the variation of α_{ii} over the lifetime of the simulation. In most cases, the contours are broader for earlier times, corresponding to a greater scatter in the offset C_i. Under the interpretation above, this may suggest that nonlinear, inhomogeneous, or diffusive components dominate the early-time dynamo.

We can see the time evolution of the spatially-averaged $\langle \alpha_{ii} \rangle$ and $\langle C_i \rangle$ in greater detail in Figures (3.8–3.9). The evolution of $\langle \alpha_{ii} \rangle$ is dominated by noise, although we can see some slight signed preferences which become clearer under Gaussian filtering—in particular, we

note that $\langle \alpha_{rr} \rangle$ shows a clear preference for positive NH, $\langle \alpha_{\phi\phi} \rangle$ shows a clear preference for negative NH, and $\langle \alpha_{\theta\theta} \rangle$ is disordered. In most cases, the noise level in $\langle \alpha_{ii} \rangle$ remains consistent with time, whereas the variation in $\langle C_i \rangle$ decreases in time—in other words, as the simulation progresses, the correlation between $\langle \boldsymbol{\mathcal{E}}'_i \rangle$ and $\langle \boldsymbol{B}_i \rangle$ is better described by the simple linear relationship $\langle \boldsymbol{\mathcal{E}}'_i \rangle = \alpha_{ii} \langle \boldsymbol{B}_i \rangle$. The only exception to this trend is `b200` under azimuthal averaging, which is not unexpected.

Given the long timescale over which the scatter in C_i decreases ($> 2000\ N_{\rm orb}$), this is unlikely to be due to initial disk settling and turbulence generation. It is also unlikely to be an artefact of the averaging method, as the same trend can be seen for `b5` under both azimuthal averaging and Gaussian filtering. Instead, the early dynamo is inherently less well-described by the kinematic mean-field theory. Under the constraints of our current fitting method, it is not possible to know which of the unfitted components (e.g. α^A_{ij}, β_{ijk}) controls this behaviour. The trend is much weaker in the `b200` simulation under Gaussian filtering and absent under azimuthal averaging. It is unclear if this results from the abnormal azimuthal variation of `b200` or if the weaker magnetic field suppresses the behaviour.

3.4 Conclusion

In this short chapter, I have presented the results of mean-field dynamo analysis on the fiducial simulated accretion disks of Chapter 2, `b200` and `b5`.

We applied the mean-field dynamo theory under the simple closure

$$\langle \boldsymbol{\mathcal{E}}'_i \rangle = \alpha_{ii} \langle \boldsymbol{B}_i \rangle + C_i$$

and solved for the diagonal dynamo coefficients $(\alpha_{rr}, \alpha_{\theta\theta}, \alpha_{\phi\phi})$ and offsets (C_r, C_θ, C_ϕ) using two different averaging schemes: (1) simple azimuthal averaging and (2) 1D Gaussian filtering.

Both methods are able to yield $\alpha_{\phi\phi}$ magnitudes of $|\alpha_{\phi\phi}| \sim (1\text{–}4) \times 10^{-4}$, tightening to $|\alpha_{\phi\phi}| \sim 2 \times 10^{-4}$ under Gaussian filtering, matching values previously found by Brandenburg & Donner (1997); Hogg & Reynolds (2018a), and slightly lower than those of e.g. Brandenburg et al. (1995); Davis et al. (2010); Dhang et al. (2020); Ziegler & Rüdiger (2001). Our results additionally confirm that all α_{ii} are antisymmetric across the midplane, although $\alpha_{\theta\theta}$ is highly disordered.

For the `b5` simulation, we consistently find:

- $\alpha_{\phi\phi}^{\mathrm{NH}} < 0$ matching previous shearing box (Brandenburg & Donner, 1997; Brandenburg et al., 1995; Davis et al., 2010; Ziegler & Rüdiger, 2001) and global (Dhang et al., 2020; Hogg & Reynolds, 2018a) simulations, and matching that expected from an $\alpha\Omega$-type dynamo,
- $\alpha_{rr}^{\mathrm{NH}} > 0$, matching that of Gressel & Pessah (2015) and Dhang et al. (2020), and
- $\alpha_{\theta\theta}^{\mathrm{NH}} > 0$, matching that of Dhang et al. (2020).

Simulation `b200` also returns a consistent $\alpha_{rr}^{\mathrm{NH}} > 0$, however its $\alpha_{\phi\phi}$ and $\alpha_{\theta\theta}$ are sensitive to the averaging method, with $\alpha_{\phi\phi}^{\mathrm{NH}} > 0$ and $\alpha_{\theta\theta}^{\mathrm{NH}} < 0$ under azimuthal averaging, and the opposite under Gaussian filtering. `b200` features a large density lump within the disk, and this azimuthal variation is expected to cause $\langle \boldsymbol{u}_\phi \rangle$ and $\langle \boldsymbol{u}'_\phi \rangle$ to deviate from their expected mean-field theory values (Ω_K and zero, respectively). For this reason, a more local average—1D Gaussian filtering—appears appropriate.

The signs and magnitudes of the spatially-averaged α_{ii} components are roughly stable in time, but the scatter in the offset C_i appears to decrease with increased simulation time for `b5`. Since we have fitted a simple linear equation of the form $\langle \boldsymbol{\varepsilon}'_i \rangle = \alpha_{ii} \langle \boldsymbol{B}_i \rangle + C_i$, the offset C_i will contain information on the unfitted dynamo components: the off-axis symmetric α_{ij}^S, the antisymmetric α_{ij}^A, the diffusive tensor β_{ijk}, and higher order terms in $\langle \boldsymbol{B}_i \rangle$. As we have only calculated diag(α_{ij}^S), it is not possible to know which unfitted component is ultimately responsible for the decreasing scatter in C_i. The signal is weaker or absent in `b200`; it is unclear if this is due to the known abnormalities of this simulation or to the weaker magnetic field. Future simulations of disks with initially weak, initially toroidal magnetic fields could help to tease apart these effects. Still, only a complete calculation of all mean-field components will determine the cause of the C_i variation within this model.

Our results support the notion that local averages (e.g. Gaussian filtering) can be used in place of azimuthal averaging to retrieve stable and consistent values for the diagonal, symmetric α_{ij} components; however, further work is needed to determine the ideal kernel or window size.

The work presented in this chapter is purely theoretical in nature, driven by the need to understand the behaviour of our simulated MHD disks at a deeper level. Direct observational probes of dynamo processes in black hole accretion disks require the ability to image magnetic field structures on sub-horizon scales, and, at minimum, require the ability to separate the magnetic field structure into separate hemispheres in order to observe the expected asymmetry in the α_{ij} coefficients. Such high resolution is not yet possible even

with the EHT observations of M87 and Sgr A, however, such observational data may become available in the future when space-based EHT capabilities are deployed.

Further theoretical progress in this field is limited both in the mathematical approach and computational resources required. Many methods used to measure the dynamo coefficients are unwieldy, impractical, and computationally expensive. Further, the fundamental theory on which such measurements rests depends on a number of assumptions that remain technically invalid in accretion disks, making it more difficult to interpret the calculated values in a physical sense. This is exemplified by the ongoing disagreement in hemisphere sign, and the corresponding disagreements over how such sign is determined in the disk.

As the dynamo processes at play are highly non-linear and turbulent, it is likely that future progress in this field will require analysis of simulated data over more dimensions than is currently achievable, although it may soon be within reach on the most powerful supercomputers. AI may additionally find some use, in particular dimensionality reduction through Principal Component Analysis (PCA) or similar methods, however any machine learning techniques must be used with care given the difficulty in verifying results.

Software: Athena++ (Stone et al., 2020), numpy (Harris et al., 2020), pandas (McKinney, 2010), scipy (Virtanen et al., 2020), seaborn (Waskom, 2021), emcee (Foreman-Mackey et al., 2013)

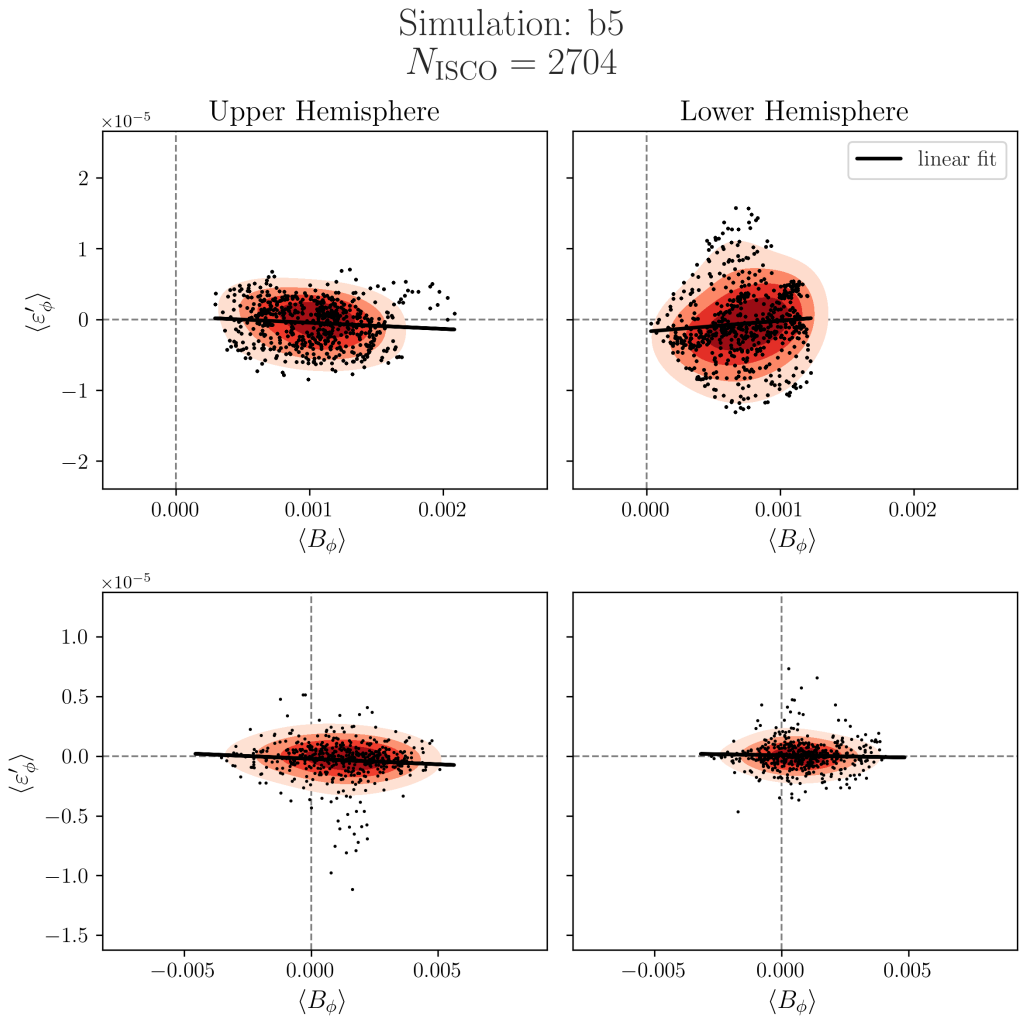


Fig. 3.3 Instantaneous $\langle \varepsilon'_\phi \rangle$-$\langle \boldsymbol{B}_\phi \rangle$ fit in the upper hemisphere (*left*) and the lower hemisphere (*right*), calculated via azimuthal averaging (*top*) and Gaussian filter averaging (*bottom*), for the `b5` simulation at $T = 2704\ T_{\mathrm{ISCO}}$. Data (black dots) is taken from the final output of the `b5` simulation, while the black line shows the best linear fit to the data. The underlying colour contours represent the point density, calculated via a kernel density estimator. The scatter in $\langle \varepsilon'_\phi \rangle$ decreases under Gaussian filter averaging, while the scatter in $\langle \boldsymbol{B}_\phi \rangle$ increases.

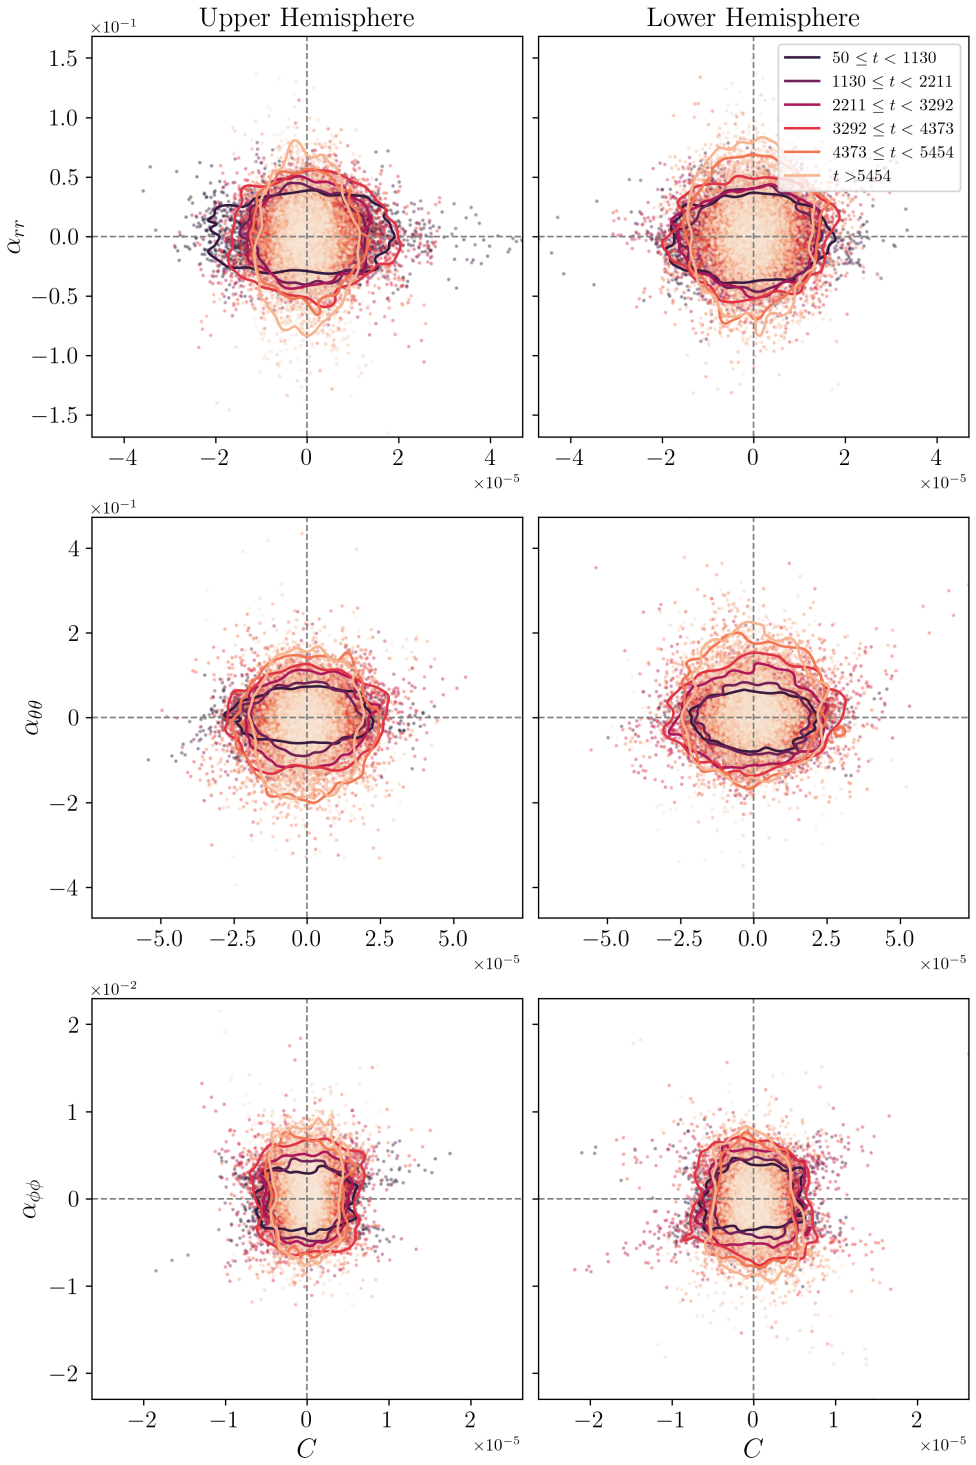


Fig. 3.4 Fitted α_{ii} and associated offset C_i values in r (*top*), θ (*middle*) and ϕ (*bottom*), for the `b200` simulation, upper (*left*) and lower (*right*) hemispheres, under azimuthal averaging. Colour denotes time, with dark points corresponding to early times and light points to late times. Contours encapsulate 95 per cent of the point density for the given time bin. In general, the scatter in α_{ii} decreases with time.

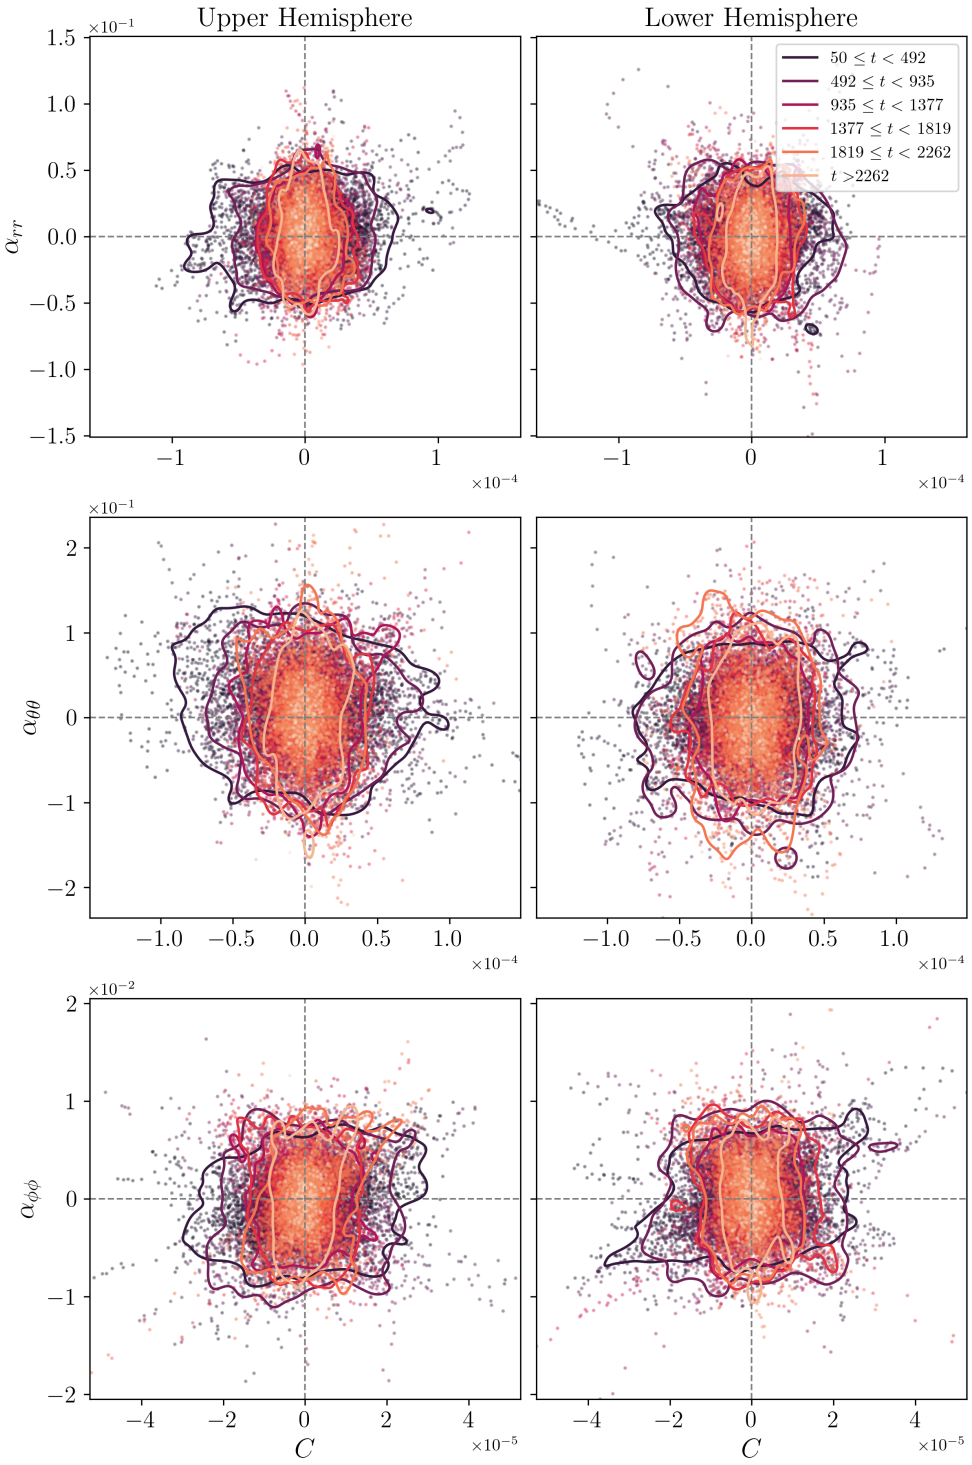


Fig. 3.5 Fitted α_{ii}, C_i values for the `b5` simulation with azimuthal averaging. Symbols as in Figure 3.4.

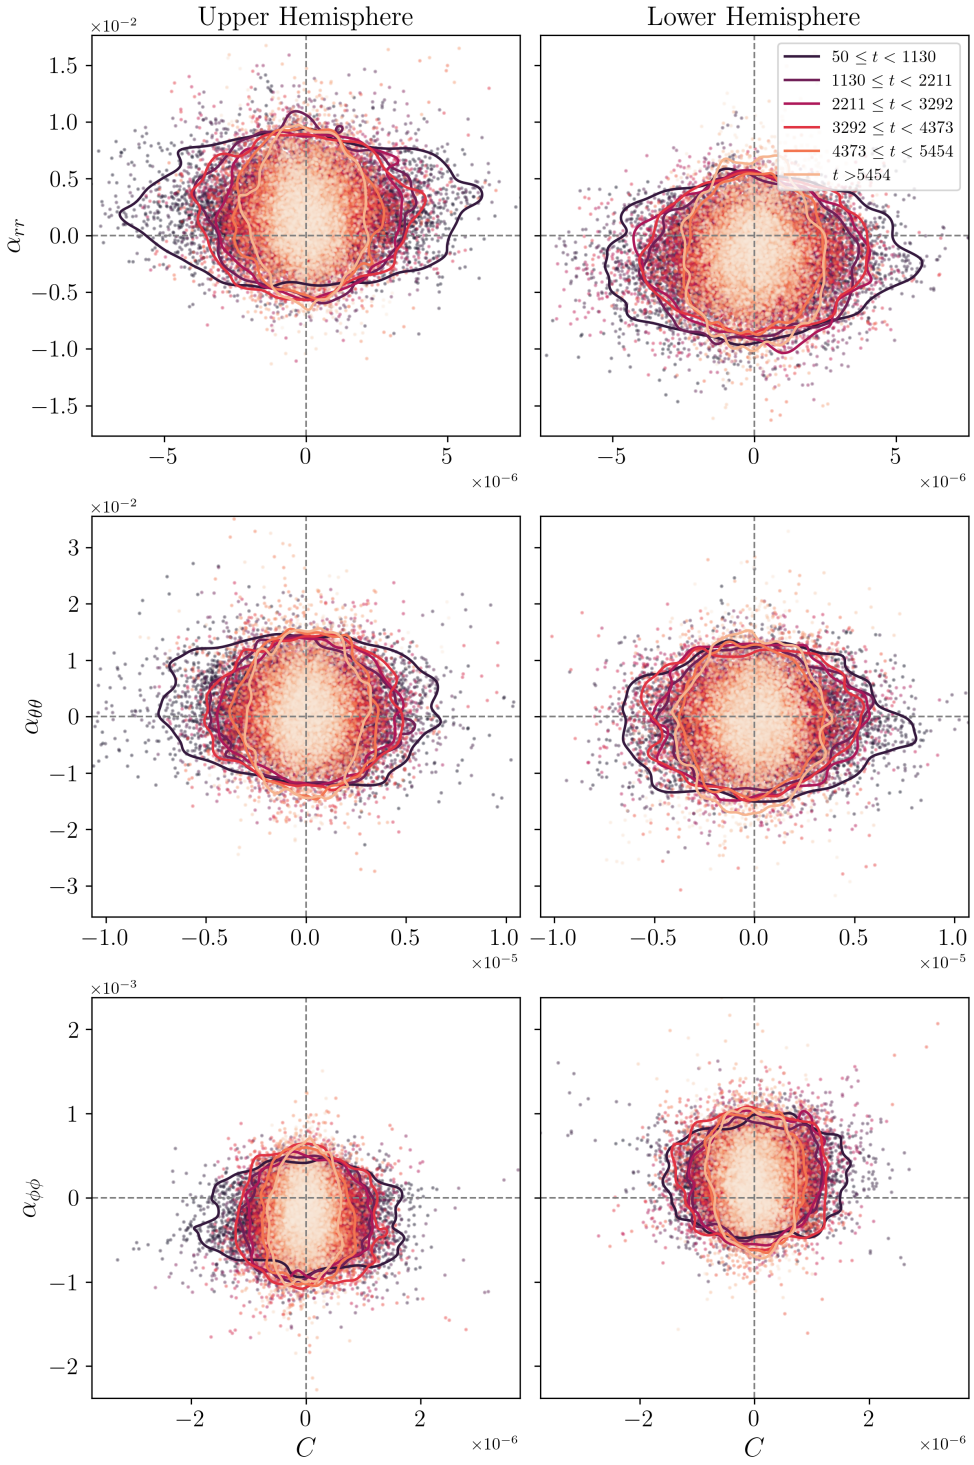


Fig. 3.6 Fitted α_{ii}, C_i values for the `b200` simulation with Gaussian filter averaging. Symbols as in Figure 3.4. Under Gaussian filter averaging, the hemisphere asymmetries in α_{ii} are more apparent, with the exception of $\alpha_{\theta\theta}$, which remains disordered.

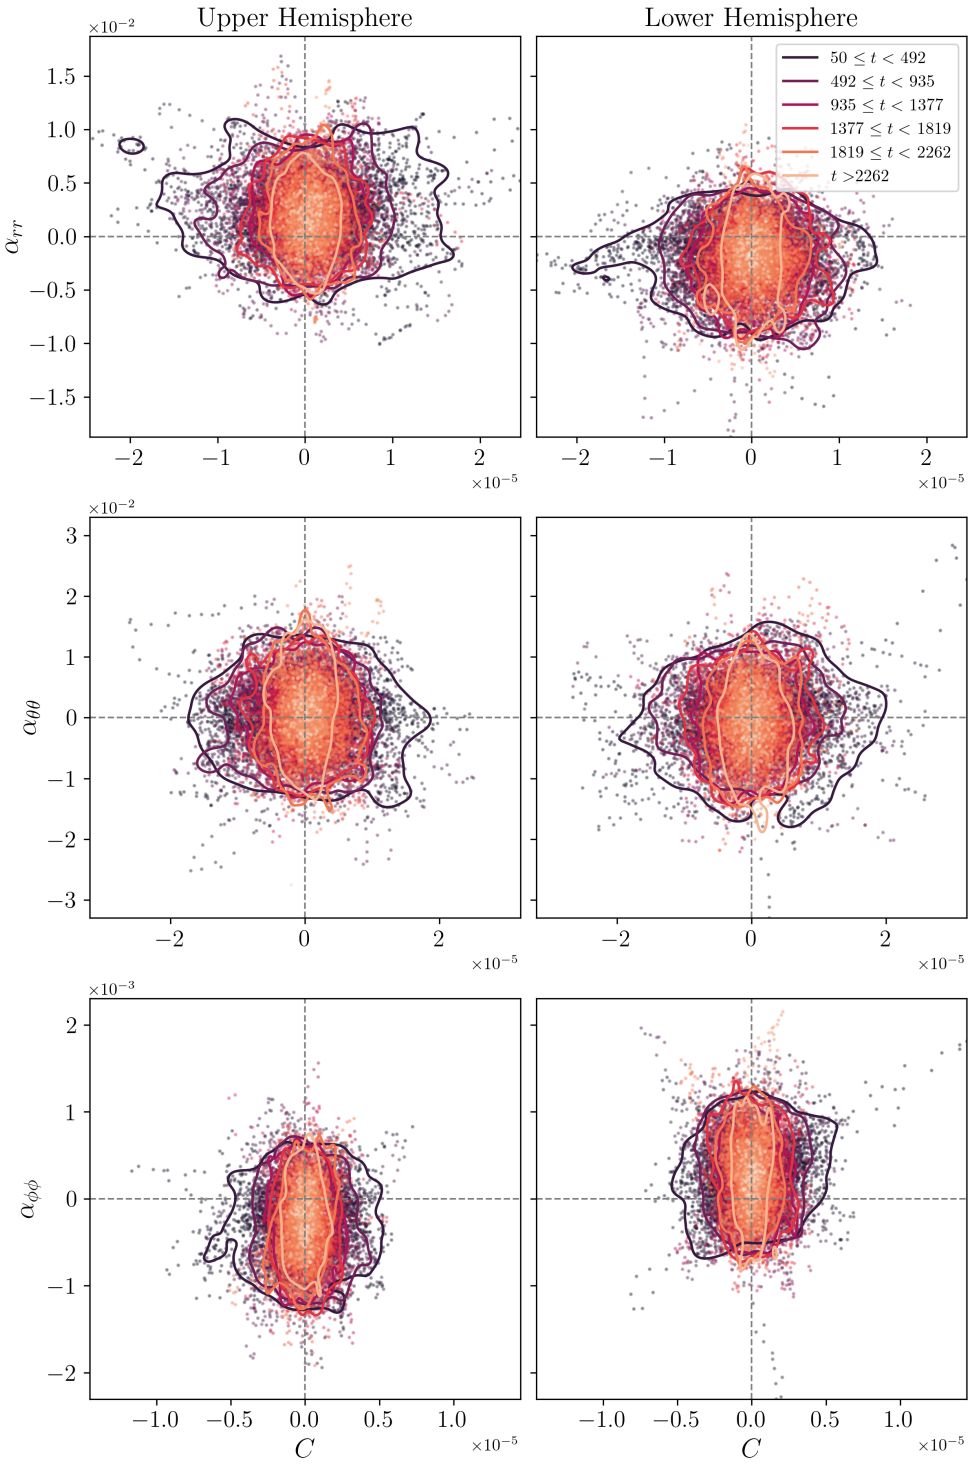


Fig. 3.7 Fitted α_{ii}, C_i values for the `b5` simulation with Gaussian filter smoothing. Symbols as in Figure 3.4.

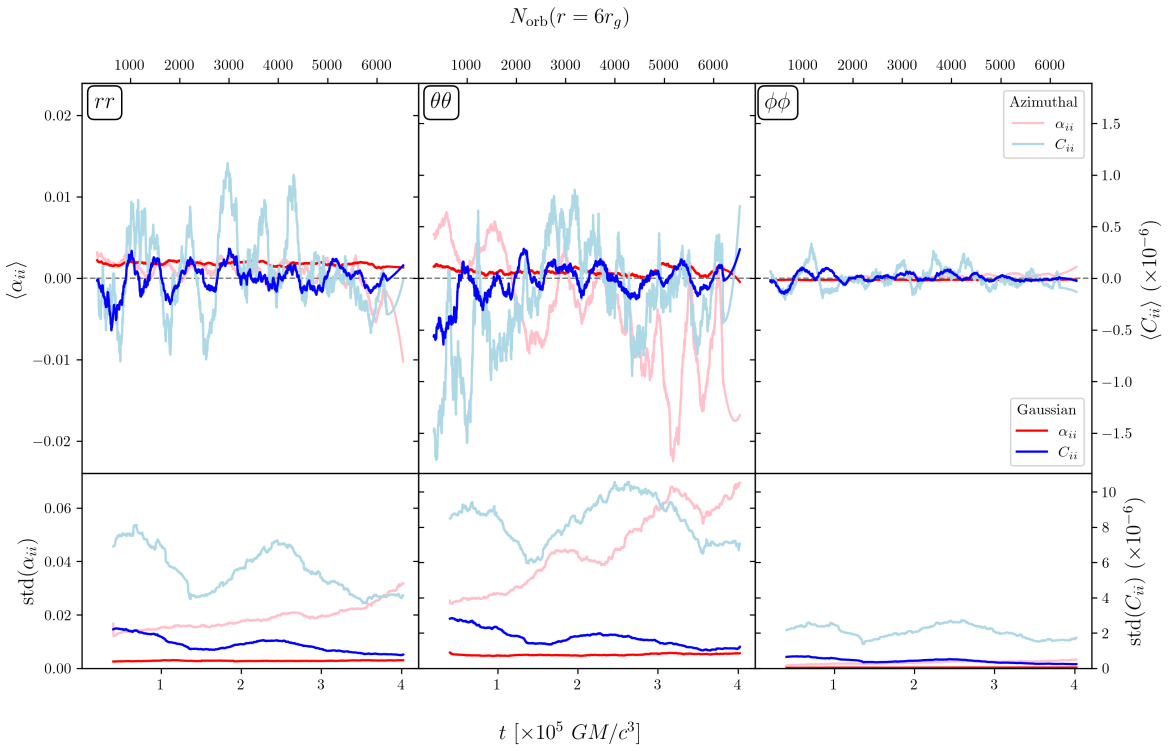


Fig. 3.8 (*Top*) Variation of $\langle\alpha_{ii}\rangle$ (blue/light blue) and $\langle C_i\rangle$ (red/pink) with time for the rr (*left*), $\theta\theta$ (*middle*), and $\phi\phi$ (*right*) components in the upper hemisphere for the `b200` simulation. Light colours correspond to measurements under azimuthal averaging, while the dark-coloured lines correspond to the Gaussian filtering method. The data has been smoothed with a Savitzky-Golay filter with window size $1/10$th $T_{\rm end}$ to aid in visualisation. A clear sign for $\langle\alpha_{ii}\rangle$ is more readily observed under Gaussian filter averaging. (*Bottom*) The variation in the standard deviation of $\langle\alpha_{ii}\rangle$ and $\langle C_i\rangle$ with time. Under both averaging methods, the variance in $\langle C_i\rangle$ decreases with time, with certain features coincident with the appearance of the overdensity/lump shortly after $\sim 10^5\ GM/c^3$. Under azimuthal averaging, the variance of $\langle\alpha_{ii}\rangle$ increases with time.

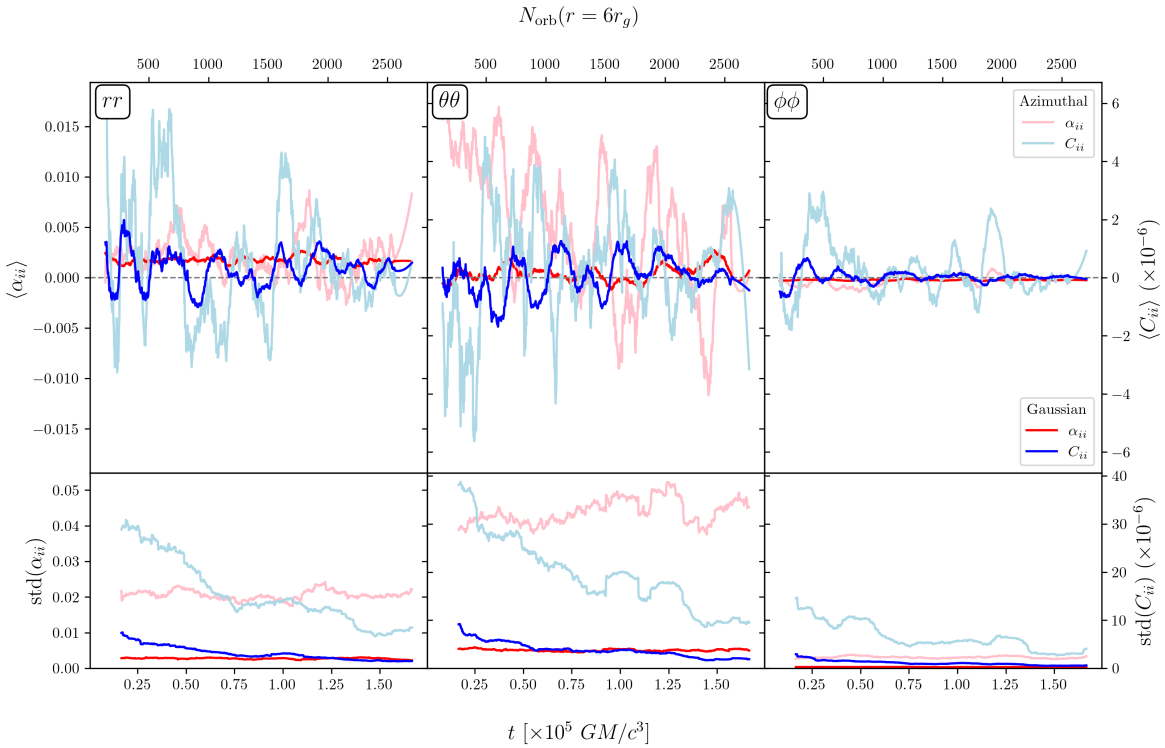
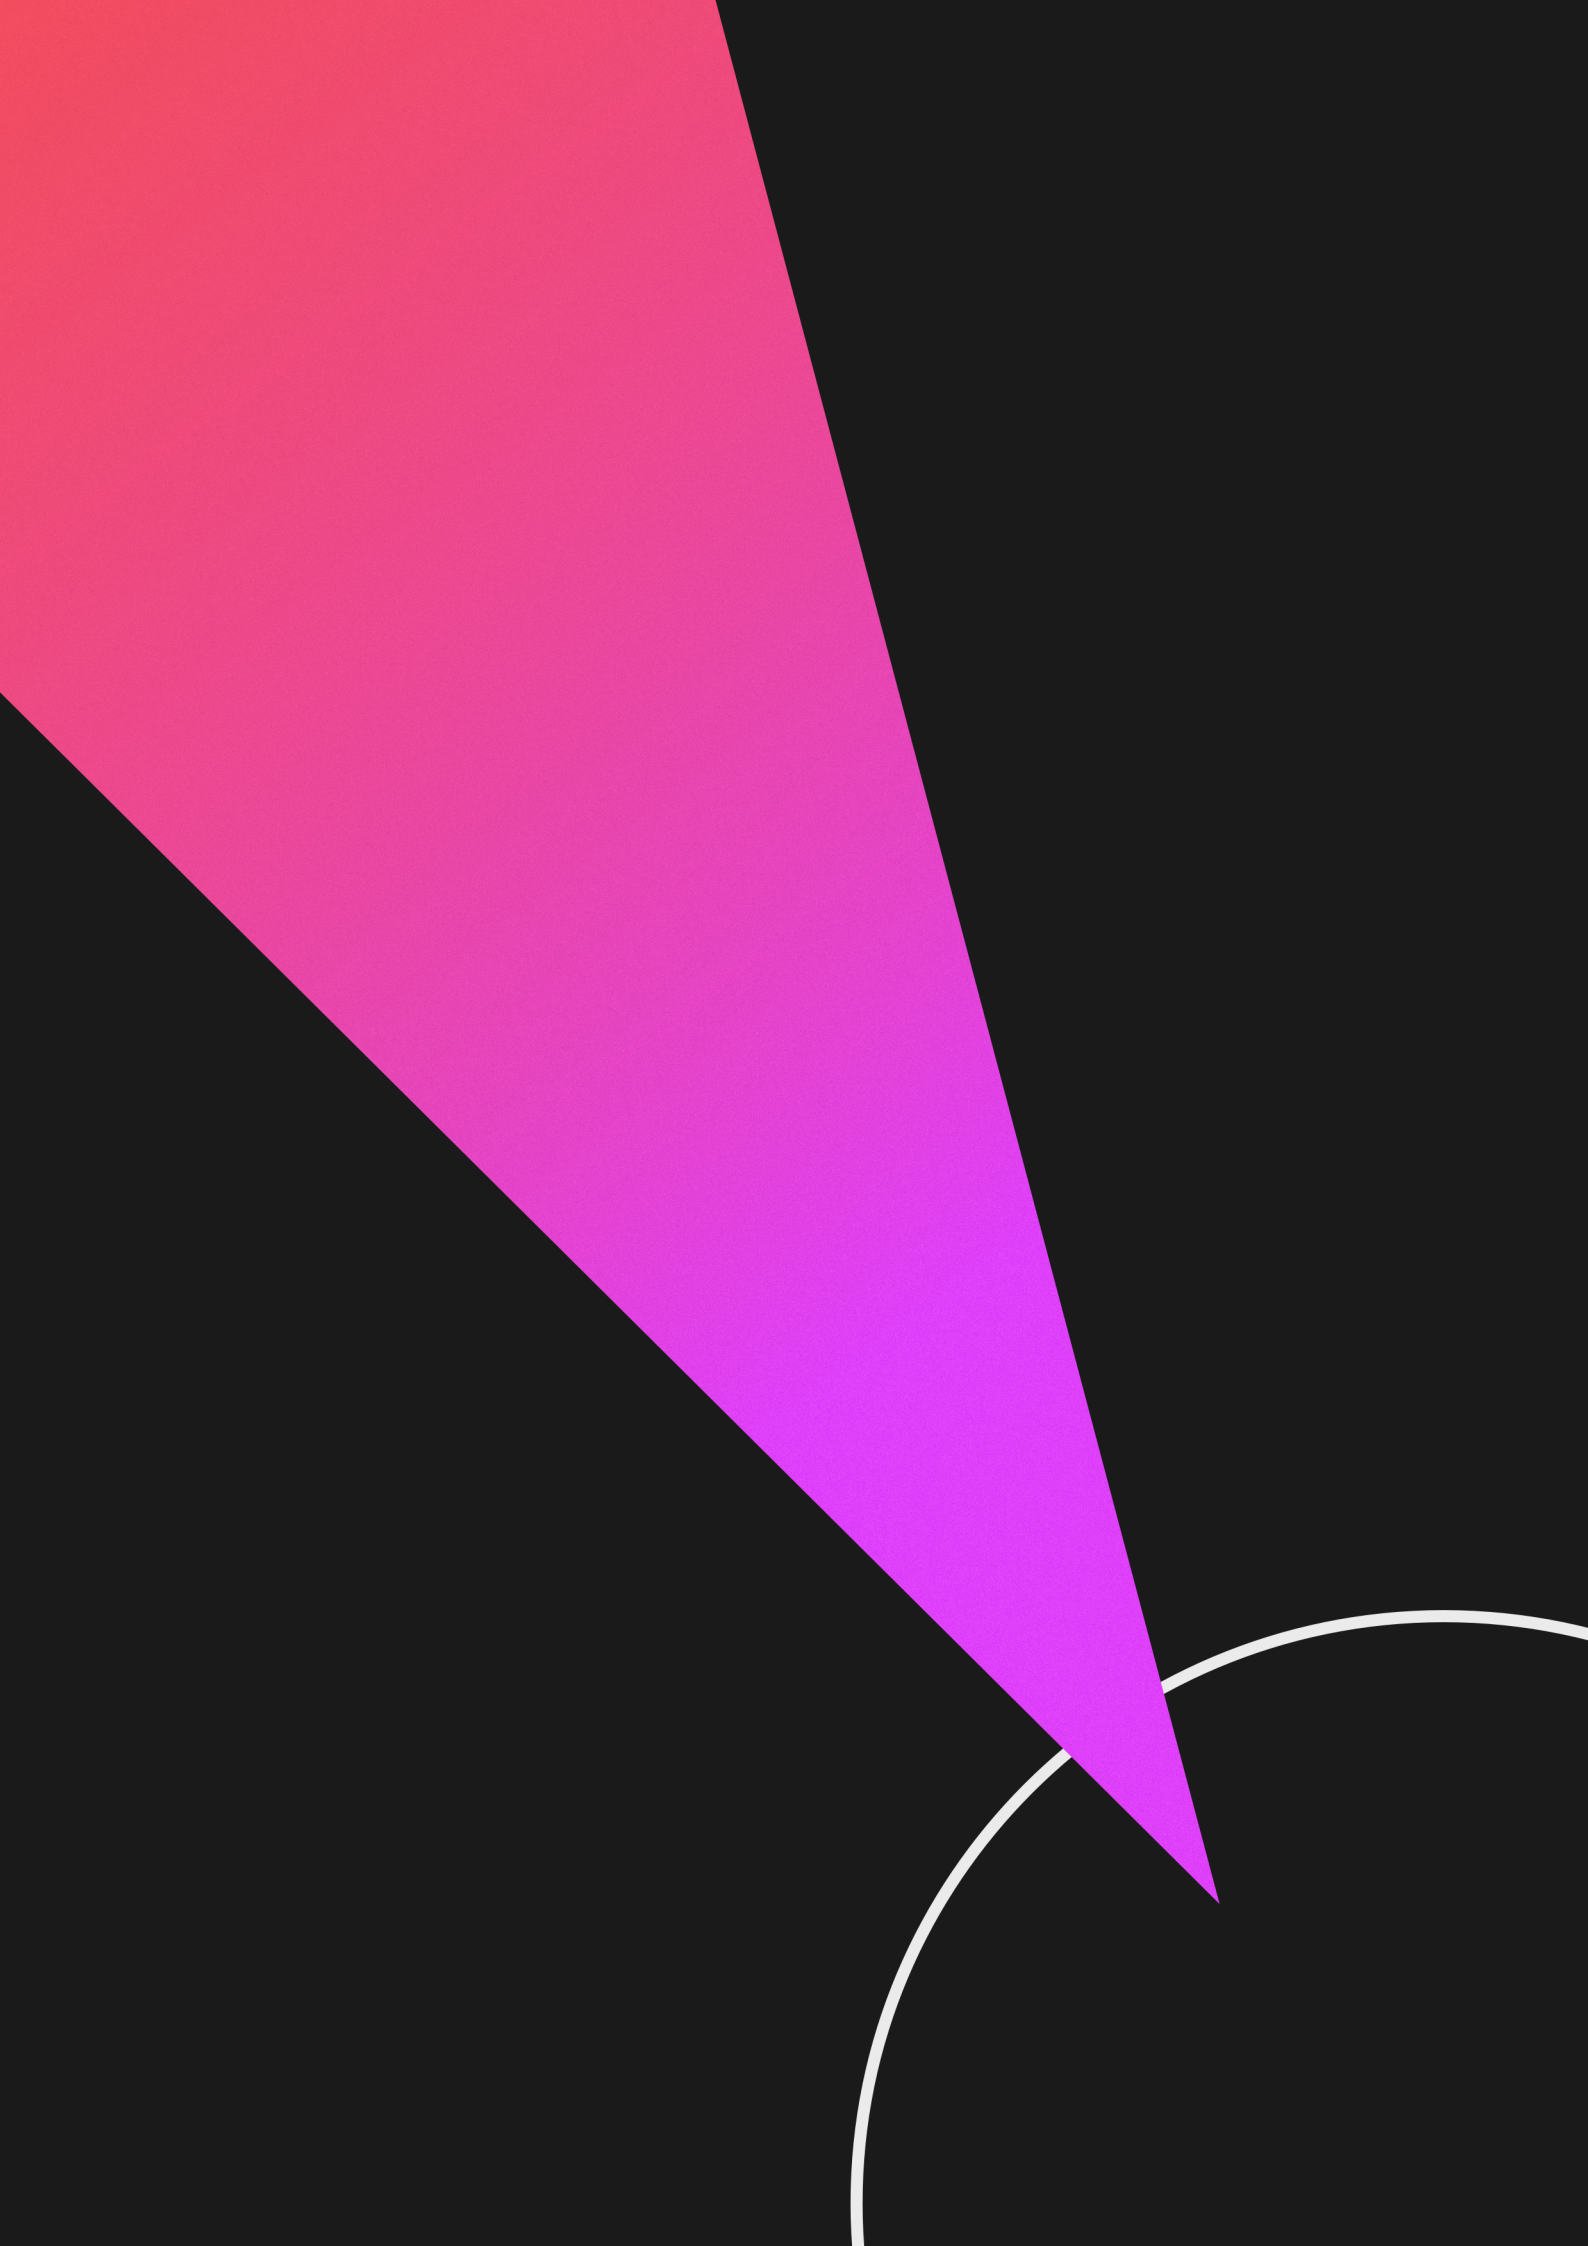


Fig. 3.9 As in Figure 3.8, but for simulation `b5`. Again, the variance in $\langle C_i \rangle$ decreases with time, while in the absence of `b200`'s overdensity/lump, the variance in $\langle \alpha_{ii} \rangle$ is now relatively consistent over the duration of the simulation.



Chapter 4

3D RMHD simulations of magnetically-linked relativistic jets

4.1 Introduction

Approximately 10–20 per cent of AGN are radio-loud (RL) and thought to play host to bipolar relativistic jets. Within this population exists a group of powerful radio sources with jets bent in a common direction away from the galaxy axis. Such sources are associated with both dominant Brightest Cluster Galaxies (BCGs) and field galaxies in groups, and their bent morphology is thought to relate to the bulk motion of the host galaxy through the intra-cluster or intra-group environment, which provides a crosswind that bends the jets via ram pressure. These bent or *head-tail* sources are classified as either Narrow Angle Tail (NAT) or Wide Angle Tail (WAT) (for review of NATs, see e.g. O'Dea & Baum, 2023) depending on the opening angle between the opposing jets, although the distinctions between them are murky due to resolution and projection effects in observation.

One such bent source is that of ESO 137-006, a WAT. Recent *MeerKAT* observations at 1030MHz by Ramatsoku et al. (2020) (Figure 4.1) show a series of highly collimated threads of synchrotron-emitting material; one filament (CS1) appears to pass through the external medium to connect the two lobes, another (CS3) appears to connect the inner part of one lobe with its outer part, and others still appear to extend from the lobe edges outwards into the external medium where they terminate or otherwise become undetectable. The authors suggest that these filaments may be magnetic in their nature and reflect either an interaction between the jet and environment magnetic field or, in the case of CS3, reconnection of the lobe field with the jet tail field. In either case, the filaments result from highly relativistic

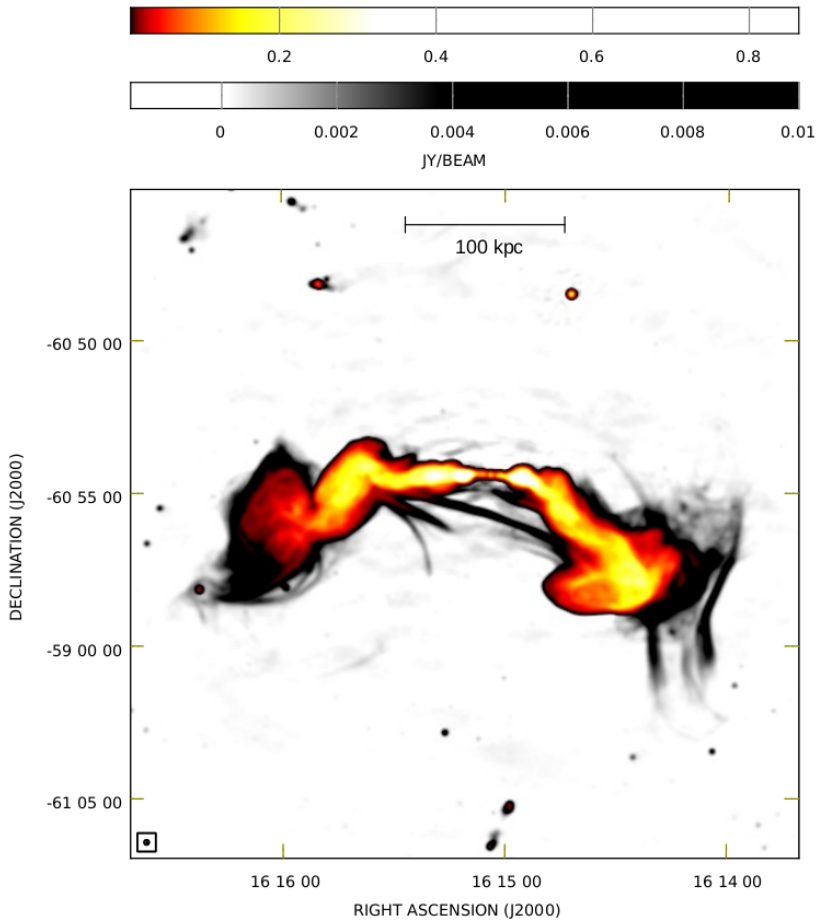


Fig. 4.1 Source ESO 137-006 imaged at 1030 MHz using MeerKAT. Filamentary structure is observed around the lobes at a radio brightness of < 10 mJy. The beam size of the observations is indicated by the circle in the lower left. A number of prominent filaments or threads appear to connect the two lobes. Reproduced from Ramatsoku et al. (2020) with permission.

synchrotron-emitting electrons that have escaped the main lobes and are flowing along (but are otherwise confined by) magnetic field lines.

Whilst filaments are well-known *within* radio lobes, a new class of low-frequency radio telescopes with increased sensitivity are now unearthing a collection of similar threads that appear to link to and extend *beyond* radio lobes, for example, as in radio galaxy 3C40B (Rudnick et al., 2022), BCG IC 711 (van Weeren et al., 2021), and radio galaxy IC 4296 (Condon et al., 2021), in addition to that of ESO 137-006 described above. As these filaments extend into the ICM, they provide a unique opportunity to explore ICM gas properties and magnetic fields and may act as a site for cosmic ray acceleration.

Previous simulations have attempted to capture the dynamics of head-tail sources, with jets that are both magnetised (e.g. Gan et al., 2017; Jones et al., 2017; Mendygral et al., 2012; Nolting et al., 2022; Ohmura et al., 2023; O'Neill et al., 2019a,b) and non-magnetised (e.g. Balsara & Norman, 1992; Loken et al., 1995; Morsony et al., 2013; Soker et al., 1988), however most simulations do not include an initial ICM field. One exception is Mendygral et al. (2012), whose ICM contains volume-averaged magnetic fields derived from the smoothed particle hydrodynamics (SPH) cluster simulations of Dolag et al. (2009). However, their derived magnetic field does not retain its filamentary structure—a feature that is likely important in producing the observed magnetic filaments. No head-tail simulation has yet included an ICM magnetic field generated self-consistently via MHD turbulence forcing, meaning the interaction between jet magnetic field and, for example, filaments in the ICM field is unexplored.

We carry out a series of (special) relativistic MHD (RMHD) simulations of light (density contrast $\rho_j/\rho_a = 10^{-3}$) jets with initially helical fields launched into a magnetised environment generated via forced turbulence and subjected to a strong transverse crosswind $v_w = 0.3$–$1c_s$ to investigate the interaction between jet and ICM magnetic fields and the production of collimated synchrotron threads (CSTs). The principal goal of this work is to determine whether magnetic connections can be made between the magnetic field in the radio lobes, and that in the ambient background medium. In essence, whether magnetic fields breach the 'contact discontinuity' that delineates the jet-injected matter and ambient matter. If so, the 'particle escape' hypothesis for the formation of the synchrotron filaments is viable.

In Section 4.2 we describe the numerical methods and simulation setup, including the generation of the turbulent background (Section 4.2.3) and jet launching (Section 4.2.3), and the main visualisation techniques employed (Section 4.2.5). We discuss the results of our simulations in Section 4.3, and summarise in Section 4.4.

Throughout this work, we adopt the standard convention $S \propto \nu^{-\alpha}$ for the spectral index α. This is not to be confused with the $\alpha_{\rm visc}$ of accretion disk theory (Chapter 2) or the α_{ij}'s of dynamo theory (Chapter 3).

Unit	Value	
ρ_0	$m_p \approx 1.673 \times 10^{-27}$	$[\mathrm{g/cm^3}]$
L_0	$1\ \mathrm{kpc} \approx 3.086 \times 10^{21}$	$[\mathrm{cm}]$
v_0	$c \approx 2.998 \times 10^{10}$	$[\mathrm{cm/s}]$
$t_0^\dagger$	1.029×10^{11}	$[\mathrm{sec}]$
$p_0^\dagger$	1.503×10^{-6}	$[\mathrm{dyne/cm^2}]$
$B_0^\dagger$	4.346×10^{-3}	$[\mathrm{Gauss}]$

Table 4.1 Code units for the jet simulation. † Value is derived from other units.

4.2 Numerical model

4.2.1 Foundational equations

Simulations are run using the astrophysical fluid dynamics code *PLUTO v4.4-patch2*[1] (Mignone et al., 2007) with an additional turbulence module (`FORCED_TURB`, see Section 4.2.3 for details) on a regular Cartesian grid with axes (x, y, z). Our setup utilises 2nd-order Runge Kutta (RK2) time-stepping and a Harten–Lax–van Leer contact (HLLC) Riemann solver with linear reconstruction using the van Leer harmonic mean limiter. For this combination of reconstruction and time-stepping, the CFL number is stable up to $\lesssim 1/N_{\mathrm{dim}} \approx 0.3$. As the stability of the simulation can vary in time, we vary the CFL number from 0.05 to 0.3 depending on need. To further improve the robustness of the code, we enable the multi-dimensional `SHOCK_FLATTENING` feature, which reverts the limiter from van Leer to the more diffusive `MINMOD` in the presence of strong shocks. The code is parallelised using the message passing interface (MPI) library over 228 Cascade Lake cores on the Cambridge Service for Data Driven Discovery (CSD3) system.

For our given setup, *PLUTO* solves the ideal relativistic MHD (RMHD) equations. In what follows, ρ, p, and E are the density, thermal pressure, and energy density, respectively. Vectors $\boldsymbol{m}$, $\boldsymbol{v}$, and $\boldsymbol{B}$ represent the momentum density, velocity, and magnetic field. As we are working within the relativistic regime, we must also define the Lorentz factor $\Gamma = (1 - v^2/c^2)^{-1/2}$ and laboratory frame density $D = \Gamma\rho$. For simplicity, all velocities are normalised to the speed of light c. The complete list of code units (defined and derived) is given in Table 4.1. Unless otherwise specified, all values are shown in code units.

[1] http://plutocode.ph.unito.it

PLUTO solves the following system of conservation laws on conservative variables $\boldsymbol{U} = (D, \boldsymbol{m}, \boldsymbol{B}, E)$ (Mignone et al., 2007),

$$\frac{\partial}{\partial t}\begin{bmatrix} D \\ \boldsymbol{m} \\ \boldsymbol{B} \\ E \end{bmatrix} + \nabla \cdot \begin{bmatrix} D\boldsymbol{v} \\ w_t \Gamma^2 \boldsymbol{v}\boldsymbol{v} - \boldsymbol{b}\boldsymbol{b} + \boldsymbol{I} p_t \\ \boldsymbol{v}\boldsymbol{B} - \boldsymbol{B}\boldsymbol{v} \\ \boldsymbol{m} \end{bmatrix}^T = \begin{bmatrix} 0 \\ \boldsymbol{f}_g \\ 0 \\ \boldsymbol{v} \cdot \boldsymbol{f}_g \end{bmatrix}, \tag{4.1}$$

with $w_t = \rho h + b_m^2$ for a specific enthalpy $h \equiv h(P, \rho)$ that is dependent on the equation of state (EOS) and $b_m^2 = |\boldsymbol{B}|^2/\Gamma^2 + (\boldsymbol{v} \cdot \boldsymbol{B})^2$, and additionally where $\boldsymbol{b} = \boldsymbol{B}/\Gamma + \Gamma(\boldsymbol{v} \cdot \boldsymbol{B})\boldsymbol{v}$, $\boldsymbol{I}$ is the identity matrix, $p_t = p + b_m^2/2$ is the total pressure, and $\boldsymbol{f}_g = \rho\Gamma^2 \left[\Gamma^2 \boldsymbol{v}(\boldsymbol{v} \cdot \boldsymbol{a}) + \boldsymbol{a}\right]$ is the acceleration term associated with acceleration vector $\boldsymbol{a}$. $\boldsymbol{a}$ represents the acceleration of the frame of reference, which, in the non-relativistic limit, reduces to $\boldsymbol{a} = -\nabla\boldsymbol{\Phi} + \boldsymbol{g}$.

In this work, it will often be simpler to reference instead the primitive variables $\boldsymbol{V} = (\rho, \boldsymbol{v}, p, \boldsymbol{B})$, which can be calculated from conservative via (Mignone et al., 2007)

$$D = \rho\Gamma, \tag{4.2}$$

$$\boldsymbol{m} = (\rho h \Gamma^2 + |\boldsymbol{B}|^2)\boldsymbol{v} - (\boldsymbol{v} \cdot \boldsymbol{B})\boldsymbol{B}, \tag{4.3}$$

$$E = \rho h \Gamma^2 - p + \frac{|\boldsymbol{B}|^2}{2} + \frac{|\boldsymbol{v}|^2|\boldsymbol{B}|^2 - (\boldsymbol{v} \cdot \boldsymbol{B})^2}{2}. \tag{4.4}$$

We note that there is an asymmetry in the difficulty of converting between primitive and conserved variables. For a given set of primitive variables, the conserved variables are trivially calculated via Equations 4.2–4.4. To proceed in the other direction, from conserved variables to primitive variables, requires these equations to be inverted and requires a numerical solution. This is the principal reason why relativistic fluid dynamics is more computationally expensive compared with Newtonian fluid dynamics.

As our simulation domain contains both non-relativistic ambient gas and highly relativistic jet plasma, we use the Taub-Mathews (TM) EOS for a relativistic perfect gas, initialised with a ratio of specific heats $\gamma = 5/3$. The *PLUTO* code employs a quadratic approximation to the Synge gas (Synge, 1957) expression that gives $\gamma \to 5/3$ in the limit of low temperature plasma, and $\gamma \to 4/3$ for high temperature (see Mignone & McKinney, 2007; Mignone et al., 2005, for details),

$$\left(h - \frac{p}{\rho}\right)\left(h - 4\frac{p}{\rho}\right) = 1 \tag{4.5}$$

where the specific enthalpy h is related to the internal energy e and pressure p as

$$h = 1 + e + \frac{p}{\rho}. \tag{4.6}$$

The $\nabla \cdot \boldsymbol{B} = 0$ condition is maintained through Powell's Eight Wave method (Powell, 1994; Powell et al., 1999), which introduces additional divergence source terms to the MHD equations to destroy monopoles to truncation error.

4.2.2 Mesh and boundary construction

Our simulation domain is rectangular cuboid extended over $x \in [-100, 200]$, $y \in [-100, 100]$ and $z \in [-100, 100]$ with a uniform resolution of 0.39 code units per cell.

We employ periodic boundary conditions in the x direction and outflowing in the y- and z directions. Utilising a periodic x allows us to extend the simulation time by sourcing new turbulent ICM material from the outgoing end of the domain, however, the total simulation runtime is still limited by growth in the x- and y-directions.

4.2.3 Initial conditions

Turbulent background

The ICM is composed of magnetised plasma and is subjected to strong perturbative events in the form of AGN feedback and galaxy mergers. As such, the ICM is expected to be turbulent, with the existence of such turbulence important for heating the ICM and offsetting the cluster cooling problem (Zhuravleva et al. 2014, although see also Bambic et al. 2018b). X-ray observations and measurements of the Sunyaev-Zeldovich (SZ) effect in the central region of clusters show turbulent density and pressure fluctuations (Khatri & Gaspari, 2016; Zhuravleva et al., 2014), however, the precise details of this turbulence are hard to pin down, with current observational constraints subject to low spectral resolution, statistical uncertainties and sampling variance (Cucchetti et al., 2019).

Faraday Rotation measurements of polarised sources lying behind galaxy clusters exhibit a 'patchiness' on scales of 5–10 kpc and up to $\sim$50 kpc, suggesting tangled magnetic fields with coherence lengths of the same order (Bonafede et al., 2010; Carilli & Taylor, 2002; Murgia et al., 2004; Vogt & Enßlin, 2005). Further, large-scale RM trends in, for example, Hydra A points to a large-scale ordered field component on the scale of $\sim$100 kpc (Carilli & Taylor, 2002). Hence it is apparent that the ICM magnetic field must contain components at both small ($\sim$5 kpc) and large ($\sim$100 kpc) scales.

For simulations where an ICM field is necessary, for example in generating simulated RM or polarisation maps, the ICM field is typically generated following the method of Tribble (1991) and Murgia et al. (2004) (e.g. Hardcastle & Krause, 2014; Huarte-Espinosa et al., 2011b; Jerrim et al., 2024; Rodman, P. E., 2018). In this method, the magnetic field vector potential $\boldsymbol{A}$ is first generated in Fourier space as $\tilde{\boldsymbol{A}}(\boldsymbol{k})$, with amplitudes A drawn from a Rayleigh distribution with scale parameter $|A_k|^2 \propto k^{-\zeta}$ for some slope ζ, and phases θ drawn from a random uniform distribution between 0 and 2π (Murgia et al., 2004)

$$P(A,\theta)dAd\theta = \frac{A}{2\pi|A_k|^2}\exp\left[\left(-\frac{A^2}{2|A_k|^2}\right)\right]dAd\theta. \tag{4.7}$$

The magnetic field in the Fourier domain can then be retrieved as

$$\tilde{\boldsymbol{B}}(\boldsymbol{k}) = i\boldsymbol{k}\times\tilde{\boldsymbol{A}}(\boldsymbol{k}) \tag{4.8}$$

and converted back to real space by the inverse Fast Fourier Transform (Press et al., 1992).

This method produces a divergence-free, multi-scale magnetic field whose components follow a power-spectrum $|B_k|^2 \propto k^{-\zeta+2} = k^n$. For $\zeta = -11/3$, a Kolmogorov-like spectrum with index $n = -5/3$ is recovered. The strength of the individual components of the magnetic field B_i follow a Gaussian distribution with $\langle B_i\rangle = 0$ and $\sigma^2_{B_i} = \langle B_i^2\rangle$, while the total magnitude B is Maxwellian with $\langle B\rangle = 2\sqrt{2/\pi}\sigma_{B_i}$ and $\sigma_B = \sqrt{(3\pi-8)/\pi}\sigma_{B_i}$ (Bonafede et al., 2010; Murgia et al., 2004).

The magnetic field generated by the Tribble (1991) method is Gaussian in each component and so, in a statistical sense, appears to match well to RM and inverse-Compton (IC) observations. As a result, this method of field generation has found much use in generating simulated RM maps. However, it falls far short for the purposes of the study of dynamical evolution in general and for any study of filamentary structures in particular. Real MHD turbulence is, in reality, poorly described by a Gaussian random field as the latter lacks intermittency and filament formation. In order to study the interaction between the jet magnetic field and possible ICM filaments, we must go beyond the approach of Tribble (1991) and others and instead turn to turbulence forcing.

Our jets are launched into a pre-evolved turbulent background that is separately generated via *PLUTO*'s `FORCED_TURB` module[2]. The underlying mesh is identical to Section 4.2.1. This portion of the study features only non-relativistic gas, relatively weak magnetic fields, and only weak shocks; hence, we use less diffusive, more restrictive computational methods

[2]With thanks to Bhargav Vaidya and Kazim Hussain for sharing the module and providing support during set-up.

here. The background setup again utilises 2nd-order Runge Kutta (RK2) time-stepping, this time with the Harten-Lax-van Leer discontinuities (HLLD) Riemann solver and linear reconstruction and a constant CFL number 0.3. We use an ideal EOS and maintain the $\nabla \cdot \boldsymbol{B} = 0$ condition to machine precision with Constrained Transport (CT).

The initial density profile is uniform and defined as $\rho_a = 1$ in code units, with a unit density of $\rho_0 = 1 m_p \ \mathrm{cm}^{-3}$ where m_p is the proton mass. The initial pressure is also uniform and recovered from the ideal gas law $p_a = n_a k_B T_a / \rho_0 v_0^2$ for a number density $n_a = 1 \ \mathrm{cm}^{-3}$, Boltzmann constant k_B, and gas temperature $T_a = 10^8$ K, and normalised to code units via the unit density ρ_0 and unit velocity $v_0 = c$ (Table 4.1).

The initial background magnetic field is purely vertical and is defined through the magnetic vector potential $\boldsymbol{A}$ in cylindrical coordinates (R, θ, Z),

$$A_R = 0 \tag{4.9}$$

$$A_\theta = 0.5 R B_{0,a} \tag{4.10}$$

$$A_Z = 0 \tag{4.11}$$

where B_0 is the initial maximum field strength set by the magnetic plasma beta, $B_{0,a} = \sqrt{2P_a/\beta_a}$ for an initial $\beta_a = 100$. When defined in this way, the magnetic field is divergence-free everywhere and, using Constrained Transport, is maintained divergence-free to machine precision at all times.

This uniform setup then develops turbulence under forcing. In the `FORCED_TURB` module, a time- and space-varying driving force $\boldsymbol{F}^{\mathrm{turb}}(\boldsymbol{x}, t)$ is added as a source term to the equations for momentum and energy, i.e.

$$\frac{\partial \boldsymbol{m}}{\partial t} + \nabla \cdot \left[w_t \Gamma^2 \boldsymbol{v}\boldsymbol{v} - \boldsymbol{b}\boldsymbol{b} + \boldsymbol{I} p_t \right]^T = \rho \boldsymbol{F}^{\mathrm{turb}} \tag{4.12}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot [\boldsymbol{m}]^T = \boldsymbol{m} \cdot \rho \boldsymbol{F}^{\mathrm{turb}}. \tag{4.13}$$

The forcing term is decomposed into Fourier components in order to control the spatial dependence of the forcing. The magnitude of each component is then evolved according to the Ornstein-Uhlenbeck (OU) process (Uhlenbeck & Ornstein, 1930), a well-defined stochastic process featuring mean-reverting and stationary effects. The OU process is a stochastic differential equation (SDE) of the form (Karatzas, 1988)

$$dX(t) = \kappa(\theta - X(t))dt + \sigma d\mathcal{W}(t), \tag{4.14}$$

where $\kappa > 0$ is the rate at which the process reverts to its mean, θ is the long-term mean, and $\sigma > 0$ is the volatility associated with the random fluctuations modelled by the Weiner process $\mathcal{W}(t)$. In the `FORCED_TURB` module, the SDE describing the evolution of the force $\boldsymbol{F}^{\rm turb}$ in Fourier space takes the form

$$d\boldsymbol{F}^{\rm turb}(\boldsymbol{k},t) = \boldsymbol{F}_0^{\rm turb}\mathcal{P}^{\zeta}(\boldsymbol{k})d\mathcal{W}(t) - \boldsymbol{F}^{\rm turb}(\boldsymbol{k},t)\frac{dt}{T}. \tag{4.15}$$

where $\mathcal{P}^{\zeta}(\boldsymbol{k})$ is the projection tensor, given in Fourier space as

$$\mathcal{P}_{ij}^{\zeta}(\boldsymbol{k}) = \zeta\mathcal{P}_{ij}^{\perp}(\boldsymbol{k}) + (1-\zeta)\mathcal{P}_{ij}^{\parallel}(\boldsymbol{k}) \tag{4.16}$$

where $\mathcal{P}_{ij}^{\perp}(\boldsymbol{k}) = \delta_{ij} - k_i k_j/k^2$ is the solenoidal projection operator and $\mathcal{P}_{ij}^{\parallel}(\boldsymbol{k}) = k_i k_j/k^2$ is the compressive projection operator, with Kronecker delta function δ_{ij}. Hence, the value of the parameter $0 \leq \zeta \leq 1$ defines the nature of the forcing, from compressive ($\zeta = 0$) to solenoidal ($\zeta = 1$).

The first term on the RHS of Equation (4.15) models diffusion of the magnitude of the Fourier component of the forcing term via the random Weiner process $\mathcal{W}(t)$, while the second term models the drift (decay) in the correlation of $\boldsymbol{F}^{\rm turb}$ with itself, with the autocorrelation timescale given by T. This method of forcing models well the process of energy injection at large scale, breaking up into small-scale turbulence (Federrath et al., 2010).

Within the `FORCED_TURB` module we enable amplitude adjustment with a desired velocity of 0.001 (approximately 30 per cent of the initial sound speed in the ambient medium) and a solenoidal weight $\zeta = 0.5$ corresponding to a mix of solenoidal and compressive turbulence. Solenoidal and compressive turbulent motions naturally generate each other over time, and so a mix of solenoidal and compressive components would be expected in the ICM even when turbulence is assumed to be subsonic (Porter et al., 2015), although the exact weighting of solenoidal/compressive modes is not currently known for ICM turbulence. We caution that the choice of ζ has a significant effect on the nature of filaments that are formed, with more compressive forcing leading to more prominent filaments.

Turbulence is driven until the simulation reaches an approximate steady state, characterised by a plateauing in both $E_{\rm mag}$ and $E_{\rm kin}$. Once this state is reached, we add a sideways velocity $v_w = fc_{s,a} = f\sqrt{\gamma p_a/\rho_a}$ where $c_{s,a}$ is the sound speed in the ambient medium, defined by the adiabatic index $\gamma = 5/3$ and the ambient pressure p_a and density ρ_a, and f is some multiplying factor.

Brightest Cluster Galaxies (BCGs)—common hosts for WAT jets—have been found with relative velocities on the order of 20 to 40 per cent the ICM sound speed ($\sim$ 200–400 km s^{-1}

O'Dea & Baum, 2023). However, they are also commonly associated with merging clusters, for which simulations suggest transonic wind velocities of up to $\sim$ 1000–2000 km s^{-1} (Roettiger et al., 1993, 1998).

In determining the ideal v_w setting, we must also consider the pre-existing turbulent motions. Our turbulence is forced with a desired RMS velocity of 30 per cent of $c_{s,a}$ and results in regions with velocities up to 50 per cent of the local sound speed. We therefore set $v_w = 0.3c_s$ for simulation `v0.24`. As the speed of jet `v0.94` is much more significant, transonic relative motions are too slow to have any noticeable effect while supersonic relative motions could introduce unrealistic effects. Instead, we opt in simulation `v0.94` to artificially increase the ambient pressure by some factor k such that the resulting sound speed is $c_{s,a} = \sqrt{k\gamma p_a/\rho_a}$. In this way, we can achieve a significantly higher relative motion while still keeping the motion transonic. The crosswind fraction f and pressure enhancement factor k are given Table 4.3 for each of the jet simulations.

This is the final state then used as the initial background for the jet launching.

Final snapshots of the turbulent background (prior to the addition of the crosswind) are shown in Figure 4.2 for variables ρ, $|v|/c_s$, $|B|$, and the magnetic plasma β. The resulting flow is clearly turbulent with large-scale filamentary structures caused by the mixed solenoidal/compressive forcing. The density remains in the neighbourhood of its initial value (unity) over most of the domain, with regions of lower density corresponding to higher $|B|$ and to lower β.

The evolution of the kinetic and magnetic energies, along with the volume-averaged magnetic plasma β is shown in Figure 4.3. The simulation displays a plateauing of $E_{\rm mag}$, $E_{\rm kin}$, and β for $t > 2.5 \times 10^5$. The magnetic field strength increases as a result of dynamo action in the turbulent fluid, with the plasma β decreasing from its initial value of 100 to $\sim$25. The power spectrum of the kinetic and magnetic energy is shown in Figure 4.4. We find an inertial range slope that is similar to that of Kolmogorov turbulence ($k^{-5/3}$). Slight deviations from the Kolmogorov spectrum are not unexpected, as the theory of Kolmogorov (1941) assumes turbulence is isotropic, homogeneous, and hydrodynamic, all of which will be violated for MHD turbulence.

We also show the distribution of one component of the magnetic field (B_y) in Figure 4.4. At first glance, the distribution appears Gaussian, with a mean and standard deviation consistent with expectations for a Gaussian distribution, namely,

$$\langle B_i \rangle = 0, \tag{4.17}$$

$$\sigma^2_{B_i} = \langle B_i^2 \rangle. \tag{4.18}$$

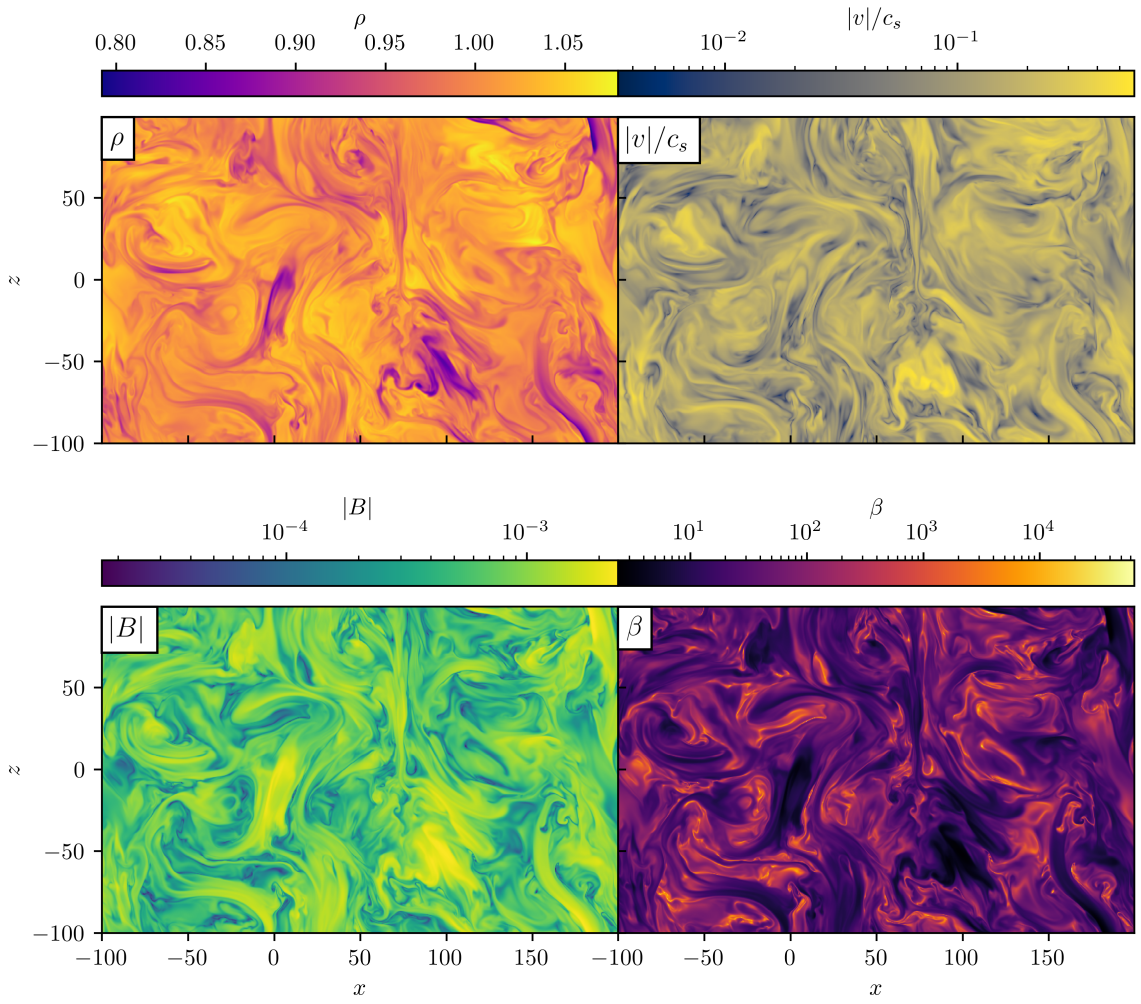


Fig. 4.2 Final snapshots of one realisation of the turbulent background: (*top left*) The gas density ρ, (*top right*) the velocity magnitude $|v|$ in units of the local sound speed c_s, (*bottom left*) the magnetic field strength $|B|$, and (*bottom right*) the magnetic plasma β. The turbulent fluid features both small and large scale structure.

However, we can observe non-Gaussianity in the distribution if we look at higher-order moments, in particular, the skewness $\mathcal{S}$ and kurtosis $\mathcal{K}$ of the variable (Federrath et al., 2010),

$$\mathcal{S}_{B_i} = \frac{\langle (B_i - \langle B_i \rangle)^3 \rangle}{\sigma_{B_i}^3}, \tag{4.19}$$

$$\mathcal{K}_{B_i} = \frac{\langle (B_i - \langle B_i \rangle)^4 \rangle}{\sigma_{B_i}^4}. \tag{4.20}$$

For a Gaussian distribution, we would expect a skewness $\mathcal{S} = 0$ and kurtosis $\mathcal{K} = 3$, instead we find $\mathcal{S}_{B_y} = -0.11$ and $\mathcal{K}_{B_y} = 3.56$. Such deviations are considered a sign of

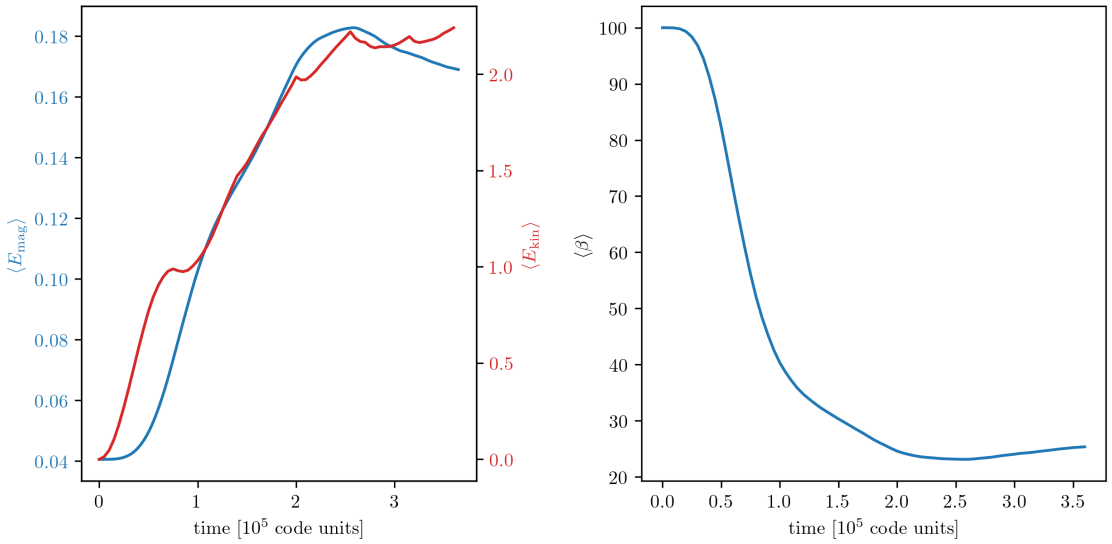


Fig. 4.3 Evolution of the turbulent background simulation in time. (*Left*) The evolution of magnetic energy $E_{\rm mag}$ and kinetic energy $E_{\rm kin}$ as a function of simulation time. (*Right*) The evolution of the volume-averaged magnetic plasma β, $\langle \beta \rangle$, as a function of simulation time.

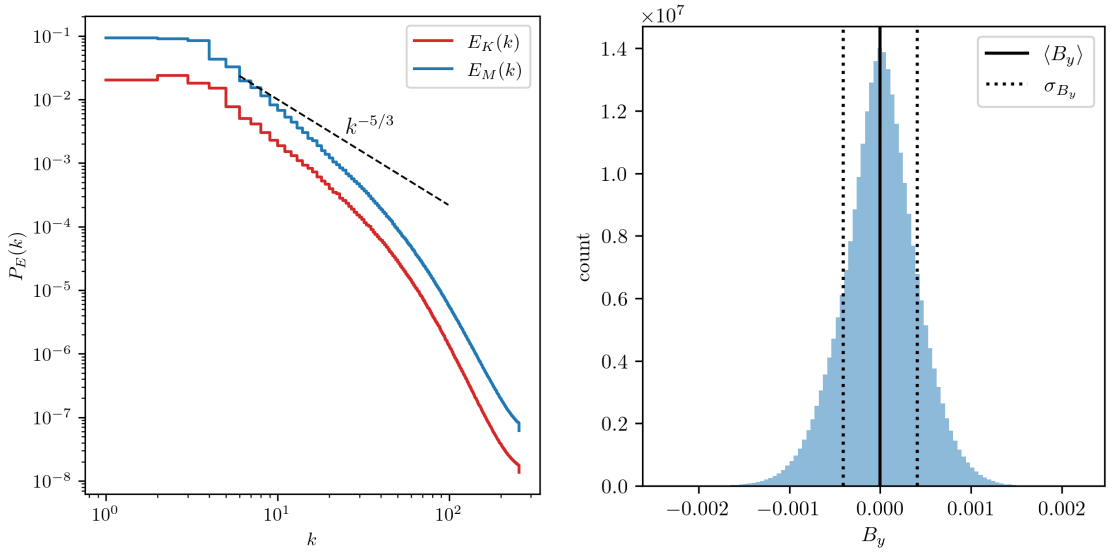


Fig. 4.4 (*Left*) The power spectrum $P(k)$ as a function of wavenumber k, for the kinetic energy E_K (red) and magnetic energy E_M (blue). The dashed line indicates the Kolmogorov slope $k^{-5/3}$. (*Right*) The distribution of B_y values, with the mean (solid black line) and standard deviation (dashed black lines) indicated. At first glance, the distribution appears Gaussian, but it features significant degrees of skewness and kurtosis.

intermittency (Federrath et al., 2010), which in the context of fluid turbulence generally refers to the tendency of a probability distribution to develop long tails with increasingly

extreme events for higher Reynolds numbers. Amplification of the turbulent field occurs primarily through stretching via solenoidal motions, which occurs most effectively if the scale of solenoidal motions l_s is much smaller than that of the largest turbulent eddies L_d, i.e. the hydrodynamical turbulent Reynolds number $\mathrm{Re}_H \sim L_d/l_s$ is large (Porter et al., 2015). The hot ICM is thought to have $\mathrm{Re}_H \gg 100$ (Porter et al., 2015), and so deviations away from perfect Gaussianity are expected.

Jet launching

The jet is launched via a cylindrical internal boundary zone centred on (0,0,0), with length $l_j = 3$ units above and below the origin and radius $r_j = 3$ units. Within this region, we set the primitive hydrodynamic variables in code units as

$$\rho_j = 0.001 \tag{4.21}$$

$$v_{j,x} = v_{j,y} = \begin{cases} 0.001, & \text{if } |z| = l_j \\ 0, & \text{otherwise} \end{cases} \tag{4.22}$$

$$v_{j,z} = (\mathrm{sign}\, z)\sqrt{1 - \frac{1}{\Gamma^2}} \tag{4.23}$$

$$p_j = p_a. \tag{4.24}$$

The small outward velocity of $v_{x,j}$ and $v_{y,j}$ is added at the injection zone caps to prevent the jet from forming a 'pencil beam' shape and to allow for physical recollimation of the jet.

In addition, we include a fluid tracer Tr which takes the value of 1 within the jet injection region and zero elsewhere. The tracer takes the form of a passive scalar that is advected according to (in conservative form)

$$\frac{\partial(\rho \mathrm{Tr})}{\partial t} + \nabla \cdot (\rho \mathrm{Tr} \boldsymbol{v}) = 0. \tag{4.25}$$

The magnetic field of the jet is initially defined in cylindrical coordinates (R, θ, Z) through the magnetic vector potential $\boldsymbol{A}$ following the 'force-balance setup' of Gaibler

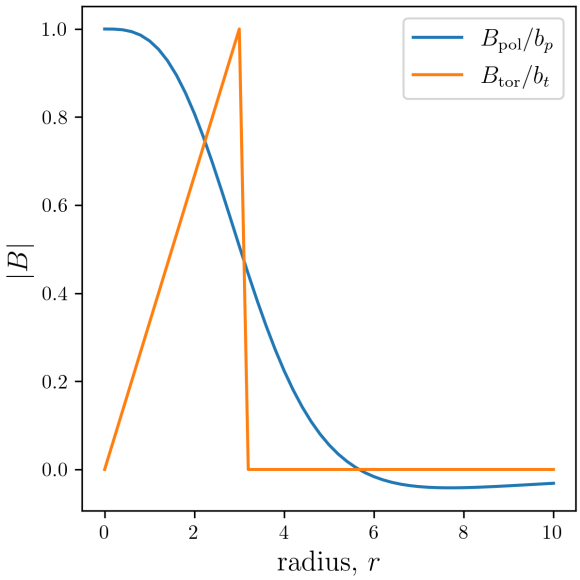


Fig. 4.5 Profile of the added jet magnetic field in the poloidal (blue line) and toroidal (orange line) components, for $y = 0$ and $z = 0$. Values are normalised to the maximum strength in the respective component.

(2008)

$$A_R = \begin{cases} b_t \left(\frac{R}{r_j}\right) Z, & \text{if } R \leq r_j \\ b_t \left(2 - \frac{R}{r_j}\right) Z, & \text{otherwise} \end{cases} \tag{4.26}$$

$$A_\theta = \frac{1}{2} b_p R \frac{1}{1+\sigma^3} \tag{4.27}$$

$$A_Z = 0 \tag{4.28}$$

where $\sigma = \sqrt{(R/a)^2 + (Z/a)^2}$, $a = 3r_j/2$ is a magnetisation radius, and b_p (b_t) is the maximum magnetic field value in the poloidal (toroidal) direction, determined as a fraction η ($1 - \eta$) of the total field strength $B_{\rm mag} = \sqrt{2p_j/\beta_j}$ for jet plasma beta $\beta_j = 10$.

This field is then converted to Cartesian coordinates through

$$A_x = (A_R \cos(\theta) - A_\theta \sin(\theta)) \tag{4.29}$$

$$A_y = (A_R \sin(\theta) + A_\theta \cos(\theta)) \tag{4.30}$$

$$A_z = A_Z \tag{4.31}$$

and the subsequent $\boldsymbol{B} = \nabla \times \boldsymbol{A}$ (Figure 4.5) is set within the nozzle region and added to the background everywhere. When defined this way, the field is continuous and divergence-free.

Parameter	Value [code units]
ρ_a	1
p_a	9.184×10^{-6}
β_a	100
ρ_j	$0.01 \rightarrow 0.001$
l_j	3
r_j	3
β_j	10

Table 4.2 Initial parameter values for the jet and turbulent background.

As we launch our jets from an internal boundary zone, the above fluid and magnetic variables are fixed within the injection region, while the rest of the simulation domain evolves as normal.

With this jet launching system, we perform two main simulations:

1. A highly relativistic jet ($v_j = 0.94c$) launched into an elevated pressure background and with a resulting crosswind speed $v_w = c_{s,a}$, and

2. A mildly relativistic jet ($v_j = 0.24c$) launched into an unenhanced pressure background with a crosswind speed $v_w = 0.3c_{s,a}$.

These simulations are termed `v0.94` and `v0.24` respectively. The parameters which uniquely define these simulations are given in Table 4.3. We additionally repeat `v0.24` without any crosswind (`v0.24w0`) and with a second turbulent background initialisation (`v0.24r2`) for comparison.

The one-sided kinetic RMHD jet power Q_j can be calculated in SI units as (English et al., 2016; Stimpson et al., 2023; Yates-Jones et al., 2021)

$$Q_j = \pi r_j^2 v_j \left[\Gamma(\Gamma - 1)\rho_j c^2 + \frac{\gamma}{\gamma - 1}\Gamma^2 p_j + \Gamma^2 \frac{B_{\text{mag}}^2}{1\mu_0} \right] \tag{4.32}$$

where μ_0 is the permittivity of free space and the rest take values derived in Tables 4.3 and 4.2, converted from code units using the conversion factors of Table 4.1. This calculation yields one-sided kinetic jet powers of $Q_j = 7.105 \times 10^{39}$ W for simulation `v0.94` and $Q_j = 1.680 \times 10^{37}$ W for simulation `v0.24`. As such, we expect simulation `v0.94` to produce a high-powered FR-II type source, while simulation `v0.24` will form a low-power FR-I.

Simulation ID	Γ_j	$c_{s,j}$	$\mathcal{M}_{\mathrm{int},j}$	$\mathcal{M}_{\mathrm{ext},j}$	k	f	Γ_a	$c_{s,a}$	$\mathcal{M}_w$
`v0.94`	3	0.11	24.75	69.99	100	1	1	0.04	1
`v0.24`	1.03	0.12	1.98	62.71	1	0.3	1	0.004	0.3

Table 4.3 Initial parameter values for the `v0.94` and `v0.24` jet simulations, from left: the simulation name, jet Lorentz factor Γ_j, sound speed in jet $c_{s,j}$, internal jet Mach number $\mathcal{M}_{\mathrm{int},j}$, external jet Mach number $\mathcal{M}_{\mathrm{ext},j}$, pressure enhancement factor k, crosswind sound speed fraction f, ambient Lorentz factor Γ_a, ambient sound speed $c_{s,a}$, and wind Mach number $\mathcal{M}_w$.

4.2.4 Jet termination

It may be interesting to track the evolution of lobes with the source turned off. Additionally, if the jet is highly relativistic, it may reach the end of the simulation domain before any significant bending occurs—to prevent this, we may need to terminate the jets.

A relativistic jet produces a significant bow shock as it expands into the ambient medium. As our jets are initially pressure-matched to the ambient background, they are under-pressured compared to the surrounding shocked medium. Hence, the injection region will implode if the jet velocity drops to zero too quickly. Implosions in such a delicate simulation lead to runaway numerical instabilities and eventual crashes, which must be avoided.

To account for this risk, we terminate our `v0.94` jet by linearly decreasing $v_{j,z}$ to zero over a period of 25 simulation time units. Once $v_{j,z} = 0$, the injection region is allowed to evolve with the rest of the simulation domain.

4.2.5 Visualisation

The main analysis of our simulation data will be performed on the basic physical variables (ρ, p, $\boldsymbol{v}$, and $\boldsymbol{B}$). In addition, our jets are injected with a jet tracer fluid with value $\mathrm{Tr} = 1$ within the injection region and zero elsewhere, which can be used to trace the evolution of jet-injected material.

In real astrophysical jets, synchrotron emission is generated by electrons orbiting about magnetic field lines. As it would be computationally prohibitive to simulate the movement of many trillions of electrons, we must instead look to other ways to approximate the emission profile of our jets. Although we do not have direct particle information, we can access information about the magnetic field and the jet tracer. With this in mind, we simulate

emissivity through the Stokes parameters (Hardcastle & Krause, 2014),

$$j_I = p(B_x^2 + B_z^2)^{0.5(\alpha-1)}(B_x^2 + B_z^2) \tag{4.33}$$

$$j_Q = \mu p(B_x^2 + B_z^2)^{0.5(\alpha-1)}(B_x^2 - B_z^2) \tag{4.34}$$

$$j_U = \mu p(B_x^2 + B_z^2)^{0.5(\alpha-1)}(2B_x B_z) \tag{4.35}$$

for pressure p, spectral index α, and maximum fractional polarisation $\mu = (\alpha+1)/(\alpha+5/3)$. As in Hardcastle & Krause (2014), we take $\alpha = 0.5$ for $\mu = 0.69$.

We take all cells with a tracer value Tr $> 10^{-3}$, multiply them by the calculated Stokes values, and integrate them through the viewing direction (here, the y-axis), producing a 2D image at the spatial resolution of the simulation.

We may additionally 'observe' our sources using the Rotation Measure, calculated as in Equation (1.82), and for code units as,

$$\text{RM} = \int n\boldsymbol{B} \cdot d\boldsymbol{l}, \tag{4.36}$$

for number density n and magnetic field $\boldsymbol{B}$ projected along the line of sight $\boldsymbol{l}$.

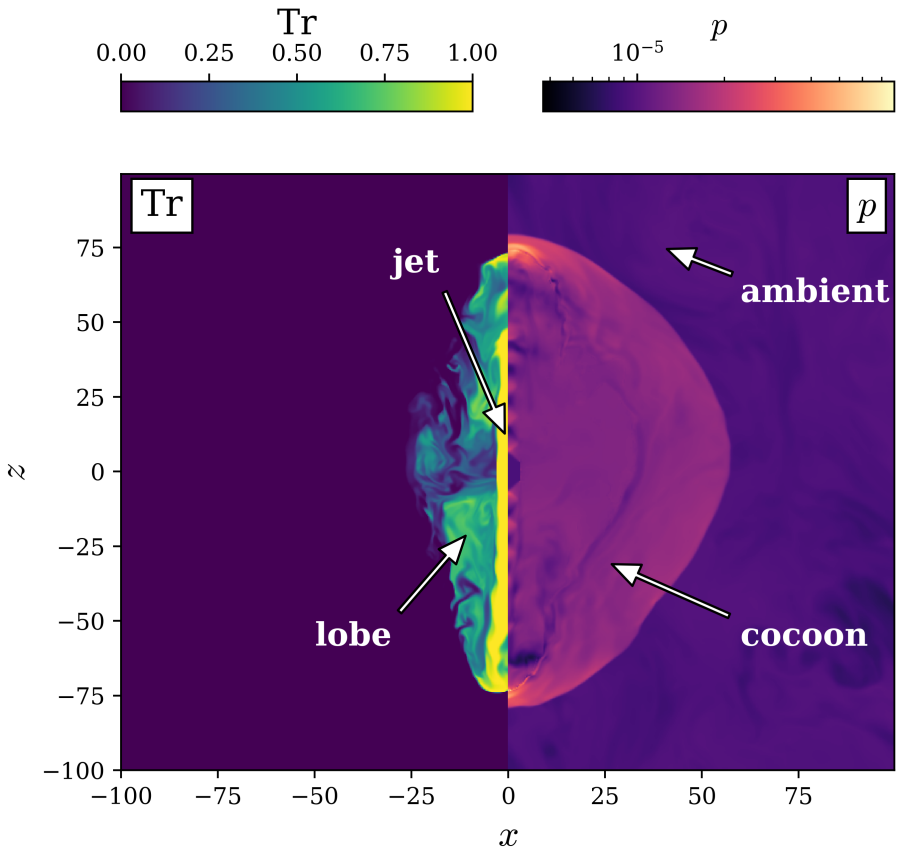


Fig. 4.6 A midplane slice of simulation `v0.24w0` at time $t = 9,200$ with components annotated. The left half of the image shows the tracer fluid Tr, the right half the pressure p.

4.3 Results

In our analysis of the data, it will be helpful to define a set of terms which describe different regions within both the wider simulation domain and the source itself. In the literature, the following terms often have slightly different definitions, and some may be used interchangeably. We define the *jet* as material of high velocity/Lorentz factor, generally with associated tracer value Tr ~ 1 that originates from the *jet injection nozzle*. This jet material slows as it impacts on the ambient medium, backflowing to fill the *lobe*, which we typically define as the region with tracer Tr $> 10^{-3}$. The expanding jet also produces a shockwave in the ambient medium which expands with the jet—we term this the (shocked) *cocoon*, and define its outer boundary by a pressure gradient exceeding three standard deviations. An annotated example is shown in Figure 4.6.

4.3.1 Jet dynamics and evolution

We show the time evolution of the four simulations through midplane slices in density in Figure 4.7. We also show various flow variables for each simulation at approximately equal jet length (~ 30) in Figure 4.8.

As the jets impact onto and move through the ambient medium, they inflate a lobe of low-density material enveloped by a cocoon of shocked ambient medium. The lobe backfills towards the jet source, and as this material is compressed in the direction of the wind, it extends the lobe and cocoon to the right. The relative movement of the low-density jet fluid against the ambient fluid leads to the development of Kelvin-Helmholtz (KH) instabilities. Characteristic KH vortices are most clearly seen in the first image of simulation `v0.24` in Figure 4.7, to the immediate left of the jet injection region. The presence of the KH instability along the boundary of the jet fluid seeds complex, turbulent structures inside the lobe, which can be seen most clearly at late times.

The difference in speed between the two main simulations leads to different apparent jet morphology, with `v0.94` taking on the early structure of a Fanaroff-Riley (FR) type II source, and `v0.24` (and its variants) taking on the structure of an FR type I source. However, we caution that without a simulated particle emission model, we cannot identify the location of the brightest point within each jet, and so a formal calculation of the Fanaroff-Riley morphological class is not possible.

We see that `v0.24` and `v0.24r2` are significantly affected by the crosswind, with apparent jet bending from $t \sim 5000$ onwards. For this simulation, we have an internal jet Mach number $\mathcal{M}_{\mathrm{int,j}} = 1.98$ and crosswind velocity $v_{x,w} = 0.3c_{s,a}$, which, per Equation (1.79), should yield an approximate bending length of $l_b \sim 130$. However, we instead observe a bending length on the order of $l_b \sim 75$–100. This disparity is because our turbulent background already features motion in the form of turbulent eddies driven at a velocity $\sim 0.3c_{s,a}$. This pre-existing velocity field combined with the added crosswind gives an actual crosswind velocity that is transonic ($\sim c_{s,a}$), similar to the work of Balsara & Norman (1992). The crosswind velocity is effectively halved inside the shocked cocoon of `v0.24` and reduced by a factor of ~ 5 in that of `v0.94`, hence leading to estimated bending lengths of $l_b \sim 75$ for simulation `v0.24` and $l_b \sim 45,000$ for simulation `v0.94`. We cannot verify the estimate for `v0.94` in the given domain; however, the updated estimate for `v0.24` matches the observations. As the `v0.94` jet cannot bend over the range studied, it is terminated at $t = 200$ and the remnant lobe is allowed to evolve passively, being swept back by the wind.

Figure 4.9 shows the jet length as a function of time. To determine the jet length, we locate the centroid of the jet spine (identified as cells with Lorentz factor $\Gamma > 0.95\Gamma_{j,\mathrm{init}}$) for

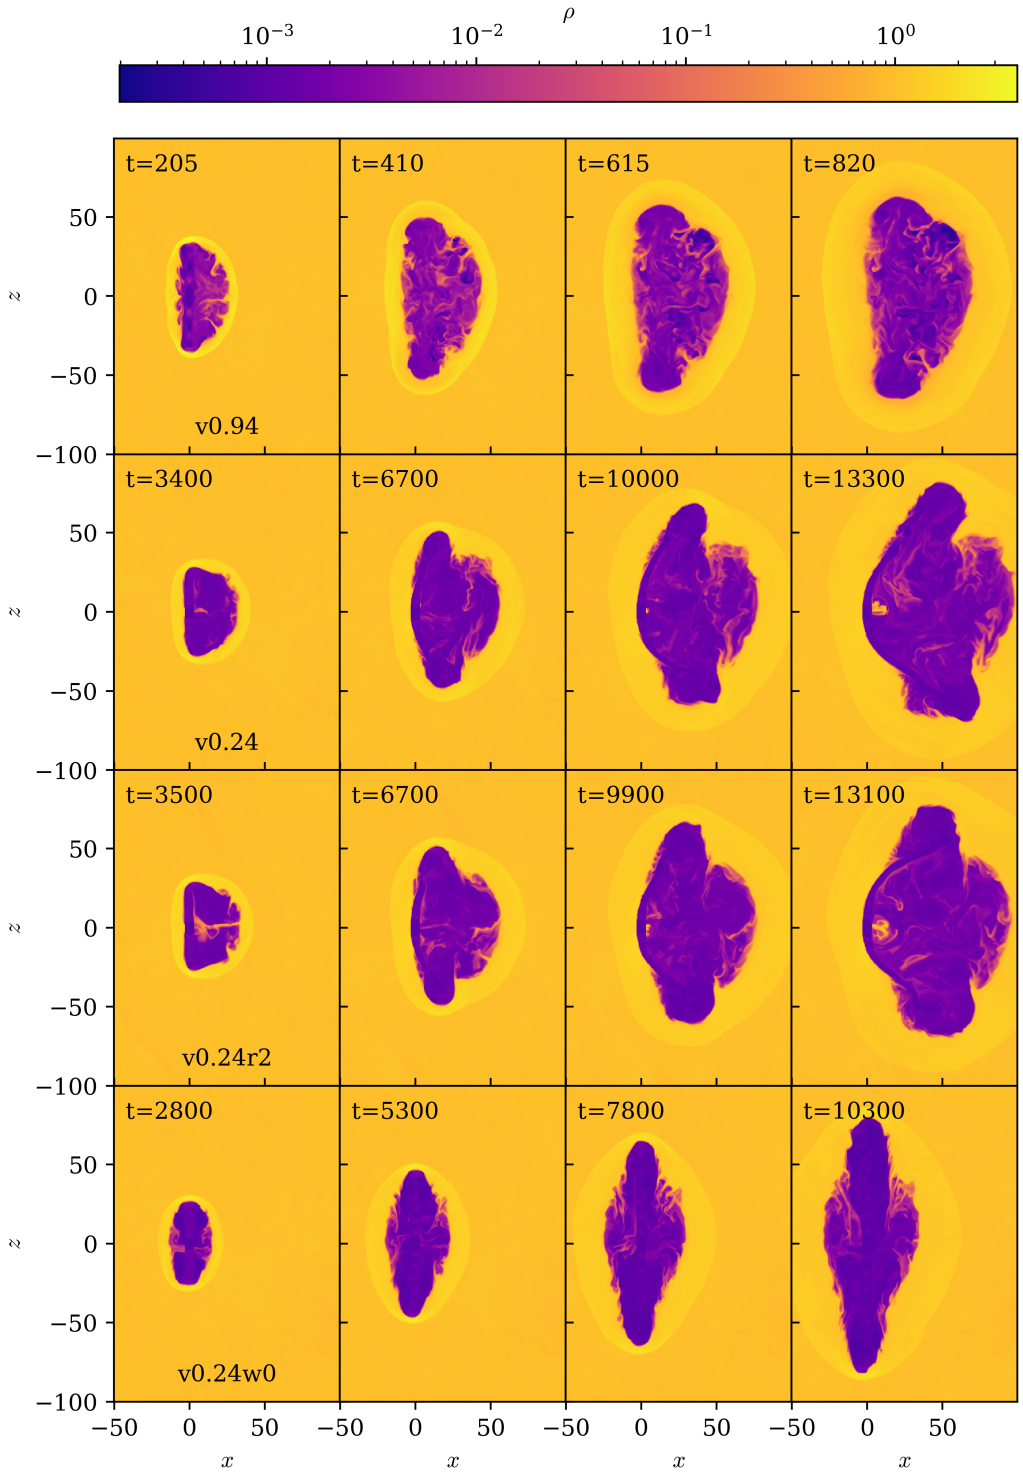


Fig. 4.7 Time evolution of the jet density for, from top to bottom, `v0.94`, `v0.24`, `v0.24r2`, and `v0.24w0`. Midplane slices in density are taken for the times indicated on each plot.

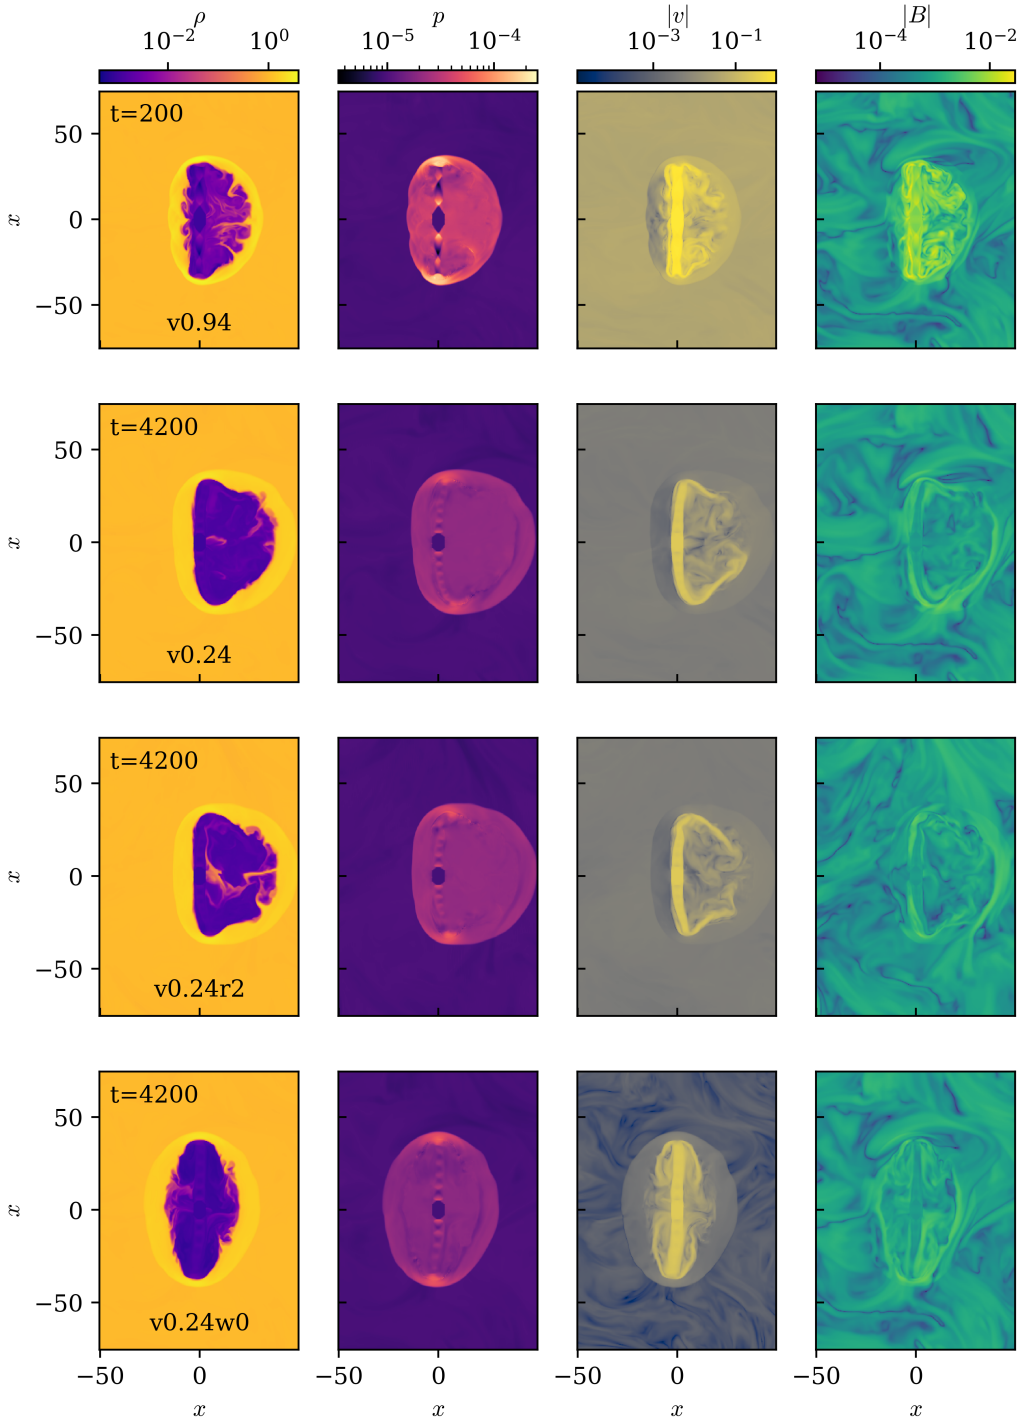


Fig. 4.8 Midplane slices of selected jet variables at the simulation end for, from top to bottom, `v0.94`, `v0.24`, `v0.24r2`, and `v0.24w0`. Quantities plotted, from left to right: density, pressure, jet (z) velocity, and magnetic field strength. Note: the pressure in `v0.94` has been rescaled to 1/100th to match that of `v0.24`.

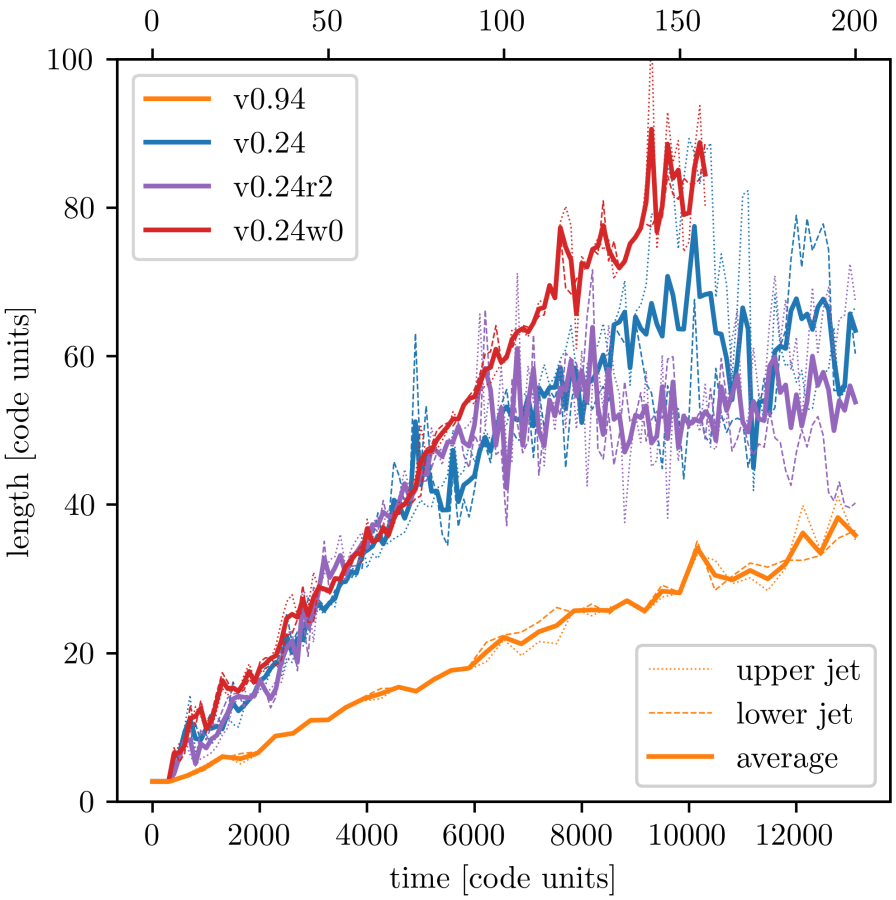


Fig. 4.9 Jet length over time for `v0.94` (orange), `v0.24` (blue), `v0.24r2` (purple), and `v0.24w0` (red). Values are shown for the upper jet (dashed lines), the lower jet (dotted lines), and the average of the two (thick solid lines). The time (horizontal) axis has been twinned to better show each simulation, with the lower axis corresponding to `v0.24` and its variants and the top axis corresponding to `v0.94`.

a given height z, and then sum the distances between centroids for each jet. This method can tolerate a slight deviation or bend in the jet but still requires the jet axis to be close to vertical.

Simulations `v0.24`/`v0.24r2` grow linearly and track the evolution of `v0.24w0` until $t \sim 5000$, when bending from the crosswind becomes apparent. Hence, values for $t > 5000$ for `v0.24`/`v0.24r2` in effect measure the bending length, not the total jet length. Simulation `v0.94` also grows linearly with time, albeit more rapidly (note the different time axes). As no jet bending is observed over the time and spatial domain covered here, there is no deviation from linear.

Simulations `v0.24` and `v0.24r2` display a characteristic 'flapping' near the jet ends, an effect seen most clearly in the jet tracer (Figure 4.10). This unstable behaviour continues

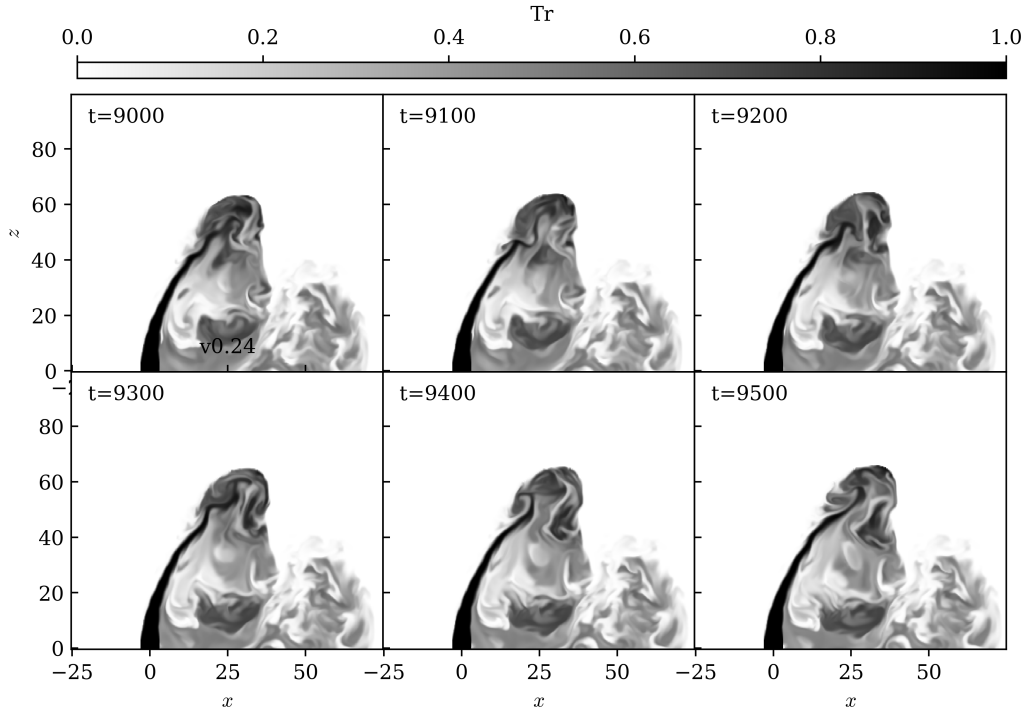


Fig. 4.10 Tracer fluid values for simulation `v0.24` from $t = 9,000$–$9,500$. The jet ends (upper jet shown here) display a chaotic 'flapping' behaviour, carving out bulbous lobes and sometimes becoming disconnected from the inflating lobe.

throughout each simulation and seeds further turbulent structures in the lobes. Similar movement has been previously seen in head-tail source simulations by O'Neill et al. (2019a) and Nolting et al. (2022), and may be the process by which a proto-FR-II is disrupted to form a plume-like FR-I. The jets of simulation `v0.24w0` also show some wandering (Figure 4.11). This behaviour is likely a feature of FR-I-type jets, as their much lower power and velocity make them more susceptible to destabilising effects like that from KH instabilities and regions of turbulent motion in the nearby ICM.

4.3.2 Magnetic field interaction

The evolution of the magnetic field strength over time for each simulation is shown in Figure 4.12. As the jets expand, they compress the ambient magnetic field against the lobe edge, leading to a 'sheath' of increased field strength enveloping the source. For simulations `v0.94`, `v0.24`, and `v0.24w0`, the initial turbulent background features a region of stronger magnetic field which hooks over the top of the upper jet and connects into the upper lobe. This connection is stable and persists over the lifetime of each simulation, providing a

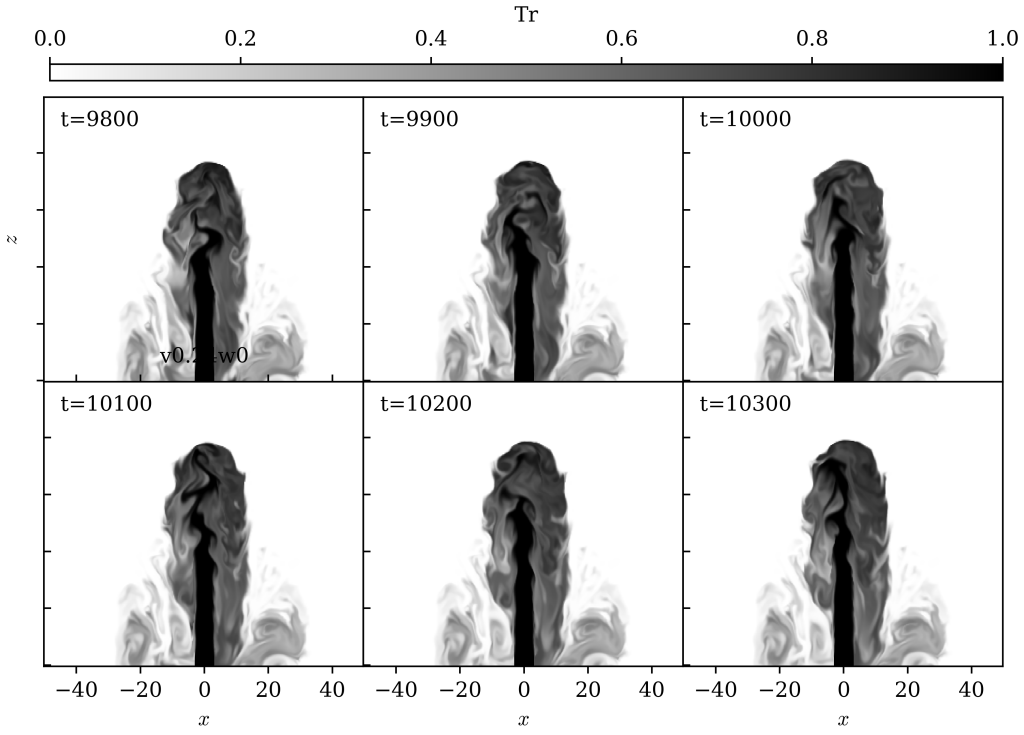


Fig. 4.11 As in Figure 4.10, but for the upper jet of simulation `v0.24w0` between $t = 9,800$ and $10,300$.

possible location for electron escape and (theoretical) filament formation. A similar feature is not visible in the background of `v0.24r2`, which has a different initial turbulent background realisation.

3D streamlines are shown in Figure 4.13 for `v0.24` at $t = 11,300$. For any charged particle looking to escape from the jet extremity along a potential fieldline, the most influential magnetic field is within the lobes. Hence, we plot a series of streamlines with starting points selected to land within each lobe (five per lobe) near the active end of the jet. We see an internal lobe field which is tanged and turbulent due to the disruption events caused by jet 'flapping', with the upper jet having been almost completely severed at $x \sim 30$ from one such event. A significant fraction of the streamlines successfully escape the lobes, with many extending back past the jet injection region. When looking side-on, we see several streamlines that may connect the opposing lobes; however, when viewed from behind, this is instead revealed to be a projection effect, with the streamlines separated from the opposing lobe by a distance of ∼20–50. A smaller number of streamlines extend a short distance into the space separating the jets before turning back—the intervening lobe material has been made turbulent by early KH instabilities along the jet edge, along with the frequent

discontinuities in the jet flows due to 'flapping', hence, there is no organised, large-scale field connecting the jet extremities in this region.

Filaments between the jet lobes and ambient medium will not be possible unless there is some degree of magnetic reconnection between the lobe and the shocked cocoon (i.e. the interface between the jet and ambient magnetic fields). In Figure 4.14 we plot regions of high magnetic current density $\boldsymbol{J} = |\boldsymbol{\nabla} \times \boldsymbol{B}|$ for simulation `v0.24` at $t = 11,300$, focussed on the extremity of the upper jet. We see elevated currents immediately adjacent to the jet (negative y), where many streamlines in Figure 4.13 exit the lobe. The elevated $\boldsymbol{J}$ values and the presence of escaping streamlines in this region strongly suggest that magnetic reconnection has occurred between the lobe magnetic field and that of the shocked ICM.

4.3.3 Synthetic observational signatures

If we assume filaments like those in source ESO 137-006 have a magnetic origin, then it is worthwhile to study our magnetic fields using the tools available to observers, namely polarisation and the Rotation Measure.

Polarisation

The Stokes parameters are shown for the final output of each simulation in Figure 4.15, as calculated in Section 4.2.5. These parameters correspond to the total polarised intensity (I), linearly polarised light at 90 degrees (Q) and 45 degrees (U), and total linear polarised intensity ($P = \sqrt{Q^2 + U^2}$), and have been weighted according to the jet tracer Tr.

The linear polarisations Q and U are patchy and irregular within the lobes, reflecting the turbulent magnetic field they contain. In the case of `v0.24` and `v0.24r2`, the jet spine is visible, being bounded by lines of high linear polarisation caused by the increased toroidal field strength along the jet edge (Figure 4.5). The change in sign of U above and below the jet injection zone reflects the change in sign of B_x that arises from the initial magnetic conditions (Equation 4.26). By contrast, the spine is not visible in the unbent `v0.24w0`—in the absence of a strong crosswind, the source expands equally in both the x and y directions, and so the jet spine is hidden behind a larger path length of turbulent lobe material which obscures it.

Rotation Measure

Simulated rotation measure maps are shown in Figure 4.16 for all simulations. We show the separate contributions to the RM from the unshocked external medium and the shocked

cocoon (the Faraday screen), as well as their combination. Simulations with a shared turbulent ICM (`v0.94`, `v0.24`, and `v0.24w0`) show similar external RM maps, the only significant difference arising from the shape of the backlighting source (the jet lobes). The screen RMs are also similar, showing the same tendency for high positive RM on the upper left and low negative RM on the lower left.

We find approximately equal contributions to the RM from the external medium and the Faraday screen. However, as the simulation domain size limits the line-of-sight extent, the path length through the external medium grows shorter as the simulation progresses and the lobes grow.

The RM screen components feature enhancements along the source edge which develop as the expanding jet lobe compresses the ICM magnetic field, in the process reordering it to lie preferentially along the cocoon surface. Along the edges of the observed source, this realigned magnetic field lies closer to the line of sight and so contributes to RM enhancement (Jerrim et al., 2024).

Real observations of radio sources have a limited resolution set by the telescope or array known as the 'beam size'. To investigate how our sources may appear to telescopes of different resolution, we additionally blur our raw data using a Gaussian filter with σ scaled to match a 10" and 30" beam observing a source of ESO 137-006's size and distance. The results are shown in Figure 4.17. While the effective 10" beam ($\sigma \sim 3$–4) can still resolve some details imprinted by the turbulence, including feathering and edge enhancements, all such details are lost at 30" ($\sigma \sim 9$–13).

We also calculate the Rotation Measure Structure Function (RMSF) (e.g. Jerrim et al., 2024; Laing et al., 2008; Leahy, 1987; Simonetti et al., 1984),

$$S(r) = \langle [\mathrm{RM}(r + r') - \mathrm{RM}(r')]^2 \rangle, \tag{4.37}$$

where r and r' are radial vectors in the xz-plane originating at $(0, 0)$ and $\langle\rangle$ indicates an average over all r'. In effect, we compute the difference in RM value between every pair of points $[\mathrm{RM}(r + r') - \mathrm{RM}(r')]$ within the jet region shown in Figure 4.16, along with the distance between those points r. Distances are then binned, and the average RM difference for a given distance is plotted as $S(r)$.

Observationally, RM values are typically only known at a low resolution and for certain parts of the radio lobe. In our simulations, however, we have many unique points of RM data, even in projection, which would make the calculation computationally tedious. To expedite the process, we randomly select a subsample of 5000 points within the 2D source outline and

calculate the structure function for these points only. Distances are binned within 0.1 spatial units, and the result of this calculation is shown in Figure 4.18.

The RMSF is broadly similar across simulations for correlation lengths $r < 60$, and are qualitatively similar to those seen in other simulations (e.g. Jerrim et al., 2024) and in observation (e.g. Laing et al., 2008). Beyond $r = 60$, the simulations diverge in part due to the different run times and corresponding jet lengths.

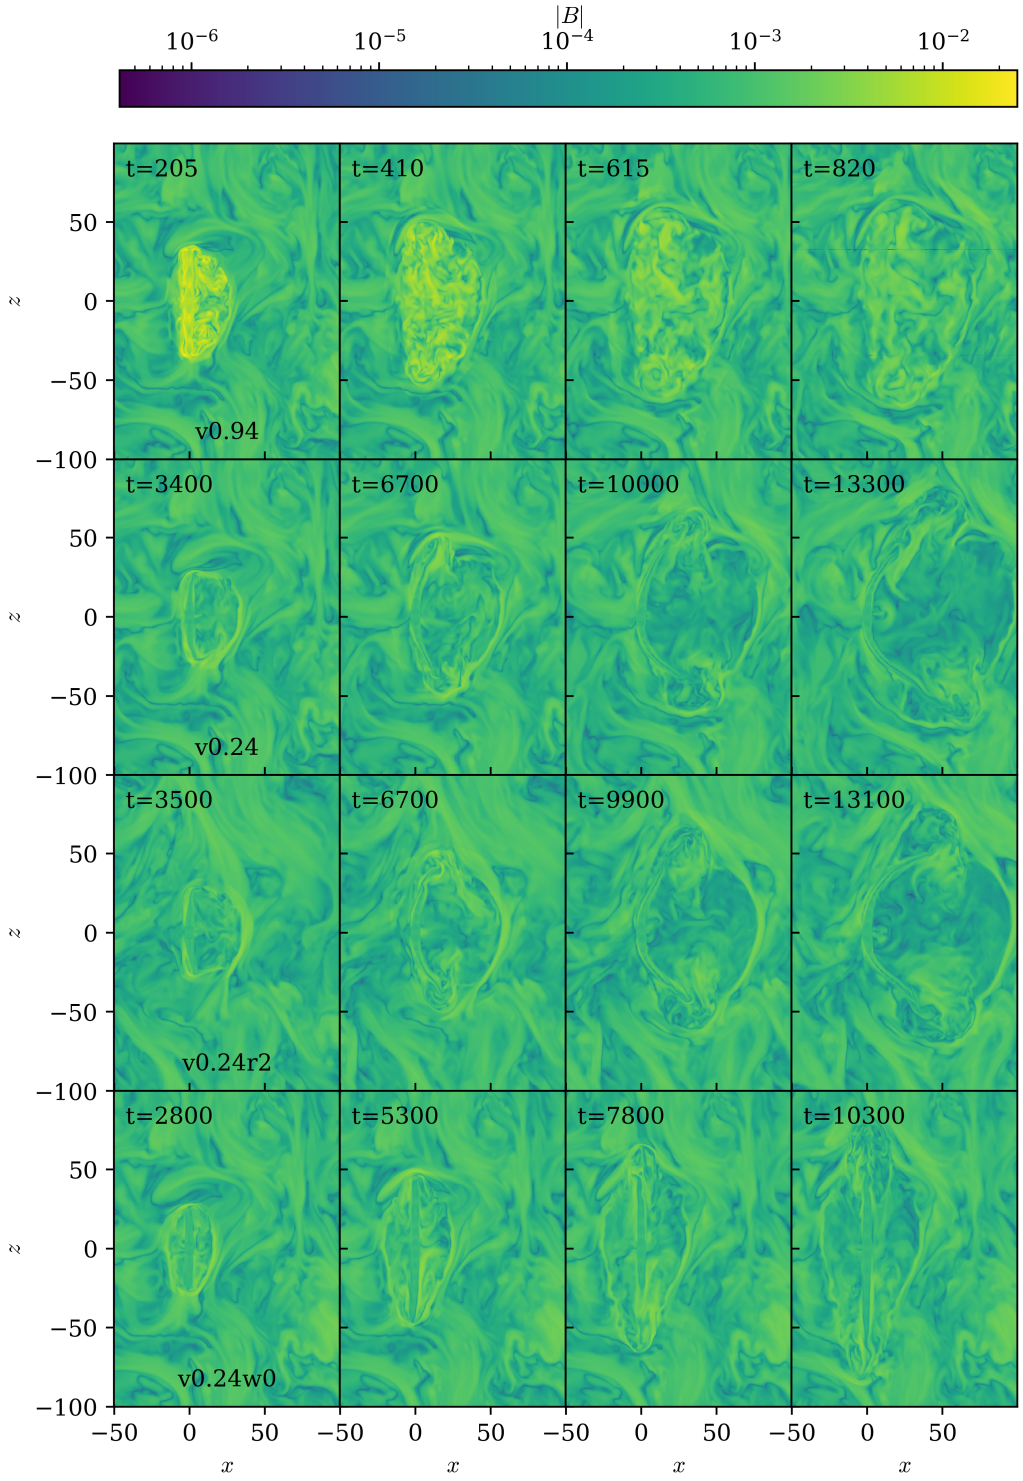


Fig. 4.12 Time evolution of the simulation magnetic field strength for, from top to bottom, `v0.94`, `v0.24`, `v0.24r2`, and `v0.24w0`. Midplane slices in $|B|$ are taken for the times indicated on each plot.

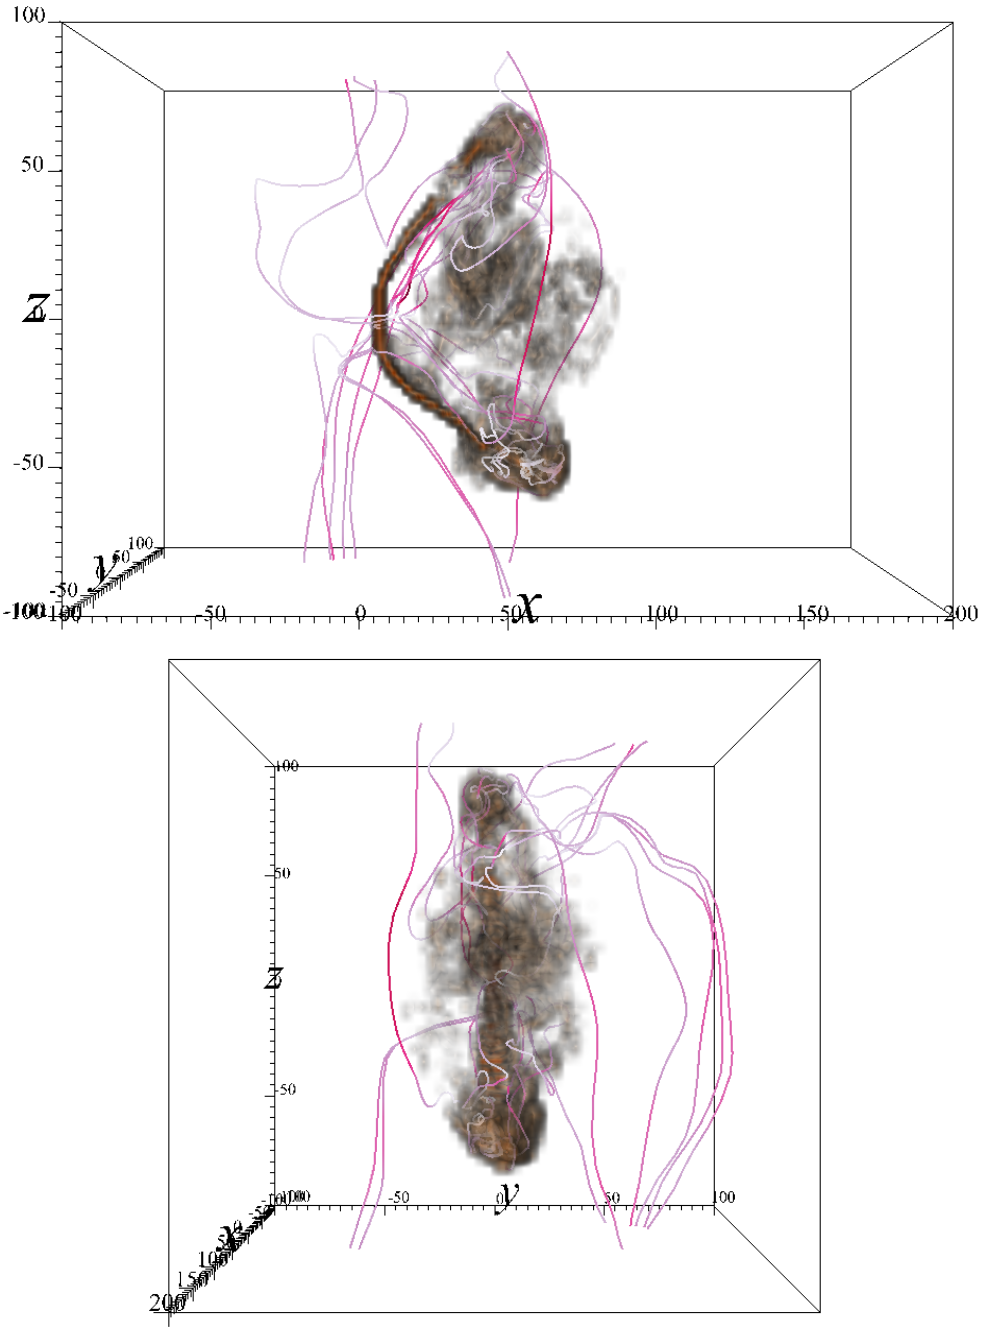


Fig. 4.13 3D magnetic field streamlines for simulation `v0.24` at time $t = 11,300$, viewed from the side (top), and from behind (bottom). Volume rendering corresponds to the fluid tracer with a cutoff below 0.3. A total of ten streamlines have been chosen, five with starting points in the upper lobe and five in the lower. They are coloured by magnitude with darker colours indicating a stronger field.

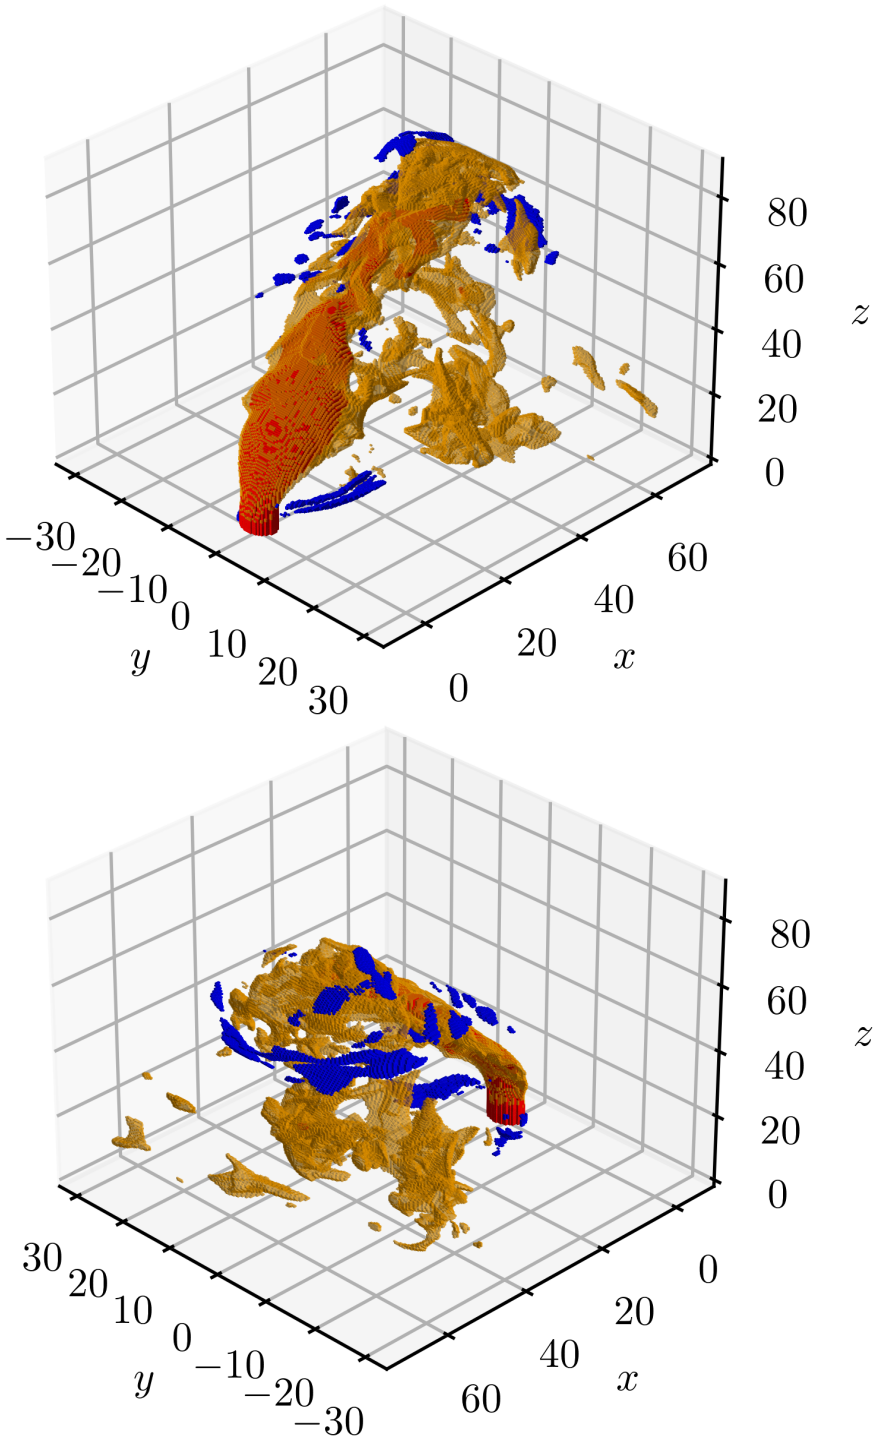


Fig. 4.14 3D voxels showing regions of elevated magnetic current density (blue) around the `v0.24` upper jet at time $t = 11,300$. Areas of high magnetic current density are defined as regions with $|\boldsymbol{\nabla} \times \boldsymbol{B}|$ more than half the maximum value. The jet (shown in yellow and orange) is defined as regions with jet tracer > 0.5 (yellow) and > 0.7 (orange).

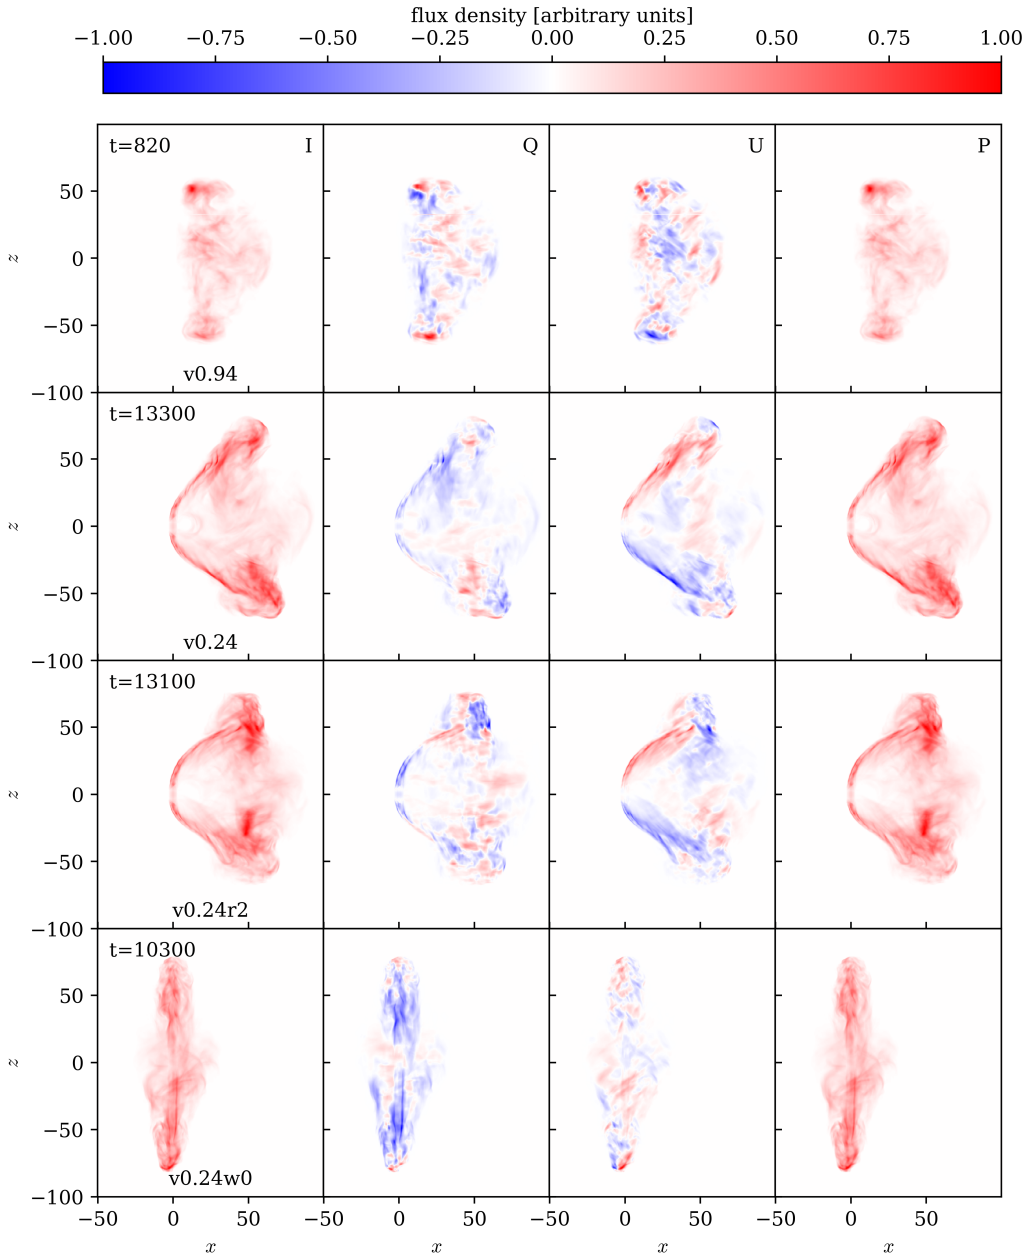


Fig. 4.15 Stokes parameters corresponding to, from left to right, the total polarised intensity I, linear polarisation at 90 degrees Q, linear polarisation at 45 degrees U, and total linear polarisation intensity P for simulations, from top to bottom, `v0.94` at $t = 820$, `v0.24` at $t = 13,300$, `v0.24r2` at $t = 13,100$, and `v0.24w0` at $t = 10,300$. A distinct asymmetry is visible for `v0.24` and `v0.24r2` in U, reflecting the change in sign of B_r with z. The same signal is not visible in `v0.24w0` due to the obscuration of the jet by the cocoon, nor in `v0.94` as its jet was terminated at $t = 200$.

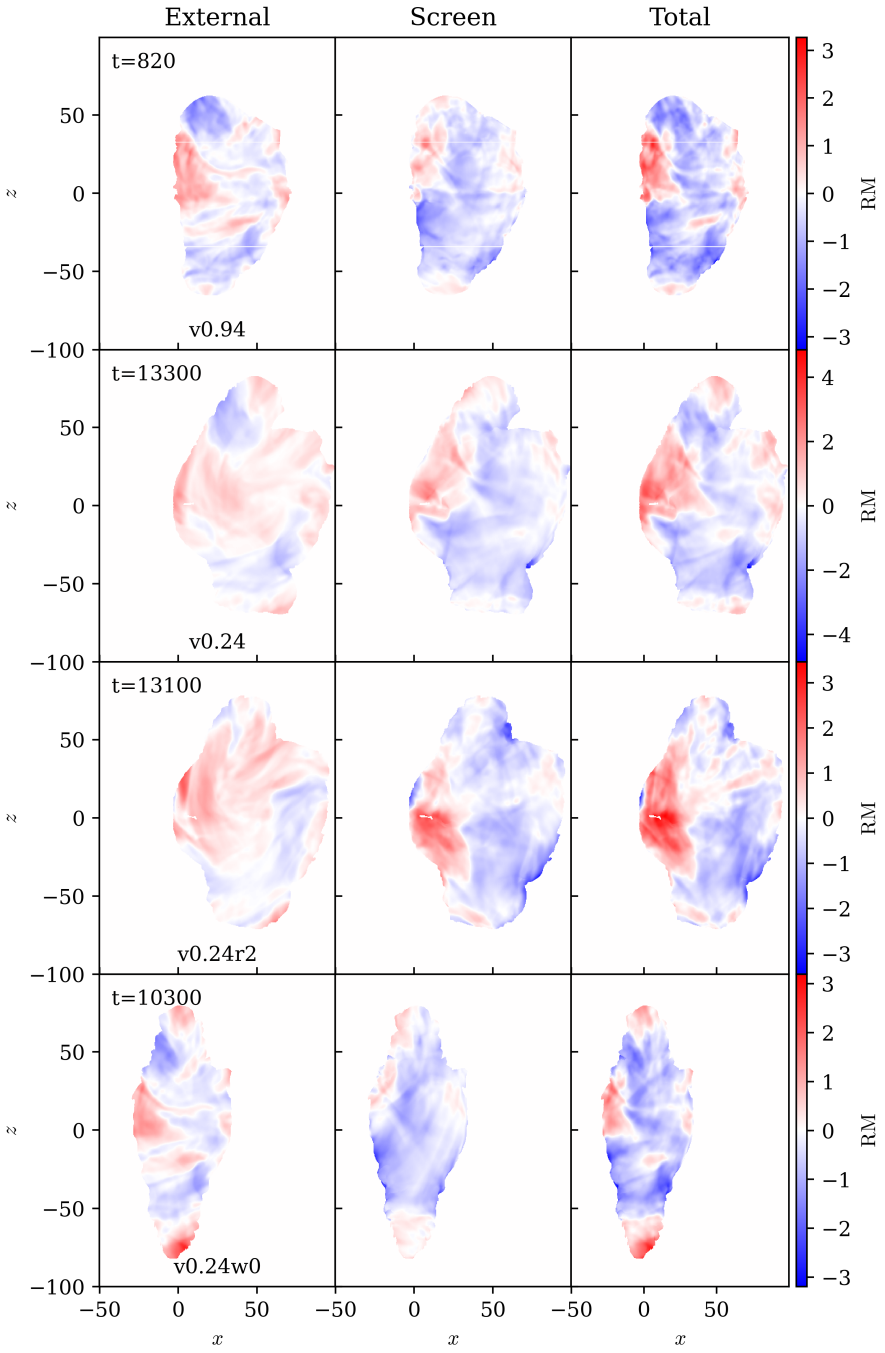


Fig. 4.16 Simulated Rotation Measure (RM) maps for (from *top* to *bottom*) `v0.94` at $t = 820$, `v0.24` at $t = 13,300$, `v0.24r2` at $t = 13,100$, and `v0.24w0` at $t = 10,300$. Maps are split into separate contributions from the external medium (*left*) and the Faraday screen immediately surrounding the jet (*middle*) and added to give the combined RM signal (*right*). The background field realisation has a significant effect on the combined RM signal despite the amplification of the field within the Faraday screen.

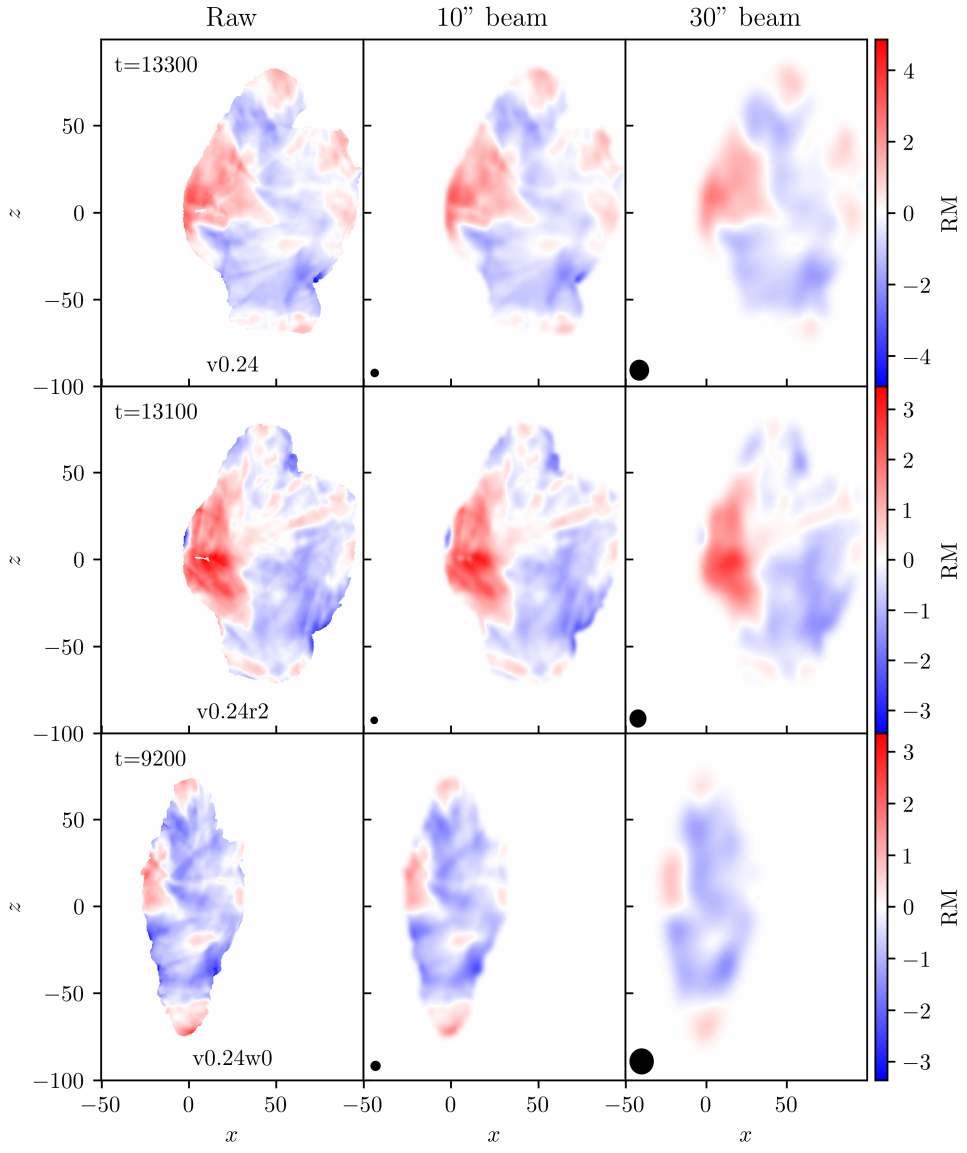


Fig. 4.17 The effect of beam size on simulated combined rotation measure (right-most column of Figure 4.16), for (from *top* to *bottom*) `v0.24` at $t = 13,300$, `v0.24r2` at $t = 13,100$, and `v0.24w0` at $t = 10,300$. Columns show the raw simulated data (*left*), convolution with an effective 10” beam (*middle*), and convolution with an effective 30” beam (*right*). As the beam size is increased, fine details are lost and edge enhancements become less obvious or disappear altogether. The effective beam size is shown in the lower left of each plot.

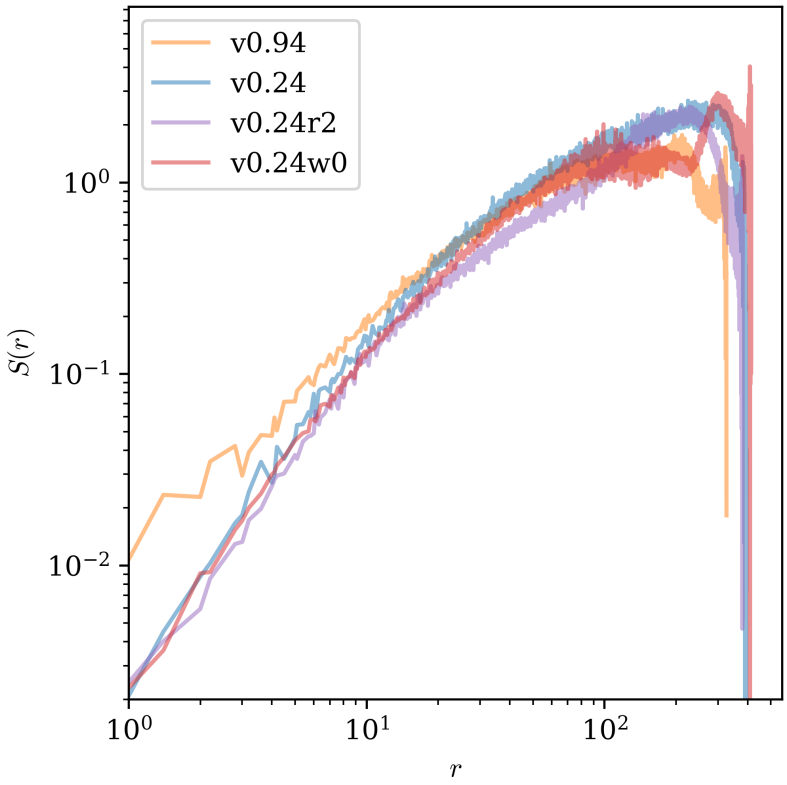


Fig. 4.18 The Rotation Measure Structure Function $S(r)$ for simulations `v0.94` at $t = 820$ (orange), `v0.24` at $t = 13,300$ (blue), `v0.24r2` at $t = 13,100$ (purple), and `v0.24w0` at $t = 10,300$ (red).

4.4 Discussion and conclusion

We have presented results from a suite of light (density contrast $\rho_j/\rho_a = 10^{-3}$), fast ($v_j = [0.94, 0.24]$), 3D RMHD simulated jets launched into pre-evolved turbulent backgrounds and subjected to transonic crosswinds.

To aid in our study of filament formation and potential particle escape along field lines, we use pre-evolved turbulent backgrounds generated using *PLUTO*'s `FORCED_TURB` module. Turbulence is driven to a desired RMS velocity of $\sim 0.3c_s$ and contains a mixture of solenoidal and compressive modes which form filamentary structures. At first glance, the distribution of magnetic field components B_i in our turbulent medium appears consistent with that generated by the Gaussian random field method; however, a more careful analysis reveals how realistic turbulence deviates from Gaussian, with filaments contributing to more extreme statistics and causing fatter distribution tails that are visible in the skewness $\mathcal{S}$ and kurtosis $\mathcal{K}$.

We find that our slower jets `v0.24` and `v0.24r2` bend to the degree predicted by theory (Jones et al., 2017; O'Neill et al., 2019a), and produce an FR-I WAT morphology, while `v0.94` produces an unbent FR-II. In the absence of a crosswind (`v0.24w0`), the source morphology is more ambiguous, supporting the idea that the WAT FR-I morphology is partly attributed to the presence of the wind and that jet morphology, in general, is affected by environmental factors. The ends of the slower jets are unstable and exhibit a periodic 'flapping' behaviour which carves out large terminating lobes, seeds further turbulent structure between the jets and, at times, causes the jet spine to become disconnected from the end of the lobe.

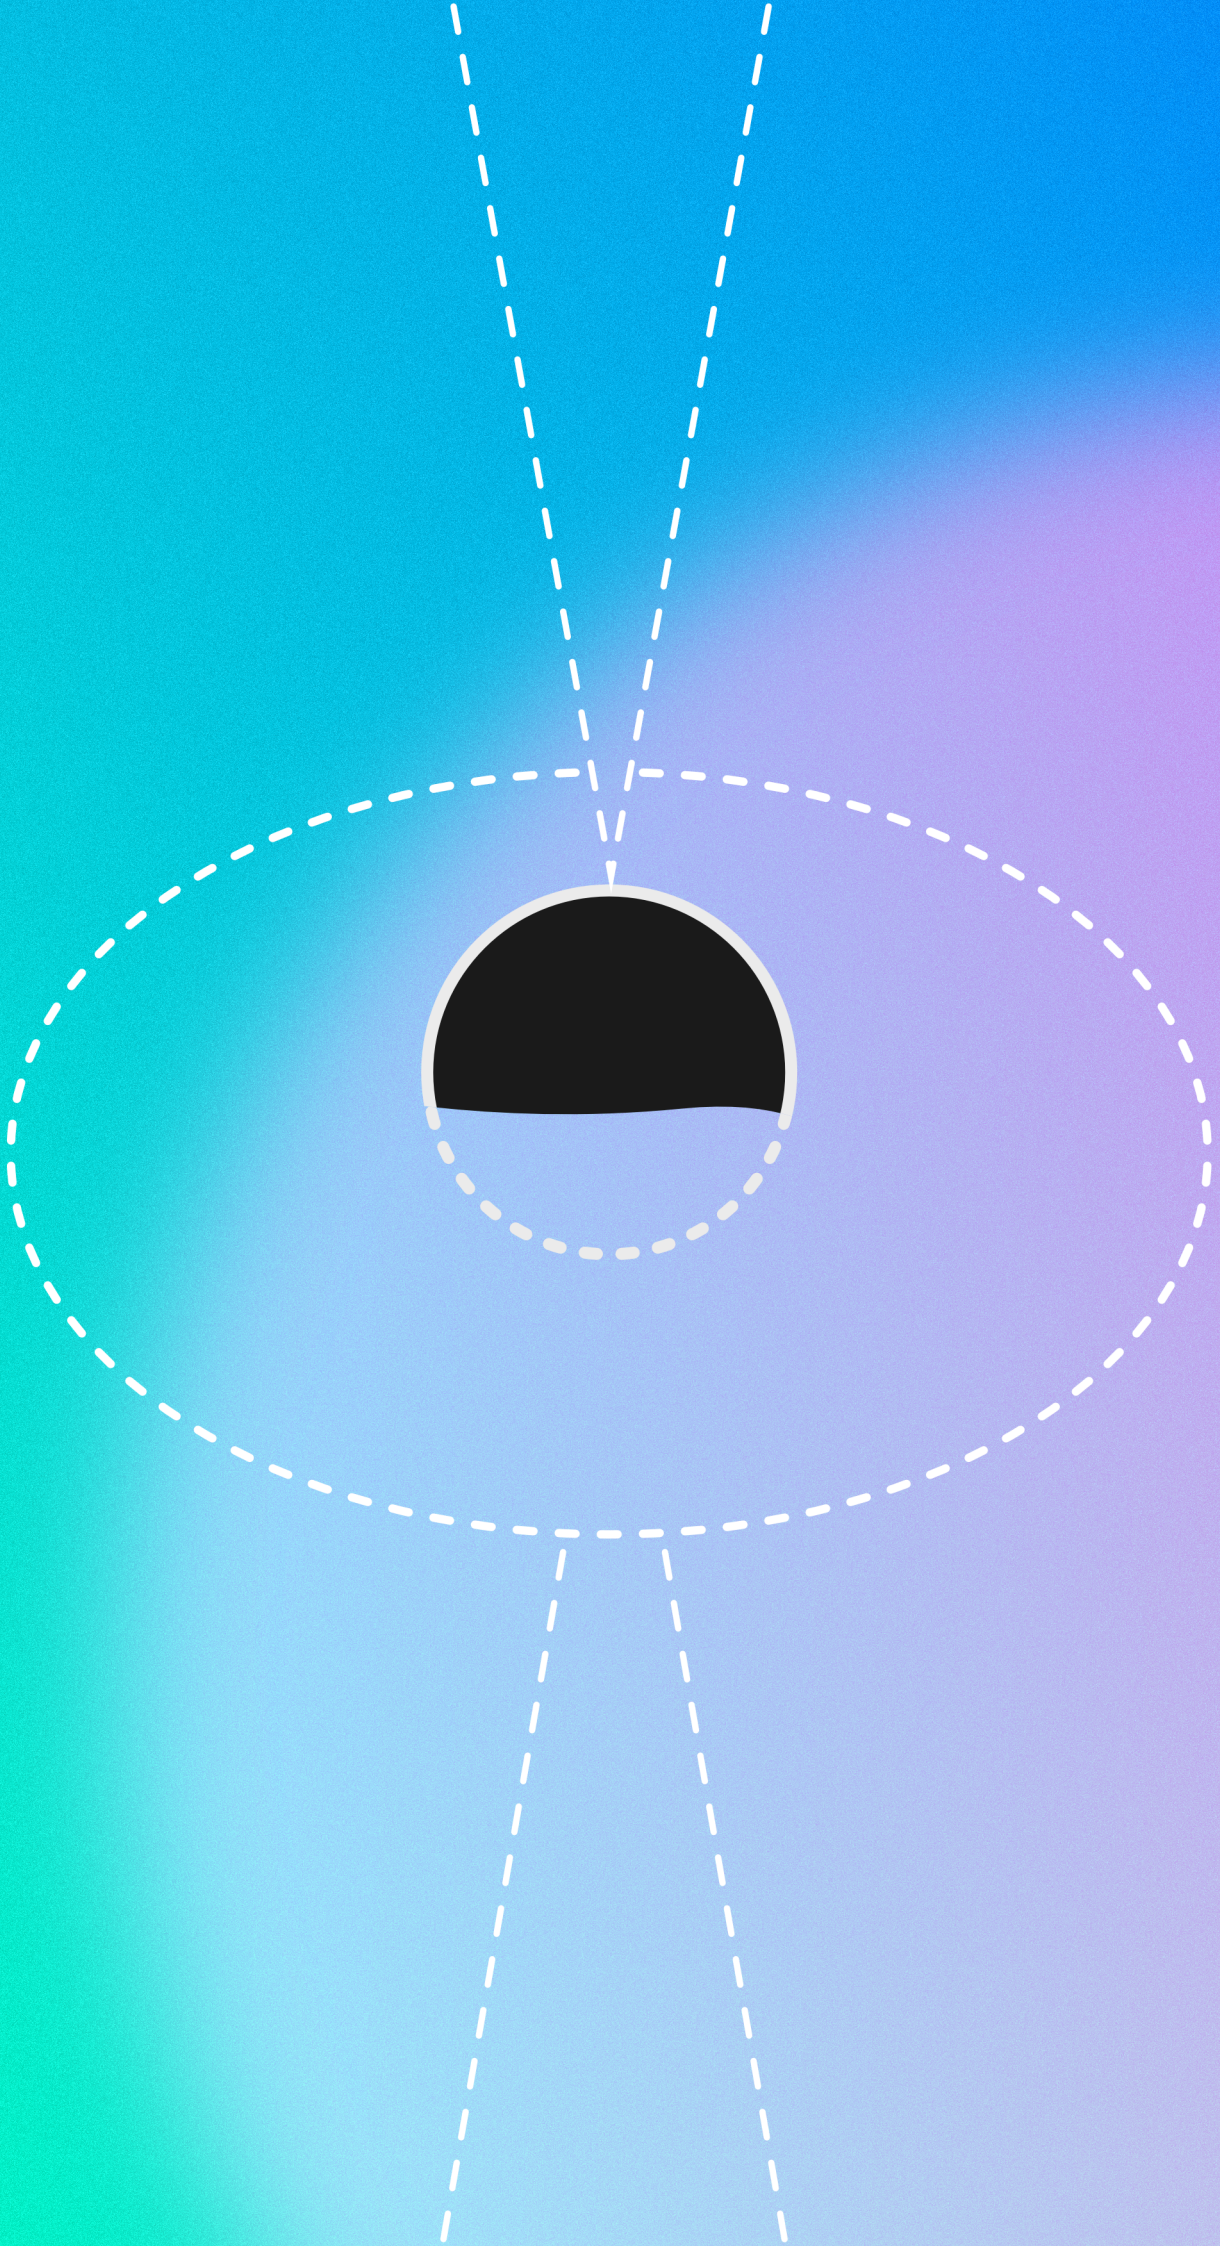
Magnetic reconnection occurs along the edge of the jet lobe and shocked ICM. This reconnection allows the formation of magnetic streamlines, which, when viewed in projection, appear to link the opposing jet lobes. However, viewing from other angles makes it clear that these filaments are displaced away from the lobes along the line of sight. The location of filaments is likely to change with time as the expanding jet/lobe intersects with different ambient field features and is different for different initial background magnetic fields. In `v0.24`, filaments emerge preferentially from the upper jet/lobe area, while in `v0.24r2` the preference moves to the lower jet/lobe area.

We also produce simulated observational signatures through polarization and Rotation Measure maps. Polarisation maps show turbulent structure within the lobes, with the back-swept `v0.24` and `v0.24r2` having clear jet spines on imaging, with strong Stokes Q and U (linear polarisation at 90 and 45 degrees, respectively) along the jet spine walls. This

signature is less apparent in the unbent `v0.24w0`, and so is likely due to the crosswind stripping the jet spine of the surrounding turbulent lobe material, leading to a near-direct observation of the intrinsic jet magnetic field. Hence, polarisation measurements of the inner jets of WAT and NAT sources may provide the most direct path to probe the jet magnetic field. Our simulation Rotation Measures indicate an approximately equal contribution to RM from the external medium and the screen (cocoon), although the domain size limits this conclusion. We reproduce the common finding of RM enhancement along the source edges caused by the compression and realignment of the ICM field within the cocoon.

In the absence of fully simulated particle emission, we cannot strictly verify the morphology of our sources nor directly observe filament formation, as the charged electrons necessary for this process are missing. We can, however, say with high confidence that magnetic streamlines can form between jet lobes and the surrounding turbulent ICM, with the exact location of these streamlines dependent on the structure of the ICM magnetic field and, in particular, the presence of large-scale ICM field structure in the form of filaments. Further, if the process of mass-loading magnetic streamlines is not too onerous, our results suggest that radio lobes should be host to many more radio filaments than are currently detected, although it may be that most are below current detection limits. Our results highlight the critical importance of using realistic ICM fields in studying jet-ICM interactions and the need for a greater understanding of the nature of ICM turbulence. The appropriate solenoidal weighting ζ of turbulence forcing is of particular interest, as this plays a large role in determining the strength and number of filaments in a turbulent fluid.

Software: PLUTO (Mignone et al., 2007), NumPy (Harris et al., 2020), pandas (McKinney, 2010), SciPy (Virtanen et al., 2020), VisIt (Childs et al., 2012)



Chapter 5

Conclusion and Outlook

5.1 Conclusion

In this thesis we have investigated the role of magnetic fields in the simulation of two parts of the Active Galactic Nucleus system—the accretion disk and the jet—finding that subtle changes in our initialisation or conceptualisation of the field leads to vastly different physical outcomes. As such, this work not only seeks to unveil the physical behaviour of real-world systems but also attempts to understand how the art of simulation itself imprints on the physics being simulated.

In Chapter 2, we sought to understand how simulated accretion disk evolution might be impacted by the initialisation of a purely toroidal field rather than one containing some poloidal component. We presented the results of four simulations, one with an initial magnetic field strength in line with current norms (magnetic plasma $\beta = 5$; `b5`), another with a much weaker initial magnetic field strength ($\beta = 200$; `b200`), and two higher-resolution counterparts. In the abstract of this thesis, I posed the question: Can a simulated accretion disk self-generate the strong, large-scale poloidal field needed to launch relativistic jets? Our results provide a tentative 'yes', but it's a qualified one, and, as is usual when dealing with magnetic fields, the situation is complicated. While poloidal fields are undoubtedly generated and the field near the BH event horizon is well-ordered and strong in the `b5` simulation, the case is less strong for the `b200` simulation, whose field growth stalls. Indeed, contrary to popular thought (and, perhaps, hope) the `b200` simulation never converges to match the stronger `b5`, despite being run for $4 \times 10^5\ GM/c^2$. Additionally, the initially weak field simulation `b200` is plagued by a strange 'lump' that forms early in the disk evolution and only becomes more prominent in time. This overdensity or $m = 1$ mode appears at a radius

associated with a down-step in resolution, meaning its formation may have been seeded or amplified by the numerics of simulation rather than physics.

We expanded our study of disk magnetic fields in Chapter 3, turning our attention to mean-field dynamo theory and calculating the α_{ij} components which define the major actions of the $\alpha\Omega$ dynamo, with a particular interest in $\alpha_{\phi\phi}$ which describes the conversion of the toroidal field into poloidal. Making the most of the data available, we calculated these coefficients directly from simulation data. We trialled two different means for our mean-field theory: (1) the typical full azimuthal average and (2) a more local Gaussian filter average. We found that these methods returned consistent results for the initially strong-field simulation `b5`, retrieving $\alpha_{\phi\phi}$ magnitudes and signs broadly consistent with previous works. Results were more mixed for simulation `b200`, with its overdensity causing difficulties for the global averaging method. For this case, the local Gaussian filter comes through and returns α_{ij} signs consistent with those of `b5`. Further, we find that the goodness-of-fit of our closure $\langle \boldsymbol{\varepsilon}'_i \rangle = \alpha_{ii} \langle \boldsymbol{B}_i \rangle + C_i$ improves with time, suggesting that dynamo processes outside of that captured by the linear mean-field theory play a more important role in early field evolution. Any direct, physical interpretation of the α components is a dangerous game, however, as the applicability of mean-field theory to accretion disks is already debatable. In particular, our findings in relation to the local Gaussian filter method raise the question of what constitutes the 'mean' in 'mean-field theory'.

In Chapter 4, we moved away from the accretion disk and on to a nearby related phenomenon: AGN jets. Again, simulated magnetic fields sat at the heart of our problem as we asked ourselves whether relativistic jets launched into a realistically turbulent background and swept back by a strong crosswind could form the magnetic structures necessary to explain the synchrotron filaments increasingly observed in radio observations of such objects. We performed a suite of low-density ($\rho_j/\rho_a = 10^{-3}$) jet simulations, testing jet speed ($v_j = 0.24c$ versus $0.94c$), wind speed (transonic versus none), and individual turbulent background realisations. We found that our pre-evolved turbulent backgrounds appear Gaussian through the first and second statistical moments (the mean and standard deviation) but actually feature fatter, more extended tails due to filamentation, highlighting the need for such realistic turbulent fields in simulation. Our jet simulations `v0.24` and `v0.24r2` successfully reproduced a head-tail morphology, and we were able to identify magnetic streamlines connecting the jet lobes to the turbulent ICM. The connection point between the lobe and the cocoon was found to be associated with regions of high magnetic field current density, suggesting magnetic reconnection. Filament-forming-streamlines were not found to connect the two lobes, as the intervening material is highly turbulent due both to the KH instability acting along the jet

edge and also due to a 'wobbling' or 'flapping' behaviour at the jet ends. We additionally produced simulated polarisation and Rotation Measure maps, finding (1) that the jet spine of bent sources is unobscured by turbulent lobe material, and hence the jet magnetic field may be observed more directly in this region, and (2) that the external ICM and shocked ICM contribute roughly equally to the RM (although constrained by the domain size), with the shocked ICM cocoon additionally contributing regions of localised RM enhancement along the source edges due to compression of the ICM field with the expanding lobes.

5.2 Future Outlook

Within this thesis, we have seen how both further simplifications *and* greater complexity are needed where magnetic fields and the art of simulation meet.

The field of accretion disk studies sits at a crossroads. With ever-increasing computing power available, there is a push to add greater complexity to simulations—higher resolution, radiation, particles, thermal decoupling, and more. However, some of the foundational physics remains uncertain, and without simple, standardised test cases, we run the risk of being increasingly unable to interpret results, understand differences between models, or even apply simulated work to observations. To tackle these issues, a two-pronged approach is needed. Firstly, the effect of simulation itself—the computational uncertainties and inaccuracies—needs to be understood. Of particular interest to magnetic fields is magnetic helicity and its conservation. No code yet exists that conserves magnetic helicity, but there is increasing interest, and the results of such future codes will undoubtedly improve our understanding of the dynamo. Secondly, as our results highlight the importance of initial conditions in simulated disk evolution, more information is needed about realistic initial conditions for real disks. Simulation of disk formation through Tidal Disruption Events will be the most likely path to answer this question, however, it will be a computationally heavy experiment as high resolution, long timescales, and large simulation domains will all be needed.

For AGN jets, understanding the interactions between jet magnetic fields and ICM fields will require a better understanding of the turbulent ICM's structure. The statistics of the ICM magnetic field are poorly constrained, and Rotation Measures (the only semi-direct form of measurement available) are historically under-utilised, with many sources suffering from sparse RM data and much of the existing RM data being fit from observations at only two frequencies (Brentjens & de Bruyn, 2005; Taylor et al., 2009; Vanderwoude et al., 2024), making it potentially unreliable. The upcoming Polarisation Sky Survey of the Universe's

Magnetism (POSSUM; Gaensler et al., 2010) from the Australian Square Kilometre Array Pathfinder (ASKAP; Hotan et al., 2021) promises to alleviate these pains to some degree by producing RM maps over $\sim 20,000$ degrees of the southern sky at ~ 1 GHz (Vanderwoude et al., 2024). The known population of well-resolved radio filaments is also set to expand greatly as more observations are undertaken with *LOFAR* and *MeerKAT*, and later with the completion of the Square Kilometre Array Observatory (SKAO; Dewdney et al., 2009) with science verification beginning 2027. The SKAO is an international project and comprises two parts: SKA-Low, operating between 50–350 MHz to be located in Australia, and SKA-mid, covering 350 MHz–15.4 GHz and located in South Africa. The steep radio spectral indices of known filaments imply that filaments are filled with older electrons undergoing synchrotron and inverse Compton losses, so future high-resolution, low-frequency observations with the SKAO will be critical in furthering our understanding of such filaments. In addition to more and better observational data, better modelling is also needed. The current state-of-the-art in ICM background field generation is that of the Gaussian random field, but this is clearly inadequate for describing the ICM field structure. With the addition of the `FORCED_TURB` module in a future release of *PLUTO*, more simulations can be performed with realistic turbulent backgrounds. Further observations of CSTs may then provide critical information about the ICM field near sources of relativistic electrons.

References

Alexander, D. M., & Hickox, R. C. 2012, NewAR, 56, 93

Antonucci, R. 1993, ARA&A, 31, 473

Asada, K., Inoue, M., Kameno, S., & Nagai, H. 2008, ApJ, 675, 79

Asada, K., Inoue, M., Uchida, Y., et al. 2002, PASJ, 54, L39

Asada, K., & Nakamura, M. 2012, ApJL, 745, L28

Bahcall, J. N., Kirhakos, S., & Schneider, D. P. 1996, ApJ, 457, 557

Bai, X.-N., & Stone, J. M. 2013, ApJ, 767, 30

Balbus, S. A. 2003, ARA&A, 41, 555

Balbus, S. A., & Hawley, J. F. 1991, ApJ, 376, 214

—. 1998, Reviews of Modern Physics, 70, 1

Balbus, S. A., Hawley, J. F., & Stone, J. M. 1996, ApJ, 467, 76

Baldi, R. D., Capetti, A., & Giovannini, G. 2019, MNRAS, 482, 2294

Balsara, D. S., & Norman, M. L. 1992, ApJ, 393, 631

Bambic, C. J., Morsony, B. J., & Reynolds, C. S. 2018a, ApJ, 857, 84

Bambic, C. J., Pinto, C., Fabian, A. C., Sanders, J., & Reynolds, C. S. 2018b, MNRAS, 478, L44

Bardou, A., von Rekowski, B., Dobler, W., Brandenburg, A., & Shukurov, A. 2001, A&A, 370, 635

Beck, R., & Krause, M. 2005, Astronomische Nachrichten, 326, 414

Beckwith, K., & Stone, J. M. 2011, ApJS, 193, 6

Begelman, M. C., Scepi, N., & Dexter, J. 2022, MNRAS, 511, 2040

Begelman, M. C., & Volonteri, M. 2017, MNRAS, 464, 1102

Bempong-Manful, E., Hardcastle, M. J., Birkinshaw, M., et al. 2020, MNRAS, 496, 676

Bicknell, G. V. 1994, ApJ, 422, 542

Blandford, R., Meier, D., & Readhead, A. 2019, ARA&A, 57, 467

Blandford, R. D., & Payne, D. G. 1982, MNRAS, 199, 883

Blandford, R. D., & Znajek, R. L. 1977, MNRAS, 179, 433

Bonafede, A., Feretti, L., Murgia, M., et al. 2010, A&A, 513, A30

Bowman, M., Leahy, J. P., & Komissarov, S. S. 1996, MNRAS, 279, 899

Boyce, P. J., Disney, M. J., Blades, J. C., et al. 1996, ApJ, 473, 760

Braginskiĭ, S. I. 1965, ZhETF, 47, 2178, [English translation: Sov. Phys.-JETP (1965), 20, 6]

Brandenburg, A., & Donner, K. J. 1997, MNRAS, 288, L29

Brandenburg, A., Nordlund, A., Stein, R. F., & Torkelsson, U. 1995, ApJ, 446, 741

Brandenburg, A., Sokoloff, D., & Subramanian, K. 2012, SSRv, 169, 123

Brandenburg, A., & Subramanian, K. 2005, PhR, 417, 1

Brandenburg, A., Svedin, A., & Vasil, G. M. 2009, MNRAS, 395, 1599

Brentjens, M. A., & de Bruyn, A. G. 2005, A&A, 441, 1217

Bridle, A. H., Hough, D. H., Lonsdale, C. J., Burns, J. O., & Laing, R. A. 1994, AJ, 108, 766

Bridle, A. H., & Perley, R. A. 1984, ARA&A, 22, 319

Brunetti, G., & Jones, T. W. 2014, International Journal of Modern Physics D, 23, 1430007

Brunetti, G., Setti, G., & Comastri, A. 1997, A&A, 325, 898

Burn, B. J. 1966, MNRAS, 133, 67

Candini, S., Brienza, M., Bonafede, A., et al. 2023, A&A, 677, A4

Carilli, C. L., & Taylor, G. B. 2002, ARA&A, 40, 319

Carroll, B. W., & Ostlie, D. A. 2017, An introduction to modern astrophysics, Second Edition (Cambridge University Press)

Chandrasekhar, S. 1960, Proceedings of the National Academy of Science, 46, 253

Charbonneau, P. 2010, Liv. Rev. Sol. Phys., 7, 3

Childs, H., Brugger, E., Whitlock, B., et al. 2012, High Performance Visualization–Enabling Extreme-Scale Scientific Insight, doi:10.1201/b12985

Condon, J. J., Cotton, W. D., White, S. V., et al. 2021, ApJ, 917, 18

Cowling, T. G. 1933, MNRAS, 94, 39

Cucchetti, E., Clerc, N., Pointecouteau, E., Peille, P., & Pajot, F. 2019, A&A, 629, A144

Curry, C., Pudritz, R. E., & Sutherland, P. G. 1994, ApJ, 434, 206

Das, H. K., & Gronke, M. 2024, MNRAS, 527, 991

Davis, S. W., Stone, J. M., & Pessah, M. E. 2010, ApJ, 713, 52

Davis, S. W., & Tchekhovskoy, A. 2020, ARA&A, 58, 407

Dewdney, P. E., Hall, P. J., Schilizzi, R. T., & Lazio, T. J. L. W. 2009, IEEE Proceedings, 97, 1482

Dhang, P., Bendre, A., Sharma, P., & Subramanian, K. 2020, MNRAS, 494, 4854

Dhang, P., & Sharma, P. 2019, MNRAS, 482, 848

Di Gennaro, G., Venturi, T., Giacintucci, S., et al. 2024, arXiv e-prints, arXiv:2408.14142

Dolag, K., Borgani, S., Murante, G., & Springel, V. 2009, MNRAS, 399, 497

Dreher, J. W., & Feigelson, E. D. 1984, Nature, 308, 43

Dursi, L. J., & Pfrommer, C. 2008, ApJ, 677, 993

English, W., Hardcastle, M. J., & Krause, M. G. H. 2016, MNRAS, 461, 2025

Evans, C. R., & Hawley, J. F. 1988, ApJ, 332, 659

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, ApJL, 875, L5

Fabian, A. C. 2012, ARA&A, 50, 455

Falanga, M., Melia, F., Tagger, M., Goldwurm, A., & Bélanger, G. 2007, ApJL, 662, L15

Fanaroff, B. L., & Riley, J. M. 1974, MNRAS, 167, 31P

Feain, I. J., Papadopoulos, P. P., Ekers, R. D., & Middelberg, E. 2007, ApJ, 662, 872

Federrath, C., Roman-Duval, J., Klessen, R. S., Schmidt, W., & Mac Low, M. M. 2010, A&A, 512, A81

Ferrarese, L., & Merritt, D. 2000, ApJL, 539, L9

Fischer, T. C., Kraemer, S. B., Schmitt, H. R., et al. 2018, ApJ, 856, 102

Fishbone, L. G., & Moncrief, V. 1976, ApJ, 207, 962

Flock, M., Dzyurkevich, N., Klahr, H., & Mignone, A. 2010, A&A, 516, A26

Flock, M., Dzyurkevich, N., Klahr, H., Turner, N., & Henning, T. 2012, ApJ, 744, 144

Fomalont, E. B., Ebneter, K. A., van Breugel, W. J. M., & Ekers, R. D. 1989, ApJL, 346, L17

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306

Fragile, P. C., & Sądowski, A. 2017, MNRAS, 467, 1838

Fricke, K. 1969, A&A, 1, 388

Gabuzda, D. C. 2021, Galaxies, 9, 58

Gabuzda, D. C., Nagle, M., & Roche, N. 2018, A&A, 612, A67

Gaensler, B. M., Landecker, T. L., Taylor, A. R., & POSSUM Collaboration. 2010, in American Astronomical Society Meeting Abstracts, Vol. 215, American Astronomical Society Meeting Abstracts #215, 470.13

Gaibler, V. 2008, PhD thesis, Ruprecht-Karls University of Heidelberg, Germany

Gan, Z., Li, H., Li, S., & Yuan, F. 2017, ApJ, 839, 14

Gebhardt, K., Bender, R., Bower, G., et al. 2000, ApJL, 539, L13

Gendron-Marsolais, M., Hlavacek-Larrondo, J., van Weeren, R. J., et al. 2020, MNRAS, 499, 5791

Gendron-Marsolais, M. L., Hull, C. L. H., Perley, R., et al. 2021, ApJ, 911, 56

Ghisellini, G. 2011, in American Institute of Physics Conference Series, Vol. 1381, 25th Texas Symposium on Relativistic AstroPhysics (Texas 2010), ed. F. A. Aharonian, W. Hofmann, & F. M. Rieger (AIP), 180–198

Ghisellini, G., & Celotti, A. 2001, A&A, 379, L1

Ghosh, S., & Mukhopadhyay, B. 2020, MNRAS, 496, 4191

Giacintucci, S., Venturi, T., Markevitch, M., et al. 2022, ApJ, 934, 49

Giannios, D., & Uzdensky, D. A. 2019, MNRAS, 484, 1378

Gilbert, A. D. 2002, Geophysical and Astrophysical Fluid Dynamics, 96, 135

Goodman, J., & Xu, G. 1994, ApJ, 432, 213

Gopal-Krishna, & Wiita, P. J. 2000, A&A, 363, 507

Gourgouliatos, K. N., & Komissarov, S. S. 2018, Nature Astronomy, 2, 167

Govoni, F., Murgia, M., Vacca, V., et al. 2017, A&A, 603, A122

Graham, J., Fabian, A. C., & Sanders, J. S. 2008, MNRAS, 386, 278

Greenstein, J. L., & Matthews, T. A. 1963, AJ, 68, 279

Gressel, O. 2010, MNRAS, 405, 41

Gressel, O., & Pessah, M. E. 2015, ApJ, 810, 59

Gruzinov, A. V., & Diamond, P. H. 1994, Phys. Rev. Lett., 72, 1651

Guan, X., Gammie, C. F., Simon, J. B., & Johnson, B. M. 2009, ApJ, 694, 1010

Hardcastle, M. J. 2001, A&A, 373, 881

Hardcastle, M. J., Birkinshaw, M., Cameron, R. A., et al. 2002, ApJ, 581, 948

Hardcastle, M. J., Birkinshaw, M., & Worrall, D. M. 2001, MNRAS, 326, 1499

Hardcastle, M. J., & Krause, M. G. H. 2014, MNRAS, 443, 1482

Häring, N., & Rix, H.-W. 2004, ApJL, 604, L89

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357

Harrison, C. M., Alexander, D. M., Mullaney, J. R., & Swinbank, A. M. 2014, MNRAS, 441, 3306

Harrison, C. M., Alexander, D. M., Swinbank, A. M., et al. 2012, MNRAS, 426, 1073

Harwood, J. J., Vernstrom, T., & Stroe, A. 2020, MNRAS, 491, 803

Hawley, J. F. 1990, ApJ, 356, 580

Hawley, J. F., Balbus, S. A., & Winters, W. F. 1999, ApJ, 518, 394

Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1995, ApJ, 440, 742

—. 1996, ApJ, 464, 690

Hawley, J. F., Guan, X., & Krolik, J. H. 2011, ApJ, 738, 84

Hawley, J. F., Richers, S. A., Guan, X., & Krolik, J. H. 2013, ApJ, 772, 102

Hazard, C., Mackey, M. B., & Shimmins, A. J. 1963, Nature, 197, 1037

Herzenberg, A. 1958, Phil. Trans. Roy. Soc. London Ser. A, 250, 543

Heywood, I., Camilo, F., Cotton, W. D., et al. 2019, Nature, 573, 235

Heywood, I., Rammala, I., Camilo, F., et al. 2022, ApJ, 925, 165

Hogg, J. D., & Reynolds, C. S. 2016, ApJ, 826, 40

—. 2018a, ApJ, 854, 6

—. 2018b, ApJ, 861, 24

Hotan, A. W., Bunton, J. D., Chippendale, A. P., et al. 2021, PASA, 38, e009

Hovatta, T., Lister, M. L., Aller, M. F., et al. 2012, AJ, 144, 105

Huarte-Espinosa, M., Krause, M., & Alexander, P. 2011a, MNRAS, 417, 382

—. 2011b, MNRAS, 418, 1621

Hughes, D. W., & Tobias, S. M. 2010, in Relaxation Dynamics in Laboratory and Astrophysical Plasmas, ed. P. H. Diamond, X. Garbet, P. Ghendrih, & Y. Sarazin (World Scientific Publishing Co. Pte. Ltd), 15–48

Husemann, B., Scharwächter, J., Bennert, V. N., et al. 2016, A&A, 594, A44

Jerrim, L. A., Shabala, S. S., Yates-Jones, P. M., et al. 2024, MNRAS, arXiv:2311.12363

Ji, H., Burin, M., Schartman, E., & Goodman, J. 2006, Nature, 444, 343

Jiang, Y.-F., Stone, J. M., & Davis, S. W. 2019, ApJ, 880, 67

Johansen, A., & Levin, Y. 2008, A&A, 490, 501

Jonas, J., & MeerKAT Team. 2016, in MeerKAT Science: On the Pathway to the SKA, 1

Jones, T. W., Nolting, C., O'Neill, B. J., & Mendygral, P. J. 2017, Physics of Plasmas, 24, 041402

Kaiser, C. R., & Alexander, P. 1997, MNRAS, 286, 215

Karatzas, I. 1988, Brownian motion and stochastic calculus / Ioannis Karatzas, Steven E. Shreve., Graduate texts in mathematics (Springer)

Kellermann, K. I., Sramek, R., Schmidt, M., Shaffer, D. B., & Green, R. 1989, AJ, 98, 1195

Khatri, R., & Gaspari, M. 2016, MNRAS, 463, 655

King, A. 2003, ApJL, 596, L27

King, A., & Pounds, K. 2015, ARA&A, 53, 115

King, A. R., Zubovas, K., & Power, C. 2011, MNRAS, 415, L6

Knowles, K., Cotton, W. D., Rudnick, L., et al. 2022, A&A, 657, A56

Kolmogorov, A. 1941, Akademiia Nauk SSSR Doklady, 30, 301

Königl, A. 1991, ApJL, 370, L39

Kormendy, J., & Ho, L. C. 2013, ARA&A, 51, 511

Krause, F., & Rädler, K.-H. 1980, Mean-field magnetohydrodynamics and dynamo theory (Pergamon Press)

Krause, M., Alexander, P., Riley, J., & Hopton, D. 2012, MNRAS, 427, 3196

Kulsrud, R. M., & Anderson, S. W. 1992, ApJ, 396, 606

Laing, R. A. 1980, MNRAS, 193, 439

Laing, R. A., & Bridle, A. H. 2002, MNRAS, 336, 328

Laing, R. A., Bridle, A. H., Parma, P., & Murgia, M. 2008, MNRAS, 391, 521

Lambrides, E., Garofali, K., Larson, R., et al. 2024, arXiv e-prints, arXiv:2409.13047

Lazarian, A., & Vishniac, E. T. 1999, ApJ, 517, 700

Leahy, J. P. 1987, MNRAS, 226, 433

Ledlow, M. J., & Owen, F. N. 1996, AJ, 112, 9

Lesur, G., & Longaretti, P. Y. 2005, A&A, 444, 25

Lin, D. N. C., & Papaloizou, J. 1985, in Protostars and Planets II, ed. D. C. Black & M. S. Matthews, 981–1072

Liska, M., Tchekhovskoy, A., & Quataert, E. 2020, MNRAS, 494, 3656

Lister, M. L., Cohen, M. H., Homan, D. C., et al. 2009, AJ, 138, 1874

Liu, G., Zakamska, N. L., Greene, J. E., Nesvadba, N. P. H., & Liu, X. 2013, MNRAS, 436, 2576

Loeb, A., & Laor, A. 1992, ApJ, 384, 115

Loken, C., Roettiger, K., Burns, J. O., & Norman, M. 1995, ApJ, 445, 80

Lovelace, R. V. E., Li, H., Colgate, S. A., & Nelson, A. F. 1999, ApJ, 513, 805

Lynden-Bell, D., & Pringle, J. E. 1974, MNRAS, 168, 603

Lyutikov, M. 2006, MNRAS, 373, 73

Macfarlane, C., Best, P. N., Sabater, J., et al. 2021, MNRAS, 506, 5888

Machida, M., Matsumoto, R., & Mineshige, S. 2001, PASJ, 53, L1

Magorrian, J., Tremaine, S., Richstone, D., et al. 1998, AJ, 115, 2285

Mahatma, V. H., Basu, A., Hardcastle, M. J., Morabito, L. K., & van Weeren, R. J. 2023, MNRAS, 520, 4427

Mahmud, M., Gabuzda, D. C., & Bezrukovs, V. 2009, MNRAS, 400, 2

Maiolino, R., Scholtz, J., Curtis-Lake, E., et al. 2023, arXiv e-prints, arXiv:2308.01230

McClure-Griffiths, N. M., Madsen, G. J., Gaensler, B. M., McConnell, D., & Schnitzeler, D. H. F. M. 2010, ApJ, 725, 275

McConnell, N. J., & Ma, C.-P. 2013, ApJ, 764, 184

McKinney, J. C. 2005a, arXiv e-prints, astro

—. 2005b, arXiv e-prints, astro

McKinney, W. 2010, in Proceedings of the 9th Python in Science Conference, ed. S. van der Walt & J. Millman, 5 –61

McNamara, B. R., & Nulsen, P. E. J. 2007, ARA&A, 45, 117

—. 2012, New Journal of Physics, 14, 055023

Mendygral, P. J., Jones, T. W., & Dolag, K. 2012, ApJ, 750, 166

Mignon-Risse, R., Aimar, N., Varniere, P., Casse, F., & Vincent, F. 2021, in SF2A-2021: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics. Eds.: A. Siebert, ed. A. Siebert, K. Baillié, E. Lagadec, N. Lagarde, J. Malzac, J. B. Marquette, M. N'Diaye, J. Richard, & O. Venot, 113–116

Mignone, A., Bodo, G., Massaglia, S., et al. 2007, ApJS, 170, 228

Mignone, A., & McKinney, J. C. 2007, MNRAS, 378, 1118

Mignone, A., Plewa, T., & Bodo, G. 2005, ApJS, 160, 199

Mingo, B., Croston, J. H., Hardcastle, M. J., et al. 2019, MNRAS, 488, 2701

Mingo, B., Croston, J. H., Best, P. N., et al. 2022, MNRAS, 511, 3250

Mishra, B., Begelman, M. C., Armitage, P. J., & Simon, J. B. 2020, MNRAS, 492, 1855

Mishra, B., Fragile, P. C., Anderson, J., et al. 2022, ApJ, 939, 31

Moffatt, H. K. 1978, Magnetic field generation in electrically conducting fluids (Cambridge University Press)

Morsony, B. J., Miller, J. J., Heinz, S., et al. 2013, MNRAS, 431, 781

Murgia, M., Govoni, F., Feretti, L., et al. 2004, A&A, 424, 429

Nandra, K., Georgakakis, A., Willmer, C. N. A., et al. 2007, ApJL, 660, L11

Narayan, R., Igumenshchev, I. V., & Abramowicz, M. A. 2003, PASJ, 55, L69

Narayan, R., Sądowski, A., Penna, R. F., & Kulkarni, A. K. 2012, MNRAS, 426, 3241

Noble, S. C., Krolik, J. H., & Hawley, J. F. 2009, ApJ, 692, 411

—. 2010, ApJ, 711, 959

Nolting, C., Lacy, M., Croft, S., et al. 2022, ApJ, 936, 130

O'Dea, C. P., & Baum, S. A. 2023, Galaxies, 11, 67

Ohmura, T., Asano, K., Nishiwaki, K., Machida, M., & Sakemi, H. 2023, ApJ, 951, 76

Oke, J. B. 1963, Nature, 197, 1040

O'Neill, B. J., Jones, T. W., Nolting, C., & Mendygral, P. J. 2019a, ApJ, 884, 12

—. 2019b, ApJ, 887, 26

Owen, F. N., Eilek, J. A., & Kassim, N. E. 2000, ApJ, 543, 611

Pacucci, F., & Loeb, A. 2024, ApJ, 964, 154

Pacucci, F., & Narayan, R. 2024, arXiv e-prints, arXiv:2407.15915

Pacucci, F., Nguyen, B., Carniani, S., Maiolino, R., & Fan, X. 2023, ApJL, 957, L3

Padovani, P., Alexander, D. M., Assef, R. J., et al. 2017, A&A Rv, 25, 2

Papaloizou, J. C. B., & Lin, D. N. C. 1995, ARA&A, 33, 505

Papaloizou, J. C. B., & Pringle, J. E. 1984, MNRAS, 208, 721

Parker, E. N. 1955, ApJ, 122, 293

—. 1966, ApJ, 145, 811

—. 1992, ApJ, 401, 137

Penna, R. F., Sądowski, A., Kulkarni, A. K., & Narayan, R. 2013, MNRAS, 428, 2255

Perley, R. A., Dreher, J. W., & Cowan, J. J. 1984, ApJL, 285, L35

Perucho, M. 2020, MNRAS, 494, L22

Perucho, M., Martí, J. M., Cela, J. M., et al. 2010, A&A, 519, A41

Perucho, M., Martí, J.-M., Quilis, V., & Ricciardelli, E. 2014, MNRAS, 445, 1462

Pessah, M. E. 2010, ApJ, 716, 1012

Pizzo, R. F., de Bruyn, A. G., Bernardi, G., & Brentjens, M. A. 2011, A&A, 525, A104

Porter, D. H., Jones, T. W., & Ryu, D. 2015, ApJ, 810, 93

Pouquet, A., Frisch, U., & Leorat, J. 1976, Journal of Fluid Mech, 77, 321

Powell, K. G. 1994, Approximate Riemann solver for magnetohydrodynamics (that works in more than one dimension), Tech. Rep. NASA-CR-194902, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton, VA 23681-0001

Powell, K. G., Roe, P. L., Linde, T. J., Gombosi, T. I., & De Zeeuw, D. L. 1999, Journal of Computational Physics, 154, 284

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical recipes in FORTRAN. The art of scientific computing (Cambridge University Press)

Pringle, J. E. 1981, ARA&A, 19, 137

Qamar, S., & Warnecke, G. 2005, Journal of Computational Physics, 205, 182

Rafferty, D. A., McNamara, B. R., Nulsen, P. E. J., & Wise, M. W. 2006, ApJ, 652, 216

Raja, R., Rahaman, M., Datta, A., & Smirnov, O. M. 2023, arXiv e-prints, arXiv:2309.14244

Ramatsoku, M., Murgia, M., Vacca, V., et al. 2020, A&A, 636, L1

Rees, M. J. 1966, Nature, 211, 468

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

Ricarte, A., Narayan, R., & Curd, B. 2023, ApJL, 954, L22

Ricci, C., Walter, R., Courvoisier, T. J. L., & Paltani, S. 2011, A&A, 532, A102

Rincon, F., Ogilvie, G. I., & Proctor, M. R. E. 2007, PhRvL, 98, 254502

Roberts, G. O. 1970, Phil. Trans. Roy. Soc. London Ser. A, 266, 535

—. 1972, Phil. Trans. Roy. Soc. London Ser. A, 271, 411

Roberts, P. H., & Soward, A. M. 1992, Annu. Rev. Fluid Mech., 24, 459

Rodman, P. 2023a, Evolution of the Magnetic Field in High- and Low-Beta Disks with Initially Toroidal Fields - Simulation code, doi:10.5281/zenodo.10061386

—. 2023b, paytonrodman/athena-analysis: Publication v1.0, doi:10.5281/zenodo.10061024

Rodman, P. E., Turner, R. J., Shabala, S. S., et al. 2019, MNRAS, 482, 5625

Rodman, P. E. 2018, Honours Thesis, University of Tasmania

Roettiger, K., Burns, J., & Loken, C. 1993, ApJL, 407, L53

Roettiger, K., Stone, J. M., & Mushotzky, R. F. 1998, ApJ, 493, 62

Rüdiger, G., & Kichatinov, L. L. 1993, A&A, 269, 581

Rudnick, L., Brüggen, M., Brunetti, G., et al. 2022, ApJ, 935, 168

Rybicki, G. B., & Lightman, A. P. 1979, Radiative processes in astrophysics (WILEY-VCH Verlag GmbH 62 Co. KGaA, Weinheim)

Sakelliou, I., Hardcastle, M. J., & Jetha, N. N. 2008, MNRAS, 384, 87

Sakelliou, I., & Merrifield, M. R. 1999, MNRAS, 305, 417

Salvesen, G., Simon, J. B., Armitage, P. J., & Begelman, M. C. 2016, MNRAS, 457, 857

Sanders, J. S., & Fabian, A. C. 2007, MNRAS, 381, 1381

Sano, T., Inutsuka, S.-i., Turner, N. J., & Stone, J. M. 2004, ApJ, 605, 321

Schmidt, M. 1963, Nature, 197, 1040

Schrinner, M., Rädler, K.-H., Schmitt, D., Rheinhardt, M., & Christensen, U. R. 2007, Geophysical and Astrophysical Fluid Dynamics, 101, 81

Seyfert, C. K. 1943, ApJ, 97, 28

Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 24, 337

Shimwell, T. W., Luckin, J., Brüggen, M., et al. 2016, MNRAS, 459, 277

Shukurov, A., Sokoloff, D., Subramanian, K., & Brandenburg, A. 2006, A&A, 448, L33

Silk, J. 2013, ApJ, 772, 112

Silk, J., & Norman, C. 2009, ApJ, 700, 262

Simonetti, J. H., Cordes, J. M., & Spangler, S. R. 1984, ApJ, 284, 126

Soker, N., O'Dea, C. P., & Sarazin, C. L. 1988, ApJ, 327, 627

Sorathia, K. A., Reynolds, C. S., & Armitage, P. J. 2010, ApJ, 712, 1241

Sorathia, K. A., Reynolds, C. S., Stone, J. M., & Beckwith, K. 2012, ApJ, 749, 189

Steenbeck, M., Krause, F., & Rädler, K. H. 1966, Zeitschrift Naturforschung Teil A, 21, 369

Stimpson, M., Hardcastle, M. J., & Krause, M. G. H. 2023, MNRAS, 526, 3421

Stix, M. 1976, A&A, 47, 243

Stokes, G. G. 1851, Transactions of the Cambridge Philosophical Society, 9, 399

Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, ApJS, 178, 137

Stone, J. M., Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1996, ApJ, 463, 656

Stone, J. M., Ostriker, E. C., & Gammie, C. F. 1998, ApJL, 508, L99

Stone, J. M., Tomida, K., White, C. J., & Felker, K. G. 2020, ApJS, 249, 4

Sur, S., Shukurov, A., & Subramanian, K. 2007, MNRAS, 377, 874

Swain, M. R., Bridle, A. H., & Baum, S. A. 1998, ApJL, 507, L29

Synge, J. L. 1957, The relativistic gas (North-Holland Publishing Company)

Taylor, A. R., Stil, J. M., & Sunstrum, C. 2009, ApJ, 702, 1230

Tchekhovskoy, A., McKinney, J. C., & Narayan, R. 2012, in Journal of Physics Conference Series, Vol. 372, Journal of Physics Conference Series, 012040

Tchekhovskoy, A., Narayan, R., & McKinney, J. C. 2011, MNRAS, 418, L79

Tobias, S. M. 2021, Journal of Fluid Mechanics, 912, P1

Tout, C. A., & Pringle, J. E. 1992, MNRAS, 259, 604

Tribble, P. C. 1991, MNRAS, 253, 147

Uhlenbeck, G. E., & Ornstein, L. S. 1930, Phys. Rev., 36, 823

Urry, C. M., & Padovani, P. 1995, PASP, 107, 803

Vainshtein, S. I., & Cattaneo, F. 1992, ApJ, 393, 165

Vainshtein, S. I., & Zel'dovich, Y. B. 1972, Usp. Fiz. Nauk, 106, 431, [English translation: Sov. Phys. Usp. (1972), 15, 159-172]

van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2

van Weeren, R. J., Shimwell, T. W., Botteon, A., et al. 2021, A&A, 651, A115

Vanderwoude, S., West, J. L., Gaensler, B. M., et al. 2024, AJ, 167, 226

Vasudevan, R. V., Fabian, A. C., Reynolds, C. S., et al. 2016, MNRAS, 458, 2012

Velikhov, E. P. 1959, ZhETF, 36, 1398, [English translation: Sov. Phys.-JETP (1959), 9, 995]

Villar-Martín, M., Arribas, S., Emonts, B., et al. 2016, MNRAS, 460, 130

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261

Visbal, E., & Haiman, Z. 2018, ApJL, 865, L9

Vogt, C., & Enßlin, T. A. 2005, A&A, 434, 67

Wardle, J. 2018, Galaxies, 6, 5

Waskom, M. L. 2021, Journal of Open Source Software, 6, 3021

White, C. J., Quataert, E., & Gammie, C. F. 2020, ApJ, 891, 63

Wilber, A., Brüggen, M., Bonafede, A., et al. 2019, A&A, 622, A25

Xie, W., Lei, W.-H., Zou, Y.-C., et al. 2012, Research in Astronomy and Astrophysics, 12, 817

Yates-Jones, P. M., Shabala, S. S., & Krause, M. G. H. 2021, MNRAS, 508, 5239

Yoshimura, H. 1975, ApJ, 201, 740

Yusef-Zadeh, F., Arendt, R. G., & Wardle, M. 2022, ApJL, 939, L21

Zavala, R. T., & Taylor, G. B. 2004, ApJ, 612, 749

Zel'dovich, Y. B. 1957, ZhETF, 31, 154, [English translation: Sov. Phys.-JETP (1957), 4, 460]

Zenati, Y., & Vishniac, E. T. 2023, ApJ, 948, 11

Zhuravleva, I., Churazov, E., Schekochihin, A. A., et al. 2014, Nature, 515, 85

Ziegler, U., & Rüdiger, G. 2001, A&A, 378, 668

Zinn, P. C., Middelberg, E., Norris, R. P., & Dettmar, R. J. 2013, ApJ, 774, 66

ZuHone, J., Ehlert, K., Weinberger, R., & Pfrommer, C. 2021a, Galaxies, 9, 91

ZuHone, J. A., Markevitch, M., Weinberger, R., Nulsen, P., & Ehlert, K. 2021b, ApJ, 914, 73