

Hydrodynamic and Magnetohydrodynamic Convection in Accretion Disks



Loren E. Held
Darwin College
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Abstract

The prevalence and consequences of convection perpendicular to the plane of accretion disks have been discussed for several decades. Recent simulations combining convection and the magnetorotational instability have given fresh impetus to the debate, as the interplay of the two processes can enhance angular momentum transport, at least in the optically thick outburst stage of dwarf novae. In this thesis we seek to isolate and understand the most generic features of disk convection, and so undertake its study in both hydrodynamical and magnetohydrodynamical models.

In the first part of this thesis we investigate hydrodynamic convection in disks. First, we investigate the linear phase of the instability, obtaining estimates of the growth rates both semi-analytically, using one-dimensional spectral computations, as well as analytically, using WKBJ methods. Next we perform three-dimensional, vertically stratified, shearing box simulations with the conservative, finite-volume code PLUTO, both with and without explicit diffusion coefficients. We find that hydrodynamic convection can, in general, drive outward angular momentum transport, a result that we confirm with ATHENA, an alternative finite-volume code. Moreover, we establish that the sign of the angular momentum flux is sensitive to the diffusivity of the numerical scheme. Finally, in sustained convection, whereby the system is continuously forced to an unstable state, we observe the formation of various coherent structures, including large-scale and oscillatory convective cells, zonal flows, and small vortices.

In the second part of this thesis we investigate magnetohydrodynamic convection in disks. First, we perform inviscid, three-dimensional, vertically

stratified zero-net-flux MHD shearing box simulations with the conservative, finite-volume code PLUTO without explicit cooling. We find that MRI heating and weak cooling facilitated through advection of material across the vertical boundaries alone results in a convectively stable disk structure. Next we explore the interaction between convection and the MRI in controlled numerical experiments in which we employ an explicit height-dependent cooling prescription and explicit uniform resistivity. We find two characteristic outcomes of the interaction between the two instabilities: MRI-dominated and MRI/convective cycles. In particular we find that MRI/convective cycles lead to alternating phases of convection-dominated quiescence (during which the turbulent transport is weak) and MRI-dominated outbursts. During these outbursts angular momentum transport is enhanced by nearly an order of magnitude. Thus convection and the MRI do not generally interact in an additive manner, though they can certainly interact in non-trivial ways. In addition we find that convection in the non-linear phase takes the form of large-scale and oscillatory convective cells, reproducing a key result of our hydrodynamic investigations in a self-consistent manner, and demonstrating that these structures are a generic feature of turbulent convection in astrophysical disks.

In the final part of this thesis (which is independent of the first two) we investigate the stress-pressure relationship in disks. The stresses accompanying MRI turbulence are related to the pressure in the disk, and have been shown to increase and decrease with the pressure. We examine the time lag associated with this dependence and discuss its implications for thermal instability.

Preface

This thesis describes work done between October 2016 and March 2020 in the Astrophysical Fluid Dynamics & Nonlinear Patterns research group at the Department of Applied Mathematics & Theoretical Physics, University of Cambridge, under the supervision of Dr. Henrik Latter.

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 before for any degree or other qualification except as declared in the preface and specified in the text.

Cambridge, June 2020

Acknowledgements

First and foremost I'd like to thank my supervisor, Dr. Henrik Latter. His attention to detail, good nature, and encouragement were constants throughout my PhD, and his advice and guidance on pretty much everything from research to postdoc applications has been invaluable. Most importantly, while I still have much to learn, I have learned a great deal over the course of my PhD thanks to him. I also want to thank my parents, whose dedication to my education has been unwavering throughout my life, from the time and effort they put into homeschooling to the encouragement and support they have offered throughout the course of my university studies. It is to them that I dedicate this thesis. Finally, I'd like to thank Milena, whose patience and moral support made the process of writing up much easier than it would have been otherwise.

Contents

1	Introduction	9
1.1	Accretion disks	9
1.2	Background on cataclysmic variables	14
1.2.1	Classification of cataclysmic variables	14
1.2.2	Structure of non-magnetic cataclysmic variables	17
1.2.3	Observations of non-magnetic CVs	19
1.2.4	Theoretical predictions: disk instability model	29
1.3	Convection	32
1.3.1	Different types of convection	34
1.3.2	Convection in nature	35
1.3.3	Linear theory of convection	44
1.4	Applications of convection in disks	52
1.4.1	Convection in dwarf novae	52
1.4.2	Convection in LMXBs and AGN	53
1.4.3	Convection in protoplanetary disks	55
1.5	Dynamics of convection in disks	56
1.5.1	Hydrodynamic convection in disks	56
1.5.2	MHD convection in disks	58
1.6	Outline of thesis	60
2	Methods	62
2.1	Governing equations	62
2.2	Shearing box formalism	65
2.2.1	Tidal expansion of effective gravitational potential	65

2.2.2	Conservation of energy in the shearing box	69
2.3	Numerical Algorithms	71
2.3.1	Finite-volume methods	71
2.3.2	Conservative form of the equations	73
2.3.3	Discretization of governing equations	74
2.3.4	Reconstruction and interpolation methods	75
2.3.5	Riemann solvers	77
2.3.6	Time-stepping algorithms	82
2.4	Codes	84
2.5	Diagnostics	86
3	Hydrodynamic convection in disks	92
3.1	Numerical set-up	93
3.1.1	Initial conditions and units	93
3.1.2	Box size and resolution	99
3.1.3	Boundary conditions	99
3.1.4	Important parameters and instability criteria	99
3.2	Linear theory	100
3.2.1	Spectral methods	101
3.2.2	WKB approach	105
3.2.3	Eigenvalues	109
3.2.4	Vertical structure of the eigenfunctions	111
3.2.5	Comparison of theory with simulations	114
3.3	Simulations of unforced compressible convection	116
3.3.1	Fiducial simulations	117
3.3.2	Simulation of hydro convection in ATHENA	123
3.3.3	Sensitivity of sign of angular momentum transport to numerical scheme	126
3.3.4	Viscous simulations	128
3.4	Structures in forced compressible convection	131
3.4.1	Set-up	132
3.4.2	Large-scale oscillatory cells	134
3.4.3	Zonal Flows	136

3.4.4	Vortices	138
3.4.5	Discussion	140
3.5	Conclusions	142
4	MHD convection in disks	145
4.1	Parameters and set-up	146
4.1.1	Numerical set-up	146
4.1.2	Criteria for convective instability in a viscous fluid . . .	148
4.2	Stratified MRI simulations without explicit cooling	151
4.2.1	Motivation	151
4.2.2	Fiducial simulation	152
4.3	Stratified MRI simulations with cooling	159
4.3.1	Inviscid simulation with uniform cooling	159
4.3.2	Resistive simulations with cooling above $ z > 0.75H_0$.	162
4.3.3	MRI/convective cycles	164
4.3.4	MRI-dominated	173
4.3.5	Parameter survey	180
4.4	Discussion	186
4.4.1	Effect of explicit resistivity	186
4.4.2	Is convection seeding net-vertical-flux MRI?	189
4.4.3	Dependence on box size	191
4.4.4	Applications to dwarf novae	191
4.5	Conclusions	196
5	The stress-pressure relationship	199
5.1	Applications	199
5.2	Set-up	203
5.2.1	Cooling Prescription	204
5.3	Results	205
5.3.1	Heating runs	205
5.3.2	Runs with cooling	208
5.4	Discussion and Conclusions	216

6	Conclusions	219
6.1	Future work	224
A	Numerical issues and convergence studies for MHD	227
A.1	Numerical issues with $\nabla \cdot \mathbf{B} = 0$ in PLUTO	227
A.2	MHD convergence studies	230
A.2.1	Dependence on resolution	230
A.2.2	Dependence on radial box size	231
B	Tables of hydro simulations	234
C	Tables of MHD simulations	238
	Bibliography	242

Chapter 1

Introduction

1.1 Accretion disks

An accretion disk is an astronomical object consisting of a disk of gas, dust, or plasma in differential rotation about a central gravitating source, such as a black hole, a white dwarf, a neutron star, or a young star. If a mechanism exists for transporting most of the angular momentum outwards (away from the central accretor), most of the mass of the disk will spiral inwards (we say it is *accreted*). Such disks have been observed in a wide range of contexts in astrophysics: around young stars (termed *protoplanetary disks*, or PPDs for short), in close binary systems whereby one of the stars in the binary accretes gas from its companion (termed *cataclysmic variables* [CVs] and *X-ray binaries* [XRBs]), and also around the supermassive black holes at the centers of galaxies (termed *active galactic nuclei* [AGN]).

Protoplanetary disks typically have radii in excess of 10^2 astronomical units (AU), and have been directly imaged (Burrows et al., 1996; Brogan et al., 2015). AGN have gaseous disks in their inner regions that have similar sizes to PPDs reaching up to 10^3 AU for M87 (EHT Collaboration, 2019), but can extend up to tens of parsecs ($\sim 10^7$ AU) if their outer dusty tori are included (Ford et al., 1994; Jaffe et al., 1996). Disks in binary systems are much smaller, typically 5×10^{-3} AU, or comparable to the solar radius (which is $\sim 4.6 \times 10^{-3}$ AU or around 7×10^{10} cm), and are too small to be

imaged directly. Nevertheless, there is strong indirect observational evidence for disks in binary systems both from their spectra and from variations in their light curves. These observations will be discussed in more detail in Section 1.2.3.

Galactic disks (such as the Milky Way), planetary rings (such as Saturn’s rings), and accretion disks are three disk-like classes of object, that seem similar at first glance but are quite different in terms of their accretion rates, composition, dynamics (especially rotation law), and ratio of radial and vertical sizes. The most important distinguishing factor between these classes is the accretion rate. Planetary rings accrete on timescales much longer than do gaseous disks: a rough estimate by Esposito (1986) yields a lifetime of 10^{12} years for Saturn’s rings, compared to typical PPD lifetimes of several millions of years (Mamajek, 2009). Galactic disks do not really ‘accrete’ in a comparable manner, except in the inner regions of AGN. Galactic disks also differ from accretion disks primarily in that the self-gravity of their baryonic components (stars and gas) completely determines the rotation law on scales much less than that of the galaxy,¹ while the rotation law of the galaxy on scales comparable to the size of the galaxy is thought to be determined by its dark matter content. In contrast, in accretion disks, which are composed primarily of gas, it is the gravity of the central object that dominates the rotation law of gas in the disk rather than the disk’s self-gravity (an exception is protoplanetary disks in their early stages of formation, which are known to be self-gravitating). At the other end of the spectrum, planetary rings consist mostly of chunks of rock and ice, and their evolution is driven primarily by inelastic collisions between their components. Finally, planetary rings are typically much thinner than gaseous accretion disks, with a typical ratio of characteristic scale height perpendicular to the plane of the disk H to cylindrical radial size r of $H/r \sim 10^{-7}$ compared to $H/r \sim 10^{-3} - 10^{-2}$ for gaseous disks, and $H/r \sim 10^{-2} - 0.1$ for the stellar component of galactic

¹As an example, the ratio of the magnitude of the gravitational force of the Sun on the Earth to the gravitational force of Sgr A* on the Earth is given by $F_{\odot}/F_{A^*} = (M_{\odot}/M_{A^*})(d_{A^*}/d_{\odot})^2 \sim 10^{11}$, where d_{\odot} is the Earth-Sun distance, and d_{A^*} is the distance of the Solar System from the galactic center.

disks (Latter et al., 2017). Accretion disks, galaxies, and planetary rings are typically treated separately in the literature, and in this thesis we restrict our focus to accretion disks, only.

Accretion disks are common in space. Disks have been deduced from their infrared excess in 60% to 80% of stars observed in nearby ($\lesssim 1400$ ly) clusters younger than 1 Myr (Mamajek, 2009; Williams & Cieza, 2011). About one-third of stars in the Milky Way are in binary or multiple systems (Lada, 2006), and Moe et al. (2019) report that the fraction of close (interacting) binaries from various surveys varies from $10\% \pm 3\%$ to $53\% \pm 12\%$ depending on the metallicity (i.e. the ratio of the abundance of iron to hydrogen). Some fraction of these interacting binaries will have disks. Finally Chiang et al. (2019) have estimated an AGN fraction of about 37% in a sample of nearly 16.5 thousand galaxies. The prevalence with which disks occur in space can be understood through the relatively simple physical ingredients needed to form a disk: gravity, rotation, and dissipation. These basic ingredients admit a variety of formation mechanisms for disks, from the gravitational collapse of a cooling rotating molecular cloud to form a torus of dust and gas (the precursor to a protoplanetary disk), to the self-intersection and viscous spreading of a stream of gas transferred from a secondary to a primary in a close binary to form cataclysmic variable and X-ray binary disks, to the accretion of gas onto supermassive black holes to form AGN.

If the disk is sufficiently cool that the radial pressure gradient is small, the disk is said to be *thin*; quantitatively, $H/r \ll 1$. In this case the dominant force balance in the cylindrical radial direction is given between the radial component of gravity and the centrifugal force. In the vertical direction (i.e. perpendicular to the plane of the disk) the dominant balance is between the vertical component of gravity and the vertical pressure gradient. Finally, the disk is in thermal equilibrium if there is a balance between heating and radiative cooling.

In order for a disk to accrete there must be a residual inward drift of material in addition to the dominant orbital motion. This can be facilitated by the viscous shear exerted between neighboring disk annuli orbiting with different periods. A faster rotating ‘inner’ annulus will drag forward

its neighboring, slower moving, ‘outer’ annulus, and vice versa, the outer annulus will retard the motion of the inner annulus. In the context of a disk where the orbiting fluid elements are sitting in the gravitational potential of a point source this will result in an increase in the orbital energy of the outer annulus (i.e. it will move outwards), whereas the orbital energy of the inner annulus will decrease (i.e. it will move inwards), and so the annulus will spread. Furthermore since the viscous shear removes net orbital energy from the annulus, the net effect of this process is a transfer of mass inwards and a transfer of angular momentum outwards. In terms of the thermodynamics, the viscous shear is responsible for extracting orbital and gravitational potential energy in the disk which is subsequently converted into thermal energy (heat), and eventually carried away by radiation. Thus thermalization of orbital energy in the disk contributes to the disk’s luminosity, leading to observational signatures, which we discuss in greater detail in Section 1.2.3).

The precise mechanism facilitating outward transfer of angular momentum in astrophysical disks remains an outstanding problem. It is known that ordinary molecular viscosity results in accretion time-scales much longer than the observed lifetimes of disks.² Thus, angular momentum must be removed through other processes. Hypotheses include (a) enhanced viscosity (i.e. diffusion of momentum) due to turbulence induced by some instability, (b) mass outflow through winds and jets, and (c) tidally induced waves (Ju et al., 2016; Kurbatov et al., 2014). Not all disks exhibit jets or are subject to strong tidal forces, however, therefore the focus for the last forty years has been on enhanced viscosity due to turbulence ν_t as providing the dominant mechanism for accretion. The simplest one-dimensional model of a disk is often expressed through the viscous-diffusion equation, which governs the time-evolution of an azimuthally and vertically-averaged annulus of surface

²To see this explicitly, consider the example given in Lin & Papaloizou (1985) for a protoplanetary disk. Taking the (cylindrical) radial size to be $r \sim 10^{15}$ cm (~ 70 AU), and estimating the molecular viscosity from kinetic theory to be $\nu \sim 10^6$ cm² s⁻¹, the viscous diffusion timescale is estimated to be $\tau \sim r^2/\nu \sim 10^{16}$ years. This is ten orders of magnitude greater than protoplanetary disk lifetimes of $\sim 10^6$ years inferred from observations (Mamajek, 2009).

density Σ and is given by

$$\frac{\partial \Sigma}{\partial t} = -\frac{1}{r} \frac{\partial}{\partial r} \left[\left(\frac{d(r^2 \Omega)}{dr} \right)^{-1} \frac{\partial}{\partial r} \left(\nu_t \Sigma r^3 \frac{d\Omega}{dr} \right) \right], \quad (1.1)$$

where $\Omega(r)$ is the angular frequency of the disk, which is a function the cylindrical radius r . For a Keplerian disk, for which the angular frequency is given by $\Omega = \Omega_K = \sqrt{GM}/r^{3/2}$, Equation 1.1 reduces to

$$\frac{\partial \Sigma}{\partial t} = \frac{3}{r} \frac{\partial}{\partial r} \left[r^{1/2} \frac{\partial}{\partial r} (\nu_t \Sigma r^{1/2}) \right]. \quad (1.2)$$

In early theoretical models of disks the turbulent viscosity ν_t was parametrized through a dimensionless parameter α through $\nu_t = \alpha c_s H$, where c_s is the sound speed of the gas and H is the disk vertical scale height (Shakura & Sunyaev, 1973). Later it was shown that the *magnetorotational instability* (MRI), an instability associated with the interplay between magnetic fields and the rotation of the disk, could lead to sustained turbulence and angular momentum transport in ionized disks (Balbus & Hawley, 1991). Nevertheless, the α constrained from observations of dwarf novae during outburst ($\alpha \sim 0.1$) remains an order of magnitude greater than that measured in simulations of MRI turbulence with zero-net magnetic flux ($\alpha \sim 0.01$) (King et al., 2007).

More recently, simulations of MRI turbulence in stratified disks have shown that the MRI is capable of generating a convectively unstable entropy gradient. In particular, there has been some indication that an interplay between convection and the magnetorotational instability might enhance angular momentum transport (Bodo et al., 2013b; Hirose et al., 2014). In addition convection transports heat from the disk mid-plane to the surface, and might therefore be important in determining the quasi-steady vertical structure of a disk.

The key aim of this thesis is to study internal flows and turbulence in accretion disks, and how these contribute to the transport of angular momentum and heat. Thus we focus our investigations on convection, the magnetorotational instability, and the interplay between the two. In Sec-

tion 1.2 we discuss observations of disks, focusing on cataclysmic variables, and on dwarf novae in particular, as these systems are the most promising ones for exhibiting convective and MRI turbulence. A general introduction to convection, including various applications to both geophysics and astrophysics, is presented in Section 1.3. Applications of convection to different disk types (dwarf novae, X-ray binaries and AGN, and protoplanetary disks) are discussed separately in Section 1.4. In Section 1.5 we review the existing literature studying the dynamics of both hydrodynamic and magnetohydrodynamic convection in accretion disks. Finally in Section 1.6 we present an outline of the remainder of the thesis.

1.2 Background on cataclysmic variables

While convection has been discussed in the literature in the context of nearly all types of disks (see Section 1.4), the most promising application is to a subset of cataclysmic variables known as the dwarf novae. It is important to emphasize, however, that there is currently no direct observational evidence for convection in dwarf novae,³ as there is for, for example, of convection in the Sun (see Figure 1.9). Nevertheless there is compelling theoretical and numerical evidence that convection should be present in dwarf novae (at least during outburst), and this theory in turn has been motivated by observations. Thus it is important to know the basic observational properties of these systems, which are reviewed below.

1.2.1 Classification of cataclysmic variables

Accretion disks can form in binary systems in which a white dwarf primary and a secondary star interact via *mass transfer* of material from the secondary to the primary. When the white dwarf is co-orbiting with a *low mass* secondary (i.e. $M_2 < M_1$, where M_1 denotes the mass of the white dwarf) that is on or near the main sequence, such that there is mass transfer from the secondary to the white dwarf through *Roche lobe overflow* (i.e. when the

³Nor of the MRI, for that matter.

secondary fills its critical equipotential surface), the system is called a *cataclysmic variable* (CV) (Knigge, 2011). About 1600 CVs are known, of which around 318 have a known orbital period (Hellier, 2001). These typically exhibit one or more outbursts in their light curves, hence their classification as a ‘cataclysmic’ event. If, on the other hand, the white dwarf is co-orbiting with a *high mass* evolved secondary (typically a red giant) such that mass transfer from the secondary to the white dwarf occurs primarily via a *stellar wind* (though there might also be some additional transfer of mass via Roche lobe overflow), the system is termed a *symbiotic star*.⁴

CVs can be further subdivided based on the strength of the magnetic field B_1 of the white dwarf primary. This field can have appreciable effects on the transfer of mass from the secondary, for example by truncating the resultant accretion disk, or precluding the formation of a disk altogether. There are *non-magnetic* CVs (for which $B_1 \lesssim 10^5$ G or $B_1 \lesssim 10$ T) and *magnetic* CVs (for which $B_1 \sim 10^5 - 10^8$ G or $B_1 \sim 10 - 10^4$ T). Magnetic CVs are further subdivided into *polars* (an archetype of which is AM Her) in which there is no disk, the material instead being accreted along the white dwarf’s magnetic field lines (a phenomenon known as ‘curtain accretion’), and *intermediate polars* (an example of which is DQ Her) in which the white dwarf is surrounded by an accretion disk. Note however, that non-magnetic CVs are, strictly speaking, still magnetic: if sufficiently ionized, the disk in a non-magnetic CV can interact with the external magnetic field of the primary, as well as with its own (internal) field generated by fluid motions in the disk through a dynamo process.

Non-magnetic CVs are subdivided into four broad classes based on their observational characteristics. The *classical novae* (CN) are CVs in which only one outburst has been observed, and during which the system brightens by ~ 12 mag over about three days. The outburst is thought to be due to accretion onto the surface of the white dwarf, which eventually leads to surface material undergoing thermonuclear runaway. *Recurring novae* (RN)

⁴Note that some symbiotic stars, too, are observed to undergo both brightness variations (roughly ~ 1 mag over several years) as well as long-lasting outbursts (roughly 4 – 5 mag over several years) (Kenyon, 1990).

are classical novae whose outburst has been observed to repeat at least once. Of the 10 RN known in our galaxy, these have been observed to brighten by ~ 8 -9 mag during outburst, with recurrence timescales ranging from decades to nearly 100 years (Schaefer, 2010). Then there are the *dwarf novae* (DN), CVs that are observed to undergo semi-regular outbursts, i.e. the outbursts (during which the system is bright) are separated by periods of quiescence (during which the system is dim compared to the outbursts). The periods of outburst and of quiescence are also referred to as the ‘high state’ and ‘low state’, respectively. Although the outburst magnitude, rise time, duration, and intervals between outbursts vary between systems (and even within the same system), dwarf novae are typically observed to brighten by 2-6 orders of magnitude over 1-2 days during outburst, and subsequently decay to quiescence over several days to a week (Knigge, 2011). Intervals between outburst typically last weeks to several months (or even years, in some cases). The outbursts themselves are thought to originate in an instability in the disk surrounding the white dwarf. This instability is described by the *disk instability model* (DIM), which we will discuss in greater detail in Section 1.2.4. Finally, there are the *nova-likes* (NLs), which are CVs that are *non-eruptive*. They are similar to dwarf novae, but appear to be permanently stuck in the high state.

Dwarf novae (of which about 900 are known⁵) are themselves generally divided into three classes based on their light-curves, each named after an actual prototypical system. Dwarf novae that exhibit standstills, i.e. that undergo a series of outbursts before becoming stuck in, or just below, the high state for several months before undergoing repeated outbursts again, are known as Z Cam type (about 20 are known). Dwarf novae whose light curves occasionally exhibit powerful outbursts known as *superoutbursts* are classified as SU UMa type (of which around 650 have been found). Superoutbursts tend to be brighter than regular outbursts by about 0.5-1 mag and last for around 14 days (compared to durations of 1-2 days for regular outbursts).

⁵Numbers for known dwarf novae and DN subtypes have been obtained from the catalogue by Ritter & Kolb (2003) (updated in 2016). Note that actual numbers vary based on catalogue, new observations, and reclassification of known systems.

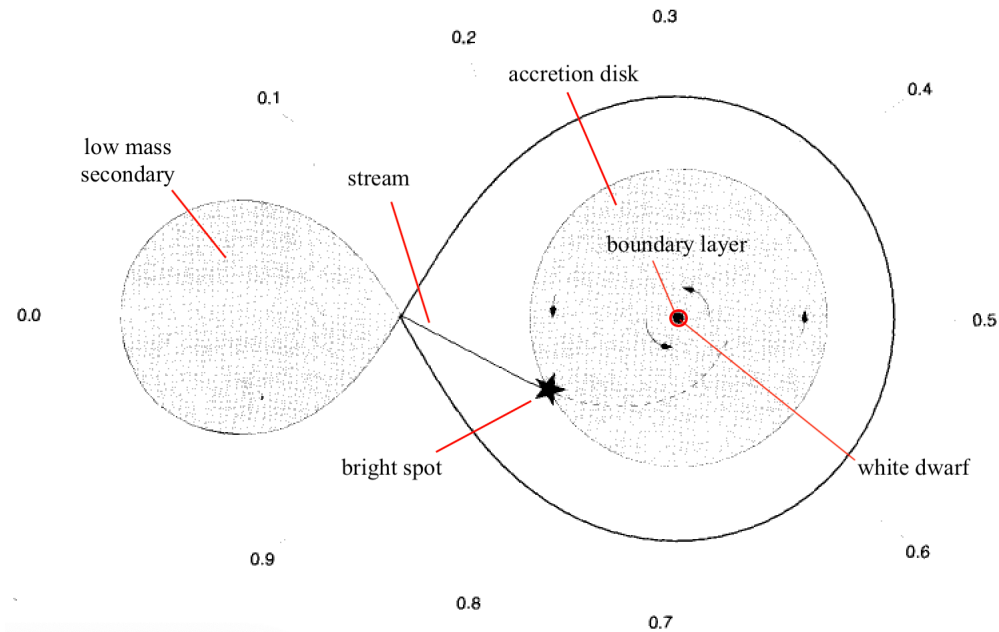


Figure 1.1: Schematic showing top down view of a typical cataclysmic variable. The numbers surrounding the figure indicate the orbital phase at which an observer viewing the binary edge-on sees a given portion of the binary. Adapted from Hellier (2001).

Finally, all other dwarf novae are classified as U Gem type (about 85 are known). A less common subclass are the AM CVn (Solheim, 2010), which are dwarf novae in which a helium rich disk is observed (about 25 are known).

1.2.2 Structure of non-magnetic cataclysmic variables

The overall structure of a typical non-magnetic CV is shown in Figure 1.1. There are six distinct components that contribute to the emission to various degrees. In high-inclination (i.e. eclipsing) systems the contribution of any individual component to the total emission varies with the orbital phase, a fact that can be used to deduce a wealth of information about the individual

components (see Section 1.2.3).⁶

The *white dwarf* primary is a stellar remnant which is supported from gravitational collapse by electron degeneracy pressure. In CVs the primary typically has a mass between $0.3M_{\odot}$ and $1.3M_{\odot}$, a radius of $R_1 \sim 10^{-3}R_{\odot}$ ($\sim 10^8$ cm), and a surface temperature of $T_{\text{eff},1} \sim (3.4\text{--}4) \times 10^4$ K. It has a near-blackbody continuum spectrum with emission peaking in the UV. Balmer lines (which are significantly pressure-broadened due to the large surface gravity of the white dwarf) signify the presence of hydrogen atoms at the surface.

The primary is surrounded by an *accretion disk*, which typically has an outer radius of $R_d \sim (0.4\text{--}0.5)a$ (around 5×10^{-3} AU) in outburst, where a is the orbital separation of the binary (Warner, 1995). The disk is somewhat smaller in quiescence ($\sim 0.3a$), which constitutes an important piece of observational evidence that the evolution of the disk in outburst is driven by viscous accretion as hypothesized by the disk instability model (see Section 1.2.4). The surface temperature of the disk decreases with radius: during outburst it varies from $T_{\text{eff},d} \sim (3\text{--}4) \times 10^4$ K in the innermost regions of the disk to $T_{\text{eff},d} \sim 5 \times 10^3$ K at the outer edge (La Dous, 1989). Finally, a narrow, hot *boundary layer* connects the disk to the surface of the white dwarf. Temperatures in the boundary layer reach as high as $T_{\text{BL}} \sim 10^5$ K, which is sufficiently hot for the boundary layer to radiate in the X-rays.

The *secondary* (typically a fully-convective red dwarf, or low mass M-type main sequence star undergoing hydrogen fusion in its core) has a mass of around $M_2 \sim 0.12M_{\odot}$,⁷ a radius of $R_2 \sim 0.15R_{\odot}$, and a surface temperature of around $T_{\text{eff},2} \sim 2900$ K, with its emission peaking in the IR (Hellier, 2001). Irradiation of the secondary by the white dwarf and the boundary layer can heat the white dwarf facing side of the secondary to around ~ 7500 K, however. Material is transferred from the secondary to the disk via Roche

⁶We define the inclination of the disk as the angle it makes with the plane of the sky relative to an observer on Earth. If the disk is viewed edge-on (‘high inclination’) its inclination is $i = 90^\circ$; if it is viewed face-on (‘low inclination’) its inclination is $i = 0^\circ$.

⁷Note that, according to some classification systems, the secondary in CVs can have a mass comparable to that of the white dwarf itself (Knigge, 2011).

lobe overflow in a *stream*. The stream impacts the disk at the *bright spot* (or hot spot), which has a temperature of around $\sim 1.3 \times 10^4$ K.

1.2.3 Observations of non-magnetic CVs

In this section we wish to address three questions: ‘What is the evidence that an accretion disk exists in the first place in cataclysmic variables?’; ‘What can be inferred about the disk from observations?’; and ‘How can disk theory be tested against observations?’.

Disks in binary systems are challenging to detect. For one, they are too small to be resolved by the current generation of telescopes. As an example, a disk of size ~ 0.005 AU ($\sim 7.48 \times 10^{10}$ cm) at a distance of around 372 ly (equivalent to the distance of SS Cyg, a well known U Gem type dwarf nova) subtends an angle of around $\theta \sim 4.38 \times 10^{-5}$ arcsec, which is four orders of magnitude smaller than the angular resolution of the Hubble Space Telescope (HST) which is $\theta_{\text{HST}} \sim 0.1$ arcsec. Conversely a disk of the same size would need to be less than ~ 0.16 ly from Earth to be resolvable by the HST, and thus significantly closer to Earth than Alpha Centauri, the nearest known star, which is located at a distance of ~ 4.37 ly from the Sun. Another limiting factor in detecting disks in binary systems is that the observer has to disentangle the disk’s emission from that of the other components in the system, i.e. the white dwarf, the secondary star, the boundary layer, and the hotspot.

Nevertheless there is compelling observational evidence that accretion disks exist in CVs from their spectra and light-curves. In eclipsing systems the disk can be inferred directly from dips in the light curves, which we discuss in more detail below. Velocity line profiles (i.e. plots of the intensity of a particular line vs. its velocity⁸) in these systems have a double-peaked structure characteristic of matter in circular motion about the primary. In high inclination systems (i.e. viewed edge-on) when the secondary eclipses

⁸For matter moving non-relativistically the shift in wavelength $\Delta\lambda$ of an emitted photon (compared to its wavelength when emitted from a fluid parcel at rest) is directly proportional to the velocity v of the emitting material along the line of sight, i.e. $\Delta\lambda/\lambda \approx v/c$.

the edge of the disk moving towards the observer, the blue-shifted component of the double-peaked line velocity profile disappears. Later, as the secondary eclipses the edge of the disk moving away from the observer, the red-shifted component of the line profile disappears and the blue-shifted component reappears. Finally the continuum spectra of CVs can be fit very well by the stretched blackbody spectrum of a disk in steady-state (see Section 1.2.3 for a more detailed discussion).

There are two tools with which to investigate CVs observationally: spectra (plots of intensity against wavelength), and light curves (plots of intensity against time). Despite the fact that CV disks are too small to be imaged directly, a wealth of information can be obtained from these systems from their spectra and light curves. From photometry this includes both long-duration light curves (taken over several months or even years) of disks at all inclinations (which reveal details about their outbursts, such as their duration, rise and decay times, and intervals between outbursts), as well as short-duration light curves (taken over several hours) in high inclination systems, which reveal the eclipse of the disk and primary by the secondary. An entire subdomain of CV photometry, known as eclipse analysis, is dedicated to extracting a large amount of information from eclipse light curves (such as the sizes of the white dwarf, the secondary, the disk, and even the hot spot), including the technique of *eclipse mapping* from which a 2D disk temperature distribution can be obtained from successive eclipse light curves. Spectroscopy yields both continuum spectra, which provide one of the key methods for testing disk theory against observations, as well as line spectra, from which the presence of a disk (through velocity line profiles), and its composition can be inferred. Furthermore, a technique known as *Doppler tomography* can be used to extract a 2D disk velocity distribution from one dimensional line velocity profiles. Each of these techniques is discussed in further detail below.




Figure 1.2: Orbital period P_{orb} distribution (in hours) of non-magnetic cataclysmic variables. The top panel shows the distribution dwarf novae (including U-Gem, Z Cam, and SU UMa subtypes), while the bottom panel shows the distribution for nova-like variables. Adapted from Warner (1995).

Orbital period distribution

In systems with sufficiently high inclination, the orbital period P_{orb} can be inferred directly from the light curve. When the secondary, which is dimmer than the disk and white dwarf primary, eclipses the disk and the white dwarf, there is a dip in the light curve. Orbital periods tend to range between 1h and 9h, and the distribution of orbital periods is bimodal with a *period gap* between $2.2\text{h} \leq P_{\text{orb}} \leq 2.8\text{h}$ (see Figure 1.2). Of the dwarf novae, the SU UMa type tend to be concentrated below this period gap, whereas U Gem and Z Cam type tend to be concentrated above the period gap. The mass ratio q of CV systems ranges from $q \sim 0.3$ for systems below the period gap to ~ 0.6 for systems above the period gap. From Kepler's third law⁹ and for

⁹Kepler's third law is given by $P_{\text{orb}}^2 = (4\pi^2 a^3)/(G M_1(1 + q))$.

a primary of mass $\sim 1M_\odot$ this implies separations a of around $a \sim 0.9R_\odot$ for systems below the period gap and a separation of $a \sim 1.2R_\odot$ above the gap.

Continuum spectra

The disk dominates emission during outburst, and continuum spectra provide one of the key methods of testing disk theory against observations. The spectral flux emitted by a disk with an inclination angle of i relative to an observer at a distance d from the disk is given by

$$F_\nu = \frac{2\pi \cos i}{d^2} \int_{R_1}^{R_d} I_\nu r dr, \quad (1.3)$$

where R_d is the outer disk radius, and I_ν is the spectral intensity (Warner, 1995). If the disk is optically thick, each annulus in the disk will radiate like a blackbody and so the spectral intensity can be approximated by Planck's law

$$I_\nu \sim B_\nu(T) = \frac{2\pi\nu^3}{c^2} (\exp[-h\nu/k_B T] - 1)^{-1}, \quad (1.4)$$

where ν is the frequency of the emitted radiation, c is the speed of light, h is Planck's constant, and k_B is Boltzmann's constant. By assuming a steady state (i.e. taking $\partial_t \Sigma = 0$) and integrating the viscous diffusion equation (cf. Equation 1.2), and by balancing viscous heating with radiative cooling, a radial temperature profile of the form

$$T(r) = \left[\frac{3GM_1 \dot{M}_d}{8\pi\sigma R_1^3} \right]^{1/4} \left(1 - \left[\frac{R_1}{r} \right]^{1/2} \right)^{1/4} \quad (1.5)$$

can be derived (see, for example, Warner (1995)), where R_1 is the radius of the white dwarf primary, σ is the Stefan-Boltzmann constant, and \dot{M}_d is the accretion rate. Substituting Equations 1.5 and 1.4 into Equation 1.3 and integrating over the radial direction yields the continuum spectrum for a disk in steady state. It has the appearance of a 'stretched blackbody' as can be seen in Figure 1.3. Most of the emission in the UV is from the disk, while the hot boundary layer (whose contributions are not shown in Figure 1.3)




Figure 1.3: Synthetic continuum spectra from successive disk annuli that are each radiating as a blackbody around a $1M_{\odot}$ white dwarf with an accretion rate of $10^{-9}M_{\odot} \text{ yr}^{-1}$. The dashed line corresponds to the integrated flux from the entire disk. From La Dous (1989).

tends to dominate emission in the X-rays. Most of the IR radiation, on the other hand, stems from the cooler outer regions of the disk, and from the secondary.

From Equation 1.5, one can see that the theoretical spectrum is parametrized by the mass and radius of the primary, M_1 and R_1 , respectively, the disk accretion rate \dot{M}_d , and the binary inclination i . Thus, given independent measurements of three of the four parameters, fitting spectra of the form given by Figure 1.3 to observations provides a method of constraining the fourth parameter.

Line spectra

An enormous amount of information about the disk can be inferred from its emission and absorption lines, including its composition, size, 2D velocity

distribution, the presence of winds, and the mass ratios and masses of the primary and secondary stars. Observed emission lines in the infrared reveal the presence of ionized hydrogen, helium, calcium, and oxygen, while in the UV emission lines for hydrogen, helium, carbon, nitrogen, oxygen, magnesium, and silicon are typically observed.¹⁰

There is some evidence for P Cygni-like profiles in certain UV lines (e.g. CIV), which suggests the existence of dipolar winds or outflows in some dwarf novae disks. These line profiles (named after the blue hypergiant, P Cygni, where they were first observed) are characterized by the feature that the same line appears both in emission and in absorption. Analysis of the blue-shifted component of these profiles suggest the radiation was emitted from gas moving with a velocity comparable to the escape velocity of the primary, thus suggesting the presence of an outflow. These winds remove angular momentum from the disk, and thus might contribute to accretion: indeed there are some arguments that wind-driven accretion can even dominate viscous (turbulent) accretion during the quiescent phase of dwarf novae (Scepi et al., 2019).

Another spectral diagnostic is that of line velocity profiles (i.e. plots of line intensity against velocity). A fluid element orbiting with Keplerian velocity $v_K(r) = r\Omega_K(r)$ and located at a position (r, θ) in the plane of a disk with inclination i will, at a given instance in time, have a velocity component $v_D = v_K(r) \sin \theta \sin i$ along the line-of-sight of the observer. For a fixed inclination angle, lines of constant projected velocity form a dipole pattern on the surface of the disk (see the upper-left-hand panel of Figure 1.4). Thus each segment of the velocity profile consists of a contribution from a certain band (in $[r, \theta]$) in the disk: material that moves perpendicular to the line of sight ($\theta = 0$) contributes to emission at the center of the line profile (region A), while the wings are dominated by contributions from the faster moving (but smaller) inner regions of the disk (region C). The truncated crescents (region B) have the greatest area and result in the double-peaked

¹⁰A more exhaustive list of emission lines observed in dwarf novae in quiescence can be found in Table 3.5 of Warner (1995).

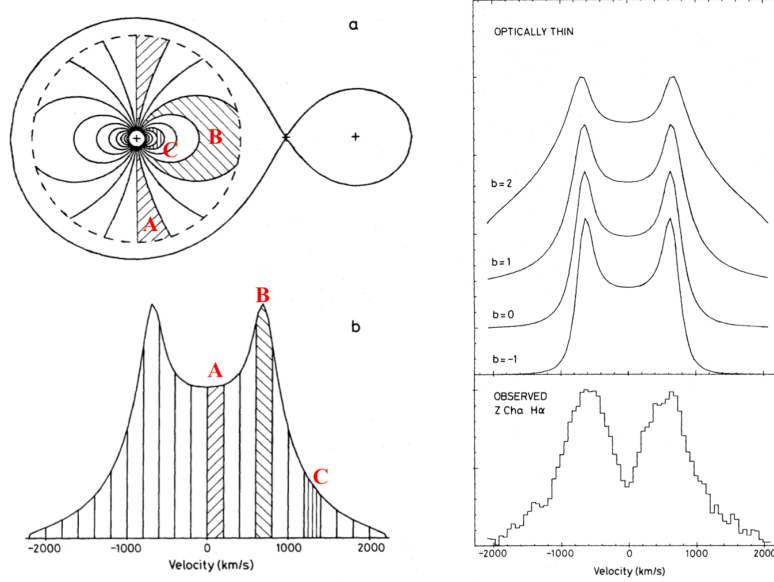


Figure 1.4: Left: schematic of line profile from an eclipsing binary system viewed at phase $\phi = 0.25$. The shaded areas at A, B, and C in the disk contribute to the local minimum, the right-hand peak, and the right-hand tail of the line profile, respectively. Right: synthetic line profiles (top), and the actual observed double-peaked $H\alpha$ line profile from the U-Gem type dwarf nova Z Cha in quiescence (bottom). Adapted from Horne & Marsh (1986).

cusps in the line profile where $v_D = v_K(R_d) \sin i$. Thus if the inclination of the disk is known, the outer disk radius R_d can be inferred directly from the cusps in the velocity profile.

Eclipse light curves

From the light-curve during eclipse (see Figure 1.5) the relative sizes of the primary, the secondary, the disk, and the bright spot can be measured directly. In addition, synthetic light curves of the eclipse can be fitted to observed eclipse light curves to deduce a wealth of parameters, including the sizes of the primary, the secondary, and the disk, the mass ratio q , and the inclination i . Such synthetic light curves were constructed for optically thick




Figure 1.5: Smoothed eclipse light curve of the SU UMa type dwarf nova Z Cha. The left-hand H above the eclipse indicates the phase at which the bright spot is eclipsed by the secondary, while the right-hand H indicates the phase at which the bright spot emerges from the eclipse. The two H's below the eclipse curve show the same but for the disk. From Hellier (2001). Original from Wood et al. (1986).

disks by Zhang et al. (1986) and Zhang & Robinson (1987), who modeled the disk temperature profile as a power law in radius. By fitting the synthetic eclipse light curves to observations of the dwarf novae U Cas and U Gem, they were able to provide an independent method of determining the radial temperature profiles of the disks, and also introduced a method for determining the uncertainty of the fitted parameters. More recently Ghoreyshi et al. (2011a,b) have constructed synthetic lights curves using their SHELLSPEC code which allows for a wider set of input parameters than those used by Zhang et al. (1986), including the effects of jets and stellar winds, and is also able to model the propagation of radiation in an optically thin medium. They applied their code to model the Algol-type¹¹ interacting binary AV Del and determined that the observational light curves were best fit by an optically thick disk.

¹¹Algol variables are binary systems in which a *main sequence* primary, rather than a white dwarf, accretes from a secondary via Roche lobe overflow.

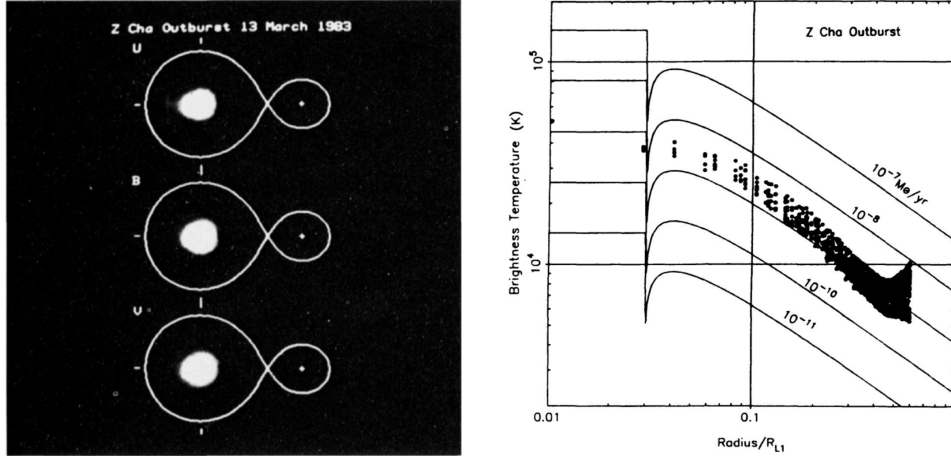


Figure 1.6: Left: surface brightness distribution of the disk in the U, B, and V bands of the SU UMa type dwarf nova Z Cha in outburst reconstructed from eclipse light curves using the technique of *eclipse mapping*. Right: comparison of observed temperature profile (black dots) with temperature profiles (solid lines) predicted by steady state disk theory. The theory is in close agreement with observations for a mass transfer rate of $10^{-9} M_{\odot} \text{ yr}^{-1}$. From Horne & Cook (1985).

An important technique related to eclipse light curves is that of *eclipse mapping*. Here 1D eclipse data is used to deduce the 2D temperature profile of the disk. As the secondary begins to eclipse the disk, successive strips of the disk are obscured from view; these strips subsequently reemerge as the eclipse progresses. Thus different strips of the disk are visible at different times during the eclipse, and different regions of the disk contribute to the overall light curve at different times. Horne & Cook (1985) developed a method for combining this data into a 2D intensity distribution. The distribution can be compared to the theoretical distribution expected for a disk in steady state, and can be used to infer the accretion rate of the disk (see Figure 1.6).




Figure 1.7: The light-curve of the U-Gem-type dwarf nova SS Cyg taken over a 10 year period showing the semi-regular outbursts characteristic of dwarf novae. Credit AAVSO.

Light curves showing outbursts

Semi-regular, repeated outbursts are the defining feature of dwarf novae, and any model attempting to explain observations of these systems needs to account for the outbursts. One such model, the disk instability model, is discussed in Section 1.2.4. Here we briefly discuss those properties of DN outbursts that have been inferred from observations.

A typical light curve for the U-Gem type dwarf nova SS Cyg is shown in Figure 1.7. From the light curve (taken over a period 10 years), it is evident that no two outbursts are ever exactly alike. Nevertheless, several patterns can be deduced from the long term light curves of SS Cyg and other dwarf novae (see Warner (1995) and references therein for a detailed discussion, from which the following figures have been obtained). It is helpful to decompose the light-curve into various intervals. *Recurrence times* τ_{rec} (i.e. times between successive outbursts) vary from system to system, but tend to coalesce around similar values in any individual system (e.g. $\tau_{\text{rec}} \sim 50$ days for SS Cyg, whereas $\tau_{\text{rec}} \sim 110$ days for U Gem). The shortest intervals observed was just 7 days, while the longest observed interval was 80 years. The *decay time* τ_{decay} tends to be of the order of several days, and is correlated

with the orbital period of the system (e.g. $\tau_{\text{decay}} \sim 1$ day for systems with orbital periods $P_{\text{orb}} \sim 1\text{h}$, and $\tau_{\text{decay}} \sim 4$ days for systems with orbital periods $P_{\text{orb}} \sim 9\text{h}$). In addition, outbursts can be classified based on their rise times τ_{rise} . The majority ($\sim 64\%$) of systems have a ‘short’ rise time of ~ 2 days (known as ‘Class A’ outbursts), while the longest rise times are ~ 10 days (known as ‘Class D’ outbursts). Finally, *outburst durations* have a bimodal distribution, with most systems clustering around ($\tau_{\text{width}} \sim 7$ days and $\tau_{\text{width}} \sim 15$ days).

1.2.4 Theoretical predictions: disk instability model

Several models have been proposed to explain the outbursts in dwarf novae, of which the leading two candidates from the seventies onwards were the *mass transfer instability* (MTI) and the *disk instability model* (DIM). The mass transfer instability hypothesized that the outbursts originated due to a variable mass transfer rate caused by an instability in the *secondary* (Paczynski, 1965; Bath, 1973), whereas the disk instability model hypothesized that the outbursts originated in an instability in the *disk* (Smak, 1971; Osaki, 1974).

By the late 1980s, however, it had become clear that the former model was increasingly in disagreement with observations (see Cannizzo (1993) for a more comprehensive discussion). The MTI was not able to account for why magnetic CVs (in which no disk is present) were not observed to undergo outburst. Another drawback of the MTI was that the hot spot (where the stream from the secondary intersects the disk) was not observed to vary appreciably in brightness over the course of an outburst, in disagreement with the MTI which predicted *enhanced* mass transfer (and therefore bright spot intensity) during outburst. Finally, observations of the disk radius during outburst provided definitive evidence that the disk *increased* in size (consistent with viscous accretion), in direct contradiction with the MTI, which predicted that the disk should shrink during outburst (Smak, 1984a). Thus the mass transfer instability was ruled out on observational grounds, leaving the disk instability model as the current favored model for explaining dwarf novae outbursts.




Figure 1.8: Comparison of vertical disk structure models with numerical simulations in the effective temperature – surface density plane for dwarf novae. The red lines show the hysteresis curve followed by a disk annulus. The solid ‘S-curves’ correspond to thermal equilibria computed using an α -disk model. Data points are the results of 3D MHD fully compressible vertically stratified numerical simulations, exhibiting an enhancement in α at the lower tip of the upper S-curve. Adapted from Hirose et al. (2014).

The disk instability model is best explained with the help of Figure 1.8. In this model outbursts are the result of a thermal-viscous instability in the accretion disk which, together with the ionization/recombination of hydrogen, enables the disk to undergo a hysteresis loop in the plane of surface density and effective temperature. To understand how the model works, first consider the outer annulus of the disk during quiescence when the disk is cool and partially ionized (e.g. at point C in Figure 1.8). Removal of gas via viscous accretion in the annulus occurs at a lower rate than addition of gas via Roche lobe overflow from the secondary. Thus mass transfer gradually results in the build up of material in the annulus and its surface density gradually rises. This is accompanied by a rise in temperature, until the tem-

perature is great enough to ionize hydrogen (at around 5000 K), which in turn greatly increases the opacity of the disk fluid. This results in runaway heating until the annulus settles into the ‘high state’ (point A in 1.8). The accretion rate in the high state is large and the annulus is drained of mass faster than it can be replenished by mass transfer from the secondary. Thus the surface density (and temperature) of the annulus gradually decreases until the annulus is cool enough such that hydrogen can recombine (point B in 1.8). The resultant drop in the opacity results in runaway cooling until the disk is back in the ‘low state’, and the cycle is repeated.

To compare the disk instability model to observations requires solving the full time-dependent viscous diffusion equation (Equation 1.2) in order to produce a light-curve that can be directly compared with observations. Many such calculations have been carried out (see, for example, Papaloizou et al. (1983); Mineshige & Osaki (1983); Smak (1984b) and Hameury et al. (1998)), and they are reviewed at length in Lasota (2001) and, more recently, in Hameury (2019). In the previous paragraph we followed the evolution in the plane of thermal equilibrium (Σ, T_{eff}) of a single annulus. In an actual dwarf nova outburst the entire disk (comprising many such annuli) undergoes this process. Initially a single annulus in the disk reaches a critical surface density at which it undergoes thermal runaway (point D in Figure 1.8). Whether this annulus is located near the outer disk radius or further inside the disk depends on the mass transfer rate from the secondary: a high mass transfer rate will result in the outburst propagating from the *outside-in* (resulting in an inward propagating heating front, in which neighboring annuli reach the critical surface density and temperature at which they enter the high state). A low mass transfer rate, on the other hand, enables material to diffuse inward due to accretion until some annulus at smaller radii becomes the first to reach its critical surface density. In this case the outburst propagates from the *inside-out*. In either case, eventually one of the outer disk annuli, where the surface density is lower to begin with compared to the rest of the disk, reaches a minimum surface density and temperature at which hydrogen is recombined, and thus transitions from the high to the low state. This results in a cooling front propagating from the outside-in, which is associated with

the decay from maximum in the light-curve.

The standard mechanism of the disk instability model described above can reproduce at least the most generic features of observed light curves: repeated outbursts. A further success of the model is that the two types of outburst it predicts (outside-in and inside-out) have successfully accounted for some of the different classes of observed outburst described in Section 1.2.3 (Hellier, 2001). During an outside-in outburst the heating front propagates rapidly to annuli of smaller specific angular momenta; this is associated with short rise times in the light-curve. Inside-out outbursts, on the other hand, must propagate against the angular momentum gradient (i.e. they must move material to radii with greater angular momenta); the heating front travels more slowly in this case, and inside-out outbursts are associated with longer rise times.

Despite these successes, however, the mechanism of dwarf novae outbursts presented above is highly idealized, and many issues remain (see Lasota (2001); Hameury (2019)). In particular Lasota (2001) points out that many time-dependent disk models used to calculate light curves are not robust to small changes in the numerical and physical set-up, such as boundary conditions and irradiation. Some models that employ a fixed outer boundary condition are known to result in weak (low amplitude), short duration outbursts (see, for example, Hameury et al. (1998)). Other models exhibit artifacts known as *reflares* in which the heating and cooling fronts are reflected, resulting in mini-bursts in the light-curve during decay that are not observed in actual systems (Menou et al., 2000). Finally, DIM models predict a gradual *rise* in the light curve during quiescence, in contrast to observations where the light curve is found to be constant or even decreasing slightly over the duration of the low state.

1.3 Convection

Convection is an extremely common and well-studied phenomenon. It is ubiquitous in nature, and is utilized widely in industry. It is found in envi-

ronments as disparate as the interior of the Earth, in oceans, in the atmosphere, the interiors of stars and giant planets, the moons of giant planets, and perhaps, as this thesis attempts to investigate, in accretion disks. It is even a part of our everyday lives, the classic example being convection in a heated pot of water.

Perhaps unsurprisingly for a topic this heterogeneous, the methods of tackling the problem of convection are as diverse as the phenomenon itself: from simple parcel arguments (Landau & Lifshitz, 1987), to mixing length theory (Kippenhahn et al., 1990), to linear stability analyses (Ogilvie, 2016), to dynamical systems approaches (Weiss & Proctor, 2014), and in numerical simulations; from Boussinesq (Chandrasekhar, 1961), to anelastic (Gough, 1969), and (as in this thesis) fully compressible; in hydrodynamics and magnetohydrodynamics. Even the specific problem of determining the criteria for the onset of convection (which we will review in Section 1.3.3) emphasizes different formulations of the instability criteria depending on the nature of the problem. For example, the study of Rayleigh-Bénard convection using the Boussinesq approximation (see, for example, Chandrasekhar (1961)), and its generalization to rotating (Boubnov & Golitsyn, 1995) and MHD (Weiss & Proctor, 2014) flows, emphasizes the role of dimensionless numbers in determining the onset of convection. In the context of convection in stars one often encounters a zoo of ‘gradients’ (Kippenhahn et al., 1990), while in the field of disks the (sign) of the square of the buoyancy frequency is commonly used (Stone & Balbus, 1996; Hirose et al., 2014; Held & Latter, 2018).

In Section 1.3.1 we provide a brief overview of the rich classification and (some might say somewhat redundant) nomenclature of convective flows in the literature, mainly based on the source of buoyancy for driving convection. In Section 1.3.2 we discuss applications of convection in nature (with an emphasis on astrophysical systems), including in stars, in the Earth’s interior, oceans, and atmosphere, and in giant planets and their moons (a detailed discussion of convection in disks is deferred to Section 1.4). Finally, in Section 1.3.3 we briefly review the linear theory of convective stability.

1.3.1 Different types of convection

There are several different types of convection. *Thermal convection* (also called free convection, thermogravitational convection, or buoyancy-driven convection) refers to convection driven by buoyancy forces which arise due to an unstable entropy gradient. A classic example of thermal convection is *Rayleigh-Bénard* convection which refers to convection in the absence of rotation and magnetic fields. Technically Bénard performed the experiments before Rayleigh developed the theory (Bénard, 1900, 1901) and Rayleigh studied true thermal convection (in the sense defined above) in a viscous fluid (Rayleigh, 1916),¹² whereas Bénard observed what is now known as *thermocapillary convection* (Block, 1956), i.e. convection in which variability of the surface-tension coefficient plays a substantial role in driving convection. This kind of surface-tension-driven convection is also occasionally referred to as *Bénard-Marangoni convection* (Lappa, 2009).

Buoyancy forces can also arise due to a compositional gradient rather than an entropy gradient, and convection driven by compositional gradients is referred to as *compositionally-driven convection* (also sometimes called gravitational convection, buoyant convection, or solutal convection). Convection can also be driven by *both* entropy and compositional gradients. If, however, one of the gradients is stabilizing while the other is destabilizing a phenomenon can arise known as *doubly diffusive convection* (Huppert & Turner, 1981). In this particular scenario, the thermal (or compositional) diffusivity can cancel out the stabilizing gradient, and thus actually serves to *destabilize* the fluid. (In astrophysics the term semi-convection (Spiegel, 1969) is sometimes used, whereas geophysicists refer to the process as thermochemical convection (Breuer et al., 2010), and in oceanography the terms thermohaline (Stern, 1960) or thermosolutal (Turner, 1974) convection are common.)

¹²Note that Karl Schwarzschild derived the criterion for convective instability in an inviscid fluid (in the context of stars) well before Rayleigh studied the problem. See Schwarzschild (1906), an English translation of which is available in Menzel & Milne (1966).

Finally, in industrial applications the term *forced convection* is employed to refer to ‘convection’ of fluid due to external forces other than gravity (or buoyancy), such as the bulk motion of fluid driven by a fan or pump, though this is really an *advective* rather than a convective process.¹³ Conversely, convection that is driven by buoyancy is referred to as *natural convection*.

1.3.2 Convection in nature

Convection in stars

Convection is thought to play a key role in heat transport and magnetic field generation in stellar interiors, where it is driven by the release of energy through nuclear fusion reactions. It is also one of the few instances of convection in astrophysics for which there are direct images, mainly the tops of convective cells in the Sun, which impart the solar surface with a distinct granular pattern (see Figure 1.9).

The Sun is not the only star in which convection is present. Low mass ($M < 0.4M_{\odot}$), low luminosity M-type stars are fully convective throughout. Intermediate mass ($0.4 < M/M_{\odot} < 1.5$) G-type stars (such as the Sun) consist of an inner radiative zone which is convectively stable (at least according to the Schwarzschild stability criterion, cf. Equation 1.18), and in which the bulk transfer of energy is carried by radiation. The radiative zone is surrounded by an outer convective zone, which comprises around 30% of the star’s radius. In general as the mass of a star is increased, the depth of the convective zone *decreases*. In high mass stars ($M > 1.5M_{\odot}$; type A, B, and O), the position of the convective and radiative zones is reversed, however, with the stars consisting of an *inner* convective zone surrounded by an outer radiative zone (this reversal occurs because high mass stars can support nuclear reaction rates, such the carbon-nitrogen-oxygen (CNO) cycle, that are particularly sensitive to temperature, and which produce very large temperature gradients as a result; see below). Note that in low and high mass stars the gravitational acceleration \mathbf{g} changes sign at the base of the convective

¹³In this thesis we use the term *forced convection* to mean *thermal convection* that is sustained through a heat source that has been added by hand, as in Section 3.4.




Figure 1.9: Image of the surface of the Sun taken in January 2020 at 789 nm with the 4 m Inouye Solar Telescope on Mt. Haleakalā, Hawaii. The bright cellular patterns are the tops of rising convective cells, while the darker lanes correspond to cooler, sinking fluid. Bright spots indicative of magnetic fields can be seen within the downflows. Credit NSO/NSF/AURA.

zone (which coincides with the stellar center), whereas in intermediate mass stars \mathbf{g} maintains the same sign throughout the convective zone.

As will be discussed in more detail in Section 1.3.3, amongst other things a fluid is unstable to thermal convection provided the temperature gradient exceeds a critical value (known as the Schwarzschild criterion for instability). This temperature gradient (and thus convective stability) is related to the internal properties of the star, mainly its opacity κ (which depends on the density, temperature, and composition of stellar material) and heat flux F (which is set by the release of nuclear energy in fusion reactions in the star). In spherical coordinates, the radial temperature gradient dT/dr can be related to the flux of radiation through a spherical surface of radius r and the opacity κ by

$$\frac{dT}{dr} = -\frac{3}{4ac} \frac{\kappa \rho}{T^3} \frac{F}{4\pi r^2}, \quad (1.6)$$

where a is the radiation constant, c is the speed of light, and T and ρ are the temperature and density, respectively (Prialnik, 2000). Here it has been assumed that stellar material is optically thick, and so radiation (photons) propagates through the fluid diffusively. Thus Equation 1.6 indicates that strong temperature gradients occur when either the heat flux F or the opacity κ are large, and thus convection is favored under those circumstances. The former condition is met in the interiors of high mass stars where hydrogen is fused into helium via the CNO cycle, which is highly temperature sensitive. Large opacities occur in stellar envelopes where cooler temperatures (compared to the core) enable partial ionization of hydrogen, which in turn leads to a large increase in the opacity.

Nucleosynthesis (or accretion) of heavier elements can also set-up unstable compositional gradients in stars, and under certain conditions stellar interiors can be unstable to various types of doubly diffusive convection (for reviews see Spiegel (1972); Garaud (2018)). This can occur when heavier (e.g. helium rich) material is accreted from a secondary onto a primary in binary systems (Stancliffe & Glebbeek, 2008), or when metal-rich planetesimals are accreted onto a single star (Vauclair, 2004). Unstable compositional gradients can also arise when nuclear burning occurs away from the center of the star, for example off-center ignition of carbon in high mass ($7 - 10M_{\odot}$) asymptotic giant branch (AGB) stars (Siess, 2009). Finally, on the outer edge of the hydrogen burning shell in red giants (which is still located within the thermally convectively *stable* radiative zone), a rare nuclear reaction of the form ${}^3\text{He} + {}^3\text{He} \longrightarrow {}^4\text{He} + \text{p} + \text{p}$ produces more particles per unit mass than it started with and thus lowers the mean molecular weight (Ulrich, 1972). The overall picture is similar to one in which (heavier) salt water lies on top of (lighter) fresh water. The rising fluid elements manifest themselves as narrow ‘fingers’ that can mix material from the hydrogen burning shell well into the convective zone.

Finally, in high mass stars the opposite situation can occur, i.e. a destabilizing thermal gradient is (over-)stabilized by a stable compositional gradient. This is analogous to the situation in oceanography when hot, salty water lies beneath cold, fresh water (see, for example, Turner (1974)). If the thermal

diffusivity is sufficiently large, an upwardly displaced ‘hot and salty’ fluid element comes into thermal – but not compositional – equilibrium with its ‘cold and fresh’ surroundings. Since salt water is denser than fresh water, the fluid element will cease its upward rise and begin to descend. Thus an unstable thermal gradient has been arrested by a stable compositional gradient. The sinking fluid element will overshoot its equilibrium position, however, and will begin to execute growing oscillations about that position, a phenomenon known as *convective overstability*. In high mass stars this situation arises when a helium-rich core is surrounded by a hydrogen-rich envelope (Spiegel, 1969, 1972). The large sensitivity of certain nuclear reaction rates (in particular the CNO cycle) to temperature results in a large heat flux that renders the core thermally convectively unstable. However, if the thermal diffusivity is sufficiently large, an upward displaced helium-rich fluid element will come into thermal equilibrium with the lighter hydrogen-rich surroundings and thus begin to descend. The resultant oscillatory (or overstable) convection efficiently mixes the helium-rich core and its hydrogen-rich surroundings.

Even compact objects such as white dwarfs and neutron stars are thought to have convective outer layers. White dwarfs harbor a convective atmosphere (comprising $< 1\%$ of the stellar radius) and consisting mostly of hydrogen (see Figure 1.10). Neutron stars have a very thin liquid ‘atmosphere’ (sometimes referred to as an ‘ocean’) just 30 m thick, enveloping a 1 km thick crust. Accretion of material onto the neutron star surface undergoes nuclear burning in which synthesis of heavy elements via rapid proton capture occurs. The heavy material sinks to the ocean-crust boundary where lighter elements are chemically separated from heavier ones. The lighter elements then rise to the top in a form of compositionally-driven convection that mixes the neutron star ocean (Medin & Cumming, 2011).

Convection on Earth: outer core

On Earth convection occurs in the outer core, in the mantle, in the crust (in magma chambers), in the ocean, and in the atmosphere. In the outer core of the Earth (which is composed primarily of liquid iron and nickel)




Figure 1.10: Velocity streamlines showing convective cells in the atmosphere of a white dwarf taken from a 2D numerical simulation. The solid contour lines are lines of constant temperature. The effective temperature at the top of the domain is $T_{\text{eff}} = 12\,200\text{ K}$ and the gravitational acceleration was taken to be $\log g = 8.0$. Credit Freytag et al. (1996).

convection is driven by both thermal and compositional gradients due to cooling of the hot inner core (specifically, crystallization of liquid iron at the inner core boundary releases latent heat and leaves behind lighter elements; Fearn & Loper (1981)).¹⁴ Although it is not definitively known whether the outer core is in differential or uniform rotation,¹⁵ to first order the convective flows are arranged in columns aligned with the rotation axis as expected from the Taylor-Proudman constraint, which restricts vertical motions in a rotating fluid to narrow columns due to the influence of the Coriolis force (Breuer et al. (2010); see also the left-hand panel of Figure 1.11). However the outer core is viscous, which breaks the restrictions imposed by the

¹⁴The decay of radioactive elements, such as potassium ($K-40$) also acts as a source of heating, though it is unclear at present whether this is significant compared to the latent heat released by solidification (Nimmo, 2007).

¹⁵At present it is thought to be more or less in uniform rotation (Chris Jones; private communication)

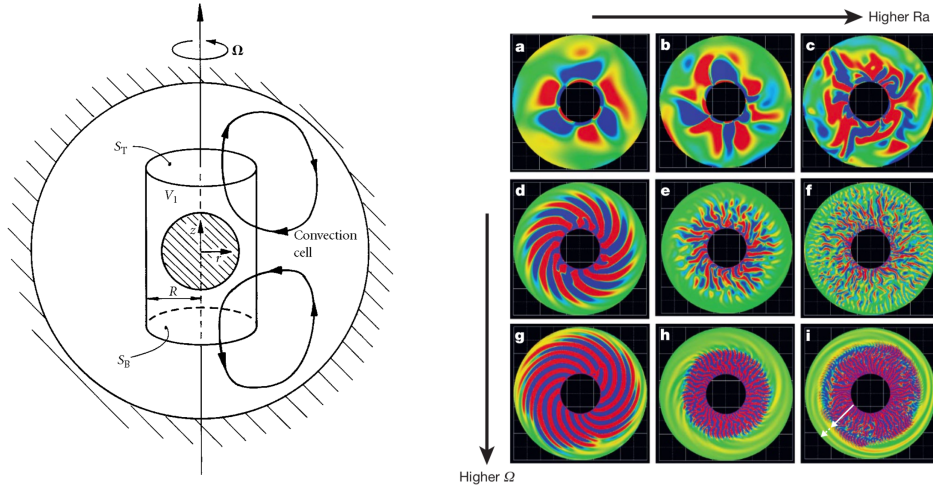


Figure 1.11: Left: schematic of convection in Earth's outer core. From Davidson (2015). Right: numerical simulations of convection in Earth's outer core exhibiting the effect of Rayleigh number Ra and angular speed Ω on convective patterns. The colorplot shows the axial vorticity and the snapshots are latitudinal slices. From Miyagoshi et al. (2010).

Taylor-Proudman theorem and results in large-scale secondary convective flows that assume a petal-like pattern. These convective fluid motions play a fundamental role in generating and maintaining the Earth's dipolar magnetic field, which varies between 0.3-0.6 G (30-60 μ T) at the Earth's surface (Aurnou et al., 2015), via the geodynamo (Davidson, 2015).¹⁶ Finally, high Rayleigh number convection in the Earth's outer core has been shown to generate coherent structures, including spirals and large-scale vortices (see right-hand panel of Figure 1.11).

¹⁶The possibility of the Earth's magnetic field being a fossil field was ruled out because the temperature in the outer core ($\sim 6 \times 10^3$ K) is hotter than the Curie temperature (1043 K) at which the spins of the electrons and nuclei that create the magnetic field of a permanent magnet are randomized. Thus a fossil field would have decayed within several tens of thousands of years due to the effects of resistivity.




Figure 1.12: Left: schematic of convective cells in the mantle of the Earth. From Lappa (2009). Right: velocity colorplot exhibiting convective plumes, taken from a high resolution 2D numerical simulation with adaptive mesh refinement of mantle convection in a quarter shell. From Kronbichler et al. (2012).

Convection on Earth: mantle

Between the edge of the outer core (at $r \sim 2.4 \times 10^3$ km) and the bottom of the crust (at $r \sim 5.3 \times 10^3$ km), convection is also found in what is known as the Earth's mantle (which consists primarily of molten rock containing iron and magnesium). The mantle evolves on timescales that are much longer than those of the outer core, and therefore the mantle can be treated as a solid compared to the outer core. The heat source for convection in the mantle is a combination of residual heat flux from the core, as well as decay of radioactive elements within the mantle itself. Unlike in the outer core, magnetic fields play no role in the convective motions in the mantle, because the mantle is a nearly perfect electrical insulator and therefore cannot support any currents. Convective flows in the mantle are, however, ultimately responsible for the drift of the continents. The convective flows are cell-like: buoyant mantle fluid rises, eventually forming a mid-ocean ridge between two oceanic plates, while cool fluid sinks at subduction zones between continental and oceanic plates (see Figure 1.12).

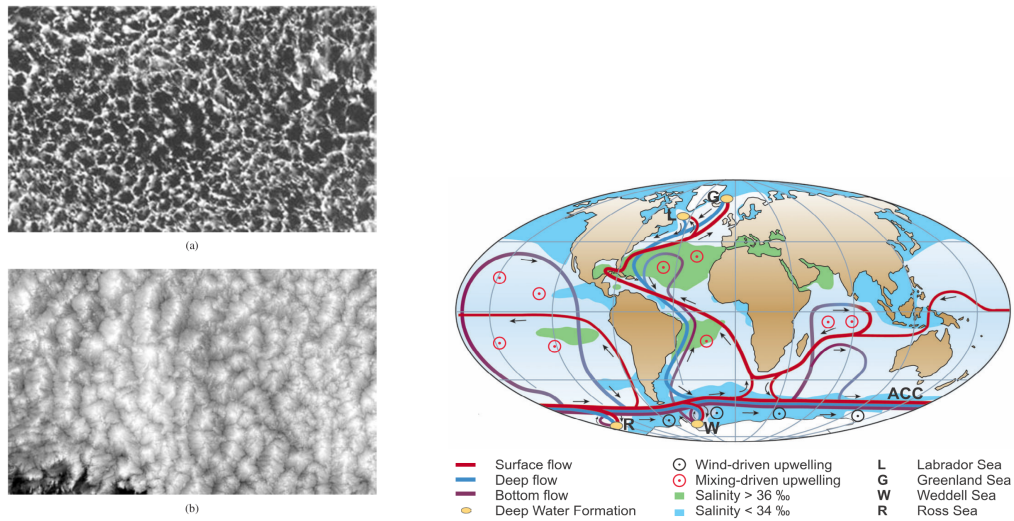


Figure 1.13: Left: convection in the Earth’s atmosphere leads to distinct cloud patterns: open-celled cumulus clouds (top panel), and closed-cell stratus clouds (bottom panel). Credit: Lappa (2009). Right: schematic showing the Earth’s ocean currents, which are manifestations of thermohaline convection. Taken from Rahmstorf (2006).

Convection on Earth: oceans and atmosphere

Convection is also present in the Earth’s oceans and atmosphere. In the oceans an example of convection is the Gulf Stream, a major ocean current (about 100 km wide and 1 km deep) that stretches from the Gulf of Mexico along the eastern coast of North America. It is an example of *thermohaline convection* (Rahmstorf, 2006), where northward-moving water cooled by wind-driven evaporation increases in density and salinity until it reaches the poles where it sinks under the lighter, less saline water of the North Atlantic. This cool downdraft then moves southward to the equator where it is heated by solar radiation, causing it to rise back up to the surface (see right-hand panel of Figure 1.13).

In the Earth’s atmosphere convection is manifested on large-scales through atmospheric circulation, and also on smaller scales inside clouds (Stevens, 2005; Emanuel, 1994). The large-scale atmospheric circulation (known as

Hadley circulation, (Hadley, 1735)) consists of convective cells that are driven by heating from the *sides*: the cells are heated by solar radiation at the equator and cooled at the poles (Wallace & Hobbs, 2006). They consist of three giant cells in each quarter of the Earth’s atmosphere. Solar radiation can also heat the Earth’s surface, which in turn warms the air above it. Thus on smaller scales (sometimes referred to as *mesoscales*) convection is also manifested as clouds, either as ‘open-cells’ (cumulus clouds) with broad downdrafts at their centers and thin updrafts at their edges, or as ‘closed cells’ (stratus clouds) with broad updrafts at their centers and thin downdrafts at their edges (see left-hand panel Figure 1.13).

Convection in giant planets and their moons

Outside of terrestrial applications, convection also plays an important role in the giant planets, particularly Jupiter and Saturn. In both planets, convection is thought to drive the *zonal flows* (or *jets*) that form the banded structure of their atmospheres, with equatorial speeds reaching 150 ms^{-1} and 300 ms^{-1} in Saturn and Jupiter, respectively (Sun et al., 1993; Jones & Kuzanyan, 2009). The lighter-colored flows (known as *zones*) correspond to convective upwellings, while the darker-colored flows (known as *belts*) correspond to downdrafts (Ingersoll et al., 2000; Gierasch et al., 2000). Rotation and magnetic fields are important: the inner parts of Jupiter are thought to be in solid body rotation, but in the outer parts the angular frequency increases with cylindrical radius and rotation is constant on cylinders (Guillot et al., 2018). Thus equatorial zonal flows rotate faster than zonal flows at higher latitudes. Furthermore, there is now some evidence that compositional gradients exist in giant planet interiors, and that these can stabilize unstable thermal gradients, leading to an oscillatory form of doubly diffusive convection (or semi-convection) similar to that found in the interiors of high mass stars (see Section 1.3.2), with consequences for both heat transfer and tidal dissipation (Pontin et al., 2020). In Saturn a particularly beautiful convective feature is the hexagonal structure (about $2.5 \times 10^4 \text{ km}$ wide and about 100 km deep) on its north pole imaged by the Cassini-Huygens mission

(Yadav & Bloxham, 2020).¹⁷

The moons of Jupiter (particularly Europa, Ganymede, and Callisto) have also been hypothesized to harbor saline liquid water oceans beneath their icy crusts. Though there is still some debate as to how deep (and how liquid) these subsurface oceans really are, the prevailing picture appears to be that there is a thin ($\sim 10\text{-}30\text{ km}$) solid crust of brittle ice, while at depths of around 100 km tidal dissipation is sufficient to completely melt the ice to form a liquid ocean. Between the solid ice surface and the liquid ocean there is a layer of warm or ‘mushy’ ice that is convecting. Note that here convection is being driven not by an intrinsic internal heat source, but by tidal dissipation (Pappalardo et al., 1998; Spohn & Schubert, 2003). A convective dynamo might be responsible for generating the magnetic fields ($\sim 20\text{-}600\text{ nT}$) measured at the moons’ equators (Zimmer et al., 2000). The Cassini mission (1997-2017) also detected geysers erupting from the icy moon Enceladus, which orbits Saturn, which have been interpreted as evidence of a subsurface ocean (Iess et al., 2014). The possibility of tidal-dissipation-driven convection in Enceladus has been discussed by Barr & McKinnon (2007).

1.3.3 Linear theory of convection

In this section we first introduce the concept of buoyancy, which is critical to understanding the dynamics of convection. We will then present more formal criteria for convection. Although later in this thesis we only consider convection driven by unstable entropy (not compositional) gradients in disks, here we will nevertheless present the criteria for convective instability in the most general form possible, i.e. taking into account both thermal and compositional gradients. This has the advantage that it will then be easy to see how these general expressions simplify for various special cases, e.g. for a uniform (i.e. chemically homogeneous) perfect gas in the local effective gravitational potential of an accretion disk.

Convection is fundamentally a *buoyancy-driven* instability. The origin

¹⁷An alternative hypothesis is that the polar hexagon is formed by baroclinic instability in the stably-stratified surface weather layer rather than convection (Aguiar et al., 2010).

of buoyancy is due to the vertical variation (stratification) of pressure in a fluid due to the effect of gravity.¹⁸ To see how this vertical variation in pressure arises, consider an open tank filled with fluid at rest and of uniform density ρ_f , and permeated by a uniform gravitational field g along the vertical direction. Let the vertical distance from the surface of the fluid be a measure of the depth $z > 0$ of a point in the fluid (let $z = 0$ at the surface). Now consider the fluid contained within a fictional thin circular slab of area A located at a depth $z = h$ in the tank. The fluid in the slab must support the weight mg of the fluid in a column of mass m above it, otherwise the fluid would not be at rest. Thus the downward force on the slab due to gravity $F_g = mg$ must be balanced by an upward force F_B . Rewriting the mass of the column of fluid directly above the slab as $m = \rho_f V$, and substituting for the volume of the column $V = Ah$, the upward force on the slab can be written as $F_B = \rho_f Ahg$. Finally dividing both sides by the area A , we see that the pressure is given by $P = \rho_f hg$. Thus pressure increases with depth in a fluid sitting in a gravitational field in hydrostatic equilibrium.

Now suppose a cylindrical object of uniform density ρ_{obj} and volume V_{obj} is immersed fully in the fluid. Let the vertical extent of the object be such that it spans $z = h_t$ from the top of the object (i.e. the point of the object closest to the surface) and $z = h_b$ at the bottom of the object; the object has a vertical extent of $\Delta h = h_d - h_t$. Since pressure increases with depth, the pressure exerted at the top of the object $P_t = \rho_f h_t g$ will be less than the pressure exerted at the bottom of the object $P_b = \rho_f h_b g$. The net upward force on the object is given by $F_{B,\text{obj}} = A(P_t - P_b) = -\rho_f A \Delta h g$ (the minus sign is because the force points upwards, or towards the surface), and is known as the *buoyancy force*. Thus we see that buoyancy arises due to the vertical variation in pressure in the fluid, which itself is due to the gravitational field permeating the fluid.

Taking the ratio of the magnitudes of the buoyancy force $F_{B,\text{obj}}$ and the weight of the object $F_{g,\text{obj}} = m_{\text{obj}} g = \rho_{\text{obj}} A \Delta h g$, we obtain $F_{B,\text{obj}}/F_{g,\text{obj}} =$

¹⁸The reader should not mistake this to mean that every pressure gradient results in a buoyancy force, however.

ρ_f/ρ_{obj} . Thus if the density of the immersed object is less than the density of the surrounding fluid $F_{B,\text{obj}} > F_{g,\text{obj}}$ and object will rise; otherwise if $\rho_{\text{obj}} > \rho_f$ the object will sink. Furthermore, we can obtain an important insight by considering the fluid within a surface whose shape (and therefore volume) is identical to that of the immersed object. Let such a parcel of fluid have a mass m_f . The fluid parcel is at rest, so its weight must be balanced by the buoyancy force on the fluid parcel, thus $F_{B,f} = -m_f g$. But the buoyancy force on the parcel of fluid must in turn be equal to that on the immersed object, since the two have the same configuration, and so $F_{B,f} = F_{B,\text{obj}}$. Thus we arrive at the key insight that the buoyancy force on the immersed object is equal to the weight of the fluid it has displaced. This is known as *Archimedes' principle*. Note that this is true irrespective of the shape of the object. To work out the buoyancy force on an object of arbitrary shape we need only replace the object with an element of fluid of the same shape and recognize that, since the fluid element is at rest, the buoyancy force on the fluid element must be equal to its weight. If the density of the fluid element is known, the problem of determining the buoyancy on the fluid element reduces to that of finding the volume of the fluid element, and since the volume of the fluid element is identical to that of the object we need only work out the volume of the object to determine the buoyancy force on the object, i.e. $F_{B,\text{obj}} = -\rho_f V_{\text{obj}} g$.

In the previous example we considered the buoyancy force on an object immersed in the fluid. To generalize this consider instead a fluid, initially in hydrostatic equilibrium, with a vertical gradient in *density* $\rho(z)$ as well as pressure $P(z)$. The fluid is permeated by a uniform vertical gravitational field. Suppose a fluid element is displaced adiabatically upwards from some depth $z = h_b$ to a depth $z = h_t$. Through the exchange of acoustic waves the fluid element will come into pressure equilibrium with its new surroundings on a dynamical timescale but it will, in general, have a different density ρ_{el} to the surroundings at its new position $\rho_{\text{sur}} = \rho(z(h_t))$. As in the previous example, the gravitationally induced stratification in pressure results in a buoyancy force on the fluid element. If $\rho_{\text{el}} < \rho_{\text{sur}}$ the fluid element will rise; otherwise it will sink. When this rising and sinking of fluid elements is

sustained by some means (which we will discuss shortly) such that there is a continuous circulation of fluid driven by the interplay between buoyancy and gravity, the fluid is said to be *convecting*. Thus gravity can be thought of as the very basic ingredient needed for convection in a fluid.¹⁹

There are two sources of buoyancy: *entropy gradients* in which case a fluid element is lighter than its surroundings by virtue of having a higher specific entropy than its surroundings, and *compositional gradients*, in which case a fluid element is lighter than its surroundings by virtue of its composition. Of course the converse can also hold: if a fluid element is cooled it will contract causing it to become denser than its surroundings, or a fluid element might be denser than its surroundings by virtue of its chemical composition. Thus thermal convection can be driven either by ‘heating from below’ (as in the traditional Rayleigh-Bénard experiment) or by ‘cooling from above’ (an example of which is the rapidly cooling surface of the exposed liquid iron core following the merger of two planetesimals (Neufeld et al., 2019)).²⁰ Compositional convection can be driven either by the synthesis of light elements from below (such as on the surface of the Earth’s inner core) or by the synthesis of heavy elements from above which subsequently sink (such as the hydrogen burning shell in red giant stars).

Consider a fluid element located at some fiducial location z_0 in a gravitational field of magnitude g , and suppose the coordinate axes are aligned such that g points along the negative z -axis. For a chemically inhomogeneous fluid with a general equation of state $P = P(\rho, T, \mu)$, the change in density $D\rho$ of a fluid element displaced by amount δz is given by (see (Kippenhahn et al., 1990))

¹⁹A straightforward table-top demonstration of the absence of convection in zero gravity (strictly speaking, microgravity) has been given by the ESA’s André Kuipers aboard the International Space Station. See <https://www.youtube.com/watch?v=GhQHXUy3tNg&frags=pl%2Cwn>. More quantitative descriptions of such experiments can be found in Naumann (1999).

²⁰By symmetry considerations, the planforms of the convective cells when convection is driven by heating from below will be the same as those when convection is driven by cooling from above in Boussinesq convection. Compressible effects or a spatially-dependent viscosity, however, can break this symmetry.

$$D\rho = \frac{\rho\delta}{H_P} \left(\nabla_{\text{el}} - \nabla_{\text{sur}} + \frac{\phi}{\delta} \nabla_{\mu} \right) \delta z, \quad (1.7)$$

where ∇_{sur} , ∇_{el} , and ∇_{μ} are various ‘gradients’ used to express the variation of a quantity with depth (for which the pressure P is typically used as proxy).²¹ Explicitly, the temperature gradient of the *surroundings* ∇_{sur} is defined as

$$\nabla_{\text{sur}} \equiv \frac{d \ln T}{d \ln P}, \quad (1.8)$$

the compositional gradient ∇_{μ} of the *surroundings* is defined as

$$\nabla_{\mu} \equiv \frac{d \ln \mu}{d \ln P}, \quad (1.9)$$

and ϕ and δ are defined as

$$\phi \equiv \left. \frac{\partial \ln \rho}{\partial \ln \mu} \right|_{T,P}, \text{ and } \delta \equiv \left. \frac{\partial \ln v}{\partial \ln T} \right|_{\mu,P} = - \left. \frac{\partial \ln \rho}{\partial \ln T} \right|_{\mu,P}, \quad (1.10)$$

respectively, where $v \equiv 1/\rho$ is the specific volume. Note that δ is related to the volume-expansivity $\alpha \equiv (\partial \ln V / \partial T)_{\mu,P}$ of the fluid by $\delta = T\alpha$. Since most fluids expand on heating δ can usually be assumed to be positive. For the special case of an ideal gas with $P = (\mathcal{R}/\mu)\rho T$, $\phi = 1$ and $\delta = 1$. Finally the pressure-scale height is defined as

$$H_P \equiv - \frac{dz}{d \ln P}. \quad (1.11)$$

Since the pressure generally increases with depth, $dz/d \ln P < 0$ and so the scale height, too, is a positive quantity, i.e. $H_P > 0$.

The gradient $\nabla_{\text{el}} \equiv d \ln T / d \ln P$ describes the variation in the temperature with depth of the *fluid element* as it moves through its surroundings. In the special case that the fluid element moves without exchanging heat with the surroundings, its motion is said to be *adiabatic* and the gradient of the fluid element is equivalent to the *adiabatic gradient*, i.e. $\nabla_{\text{el}} \rightarrow \nabla_{\text{ad}}$,

²¹In the literature, the temperature gradient of the surroundings ∇_{sur} is usually denoted simply by ∇ .

where the adiabatic gradient is defined as the derivative of temperature with respect to pressure (or depth) at constant entropy s , i.e.

$$\nabla_{\text{ad}} \equiv \left. \frac{\partial \ln T}{\partial \ln P} \right|_{s, \mu}. \quad (1.12)$$

Instability occurs when $D\rho < 0$. Thus the general instability criterion for convection (known as the *Ledoux criterion* after Ledoux (1947)) is given by

$$\boxed{\nabla_{\text{el}} - \nabla_{\text{sur}} + \frac{\phi}{\delta} \nabla_{\mu} < 0.} \quad (1.13)$$

The instability criterion given by Equation 1.13 is valid for a fluid of non-uniform composition and for a general equation of state $P = P(T, \rho, \mu)$.

It is often more intuitive to think of convective instability in terms of the motion of the displaced fluid element, rather than purely in terms of thermal and compositional gradients. A fluid element that is displaced by an amount δz in a gravitational field g will feel a buoyancy force density of $f_B = -|g|D\rho$. From Newton's second law the acceleration of the fluid element is given by

$$\rho \frac{d^2 \delta z}{dt^2} = -|g|D\rho. \quad (1.14)$$

Substituting Equation 1.7 into Equation 1.14 and dividing through by ρ we obtain

$$\frac{d^2 \delta z}{dt^2} = - \left[\frac{|g|\delta}{H_P} \left(\nabla_{\text{el}} - \nabla_{\text{sur}} + \frac{\phi}{\delta} \nabla_{\mu} \right) \right] \delta z. \quad (1.15)$$

The quantity in the brackets has dimensions of frequency squared. It is defined as the (square) of the *buoyancy-frequency* N_B^2 , i.e.

$$\boxed{N_B^2 \equiv \frac{|g|\delta}{H_P} \left(\nabla_{\text{el}} - \nabla_{\text{sur}} + \frac{\phi}{\delta} \nabla_{\mu} \right).} \quad (1.16)$$

The equation of motion of the fluid element is thus given by the second-order ordinary differential equation $d^2 \delta z / dz^2 = -N_B^2 \delta z$. A displacement δz will grow exponentially when $N_B^2 < 0$ (corresponding to *instability*), $N_B^2 = 0$ corresponds to *marginal* stability, and $N_B^2 > 0$ results in the displaced

element undergoing oscillations (stability). These oscillations lead to waves known *internal gravity waves*.

So far we have written the instability criterion for convection purely in terms of ‘gradients’ (Equation 1.13), and in terms of a buoyancy frequency (Equation 1.16). An equivalent way of considering convective instability is to express the instability criterion in terms of an *entropy gradient*

$$\boxed{\frac{1}{c_P} \frac{ds}{dz} - \frac{\phi}{\delta} \frac{d \ln \mu}{dz} < 0.} \quad (1.17)$$

Thus negative entropy and positive compositional gradients *destabilize* the fluid, whereas positive entropy and negative compositional gradients *stabilize* the fluid. Some interesting scenarios arise, however, when the entropy and compositional gradients are of the same sign. A fluid with a negative entropy gradient (i.e. thermally unstable) and a negative compositional gradient (i.e. compositionally stable) might be convectively stable if the magnitude of the compositional gradient exceeds the magnitude of the entropy gradient. Conversely, a fluid with a positive entropy gradient (i.e. thermally stable) and a positive compositional gradient (i.e. compositionally unstable) might be convectively unstable if the magnitude of the compositional gradient exceeds the magnitude of the entropy gradient.

Two special cases arise from Equation 1.13. In the case of a fluid of uniform composition $\nabla_\mu = 0$, and the Ledoux criterion for instability reduces to the *Schwarzschild criterion* for instability (Schwarzschild, 1906; Menzel & Milne, 1966).²²

$$\nabla_{\text{sur}} > \nabla_{\text{el}}. \quad (1.18)$$

In the case where the fluid element moves adiabatically $\nabla_{\text{el}} \rightarrow \nabla_{\text{ad}}$ and the Schwarzschild criterion becomes

²²The criterion was first derived in 1906 by Karl Schwarzschild (1873-1916), who is also known for deriving the first exact solutions to the Einstein Field Equations for a static, non-rotating, uncharged spacetime outside a spherical mass. His son, Martin Schwarzschild (1912-1997), also worked on convection in stars and published a treatise on stellar structure (Schwarzschild, 1958).

$$\nabla_{\text{sur}} > \nabla_{\text{ad}}, \quad (1.19)$$

Thus convection is said to occur when the temperature gradient of the surroundings is *super-adiabatic*. A further simplification can be made when the fluid is described by a perfect gas equation of state. In this case $\nabla_{\text{ad}} = (\gamma - 1)/\gamma$ and thus the instability criterion reduces to the simple expression

$$\nabla_{\text{sur}} \equiv \frac{d \ln T}{d \ln P} > \frac{(\gamma - 1)}{\gamma}, \text{ (valid for perfect gas)}. \quad (1.20)$$

This can be expressed equivalently in terms of the buoyancy frequency as

$$N_B^2 \equiv |g| \left[\frac{1}{\gamma} \frac{d \ln P}{dz} - \frac{d \ln \rho}{dz} \right] < 0, \text{ (valid for perfect gas)}. \quad (1.21)$$

For the specific case of a thin accretion disk in a local (co-orbiting) reference frame with angular frequency Ω_0 (see Section 2.2), the vertical component of the gravitational acceleration is given by $g = -z\Omega_0^2$.

The second case occurs when the entropy gradient is zero but the compositional gradient is non-zero. In this case the Ledoux criterion given by Equation 1.17 reduces to

$$-\frac{\phi}{\delta} \frac{d \ln \mu}{dz} < 0. \quad (1.22)$$

In the special case of an ideal gas $\delta = 1$ and $\phi = 1$, and the criterion becomes

$$\frac{d \ln \mu}{dz} > 0, \quad (1.23)$$

In other words the fluid is *unstable* when its mean molecular weight decreases with depth (i.e. when ‘heavy’ fluid lies on top of ‘lighter’ fluid). This is a general form of the *Rayleigh-Taylor instability*, except it involves a gradual change in composition rather than a discontinuity.

Finally, the criteria discussed above constitute both necessary and sufficient criteria for the onset of convection only in an *inviscid* fluid. In a viscous fluid, the destabilizing buoyancy force is countered by the effects of viscosity ν and thermal diffusivity χ . The ratio of destabilizing and stabilizing

processes is quantified by the *Rayleigh number*

$$\text{Ra} \equiv \frac{|N_B^2| H^4}{\nu \chi}, \quad (1.24)$$

where H is the disk's scale height. In a viscous fluid, the Rayleigh number must exceed some critical value in order for the onset of convection to occur. Note that when magnetic fields and rotation are present the critical Rayleigh number itself is a function of other dimensionless numbers involving the magnetic field strength, and the rotation frequency, respectively (see Chandrasekhar (1961) for a detailed analysis). We will discuss other dimensionless numbers that are relevant to convection in Section 3.1.4, and a more detailed discussion of the Rayleigh number criterion in turbulent flows is given in Section 4.1.2.

1.4 Applications of convection in disks

1.4.1 Convection in dwarf novae

Convection (perpendicular to the plane of the disk) has long featured in theoretical models of dwarf novae. As discussed in Section 1.2.3, a defining observational feature of dwarf novae is the presence of repeated outbursts in their light curves. A key prediction of the disk instability mode (DIM) used to explain these outbursts is the existence of two thermal equilibrium solutions in the shape of an 'S-curve' in the plane of temperature against surface density (Smak, 1971; Osaki, 1974). The exact instability mechanism enabling the disk to transition between the two states remained unclear, however, until Hōshi (1979) suggested that the transition could be facilitated through ionization of hydrogen at temperatures around 10^4 K. Such ionization transitions result in an opacity with a very strong temperature dependence, allowing for a transition between the two thermal equilibrium solutions proposed in the disk instability model. Hoshi's refined model, however, utilized vertically averaged equations which necessitated a separate calculation of the lower and upper branches, and thus failed to account for the cyclical nature of the outbursts.

Convection appeared to solve this problem. The inclusion of a convective heat flux in the thermal energy equation in 1D steady state disk models correctly predicted ‘S-curves’ for thermal equilibrium, thus enabling the disk to undergo *cycles* between the low and the high states (Meyer & Meyer-Hofmeister, 1981, 1982; Smak, 1982; Cannizzo et al., 1982). The significance of convection in these cycles was called into question, however, by Faulkner et al. (1983), who did not include an explicit convective heat flux in their 1D steady state models and demonstrated that ionization/recombination of hydrogen and the subsequent strong temperature dependence of the opacity was sufficient for the models to produce ‘S-curves’ for thermal stability, and therefore cyclical outbursts. Although Faulkner et al. (1983) demonstrated that convection was not the *cause* of cyclical outbursts in dwarf novae, their calculations indicated that convection was a natural *consequence* of the vertical variation in ionization needed to produce these outbursts. More specifically, convection shifts the position of the local maximum on the upper branch of the S-curve, although the quantitative details of this shift depend on the particular prescription used for the convective heat flux (Cannizzo, 1993).

1.4.2 Convection in LMXBs and AGN

Convection might also be present in low mass X-ray binaries (LMXBs), systems in which a black hole (BH) or neutron star (NS) primary accretes via Roche lobe overflow from a low mass secondary. Like dwarf novae, LMXBs are observed to undergo repeated outbursts. The outburst phase, however, is much more complicated than in dwarf novae, consisting of multiple states. ‘Super soft’ and ‘high/SPL’ states are characterized by very high accretion rates and are modeled by a geometrically thin, optically thick disk which extends all the way down to the innermost stable circular orbit (ISCO). The disk in these states is likely radiation pressure dominated and hence convectively unstable, at least according to simple alpha disk models (Bisnovatyi-Kogan & Blinnikov, 1977; Tayler, 1980).

The disk morphology in the ‘low’ and ‘hard’ states, and in quiescence, is more complicated than in the high states. The relatively cool, geomet-

rically thin, optically thick outer regions of these disks are very similar to those of dwarf novae. The inner regions, however, differ substantially from dwarf novae disks. They are very hot, optically thin, and geometrically thick (bearing greater similarity to tori than to thin disks). General relativistic effects and radiation pressure are important. These thick accretion flows (tori) also carry over to low-luminosity active galactic nuclei (AGN), systems in which an accretion disk surrounds a supermassive black hole at the center of a galaxy.

Several theoretical models have been developed to describe the inner regions of LMXBs and AGN. Narayan & Yi (1994) constructed models called advection-dominated accretion flows (ADAFs). These models were shown to be *very* convectively unstable in the plane of the disk, in fact so convectively unstable that the disk fluid is advected inwards faster than it can thermalize. Because rotation exerts a stabilizing influence on convection (via the *Solberg-Hoiland criterion*, see, for example, Kippenhahn et al. (1990)), ADAF disks are likely to be even more convectively unstable in the plane *perpendicular* to the disk (where the Solberg-Hoiland criterion reduces to the less stringent *Schwarzschild criterion*). Various extensions of, and modifications to, these models have been introduced, including advection-dominated inflow outflow solutions (ADIOS) (Quataert & Gruzinov, 2000), and convection-dominated accretion flows (CDAFs) (Blandford & Begelman, 1999). The CDAF models, in particular, enforce a constant entropy gradient in the radial direction, the motivation being that an ADAF will settle into a marginally unstable state due to the process of convection. However, the interiors of stars, too, are only very marginally convectively unstable (formally defined as a very small but negative N_B^2), yet convection is thought to significantly influence their structure. Therefore convection in CDAFs, though tame compared to convection in ADAFs, likely plays an important role in determining the vertical as well as radial structure of the disk.

1.4.3 Convection in protoplanetary disks

The existence of thermal convection in cooler protoplanetary disks is not clear. Steady-state models in which heating is provided solely via an α viscosity have found that the disk can be convectively unstable in the inner regions, generally within 10 AU, although the exact instability zones vary depending on stellar mass, disk mass, and the magnitude of α (Pfeil & Klahr, 2019). Convection is facilitated through opacity transitions in these models, in particular the change in opacity due to the melting of ice grains (at $T \sim 160$ K) and evaporation of metal/silicate grains (at $T \sim 1000$ K). Irradiation of the disk by the central star might mitigate convection because heating of the disk surface can result in the disk having an inverted temperature profile (i.e. hot at the surface, cool at the mid-plane) which is not conducive to convection. D’Alessio et al. (1998) performed calculations that included contributions to heating both from irradiation and from viscous dissipation via an α viscosity, and determined that the α viscosity dominated heating only in the very innermost regions of the disk ($r < 2$ AU).

Of course the aforementioned models all simply prescribe the heat source via an α viscosity. Thus identifying a physical mid-plane heat source that could drive convection in protoplanetary disks remains an outstanding problem. Disks around young stars are typically cold, and insufficiently ionized to support the MRI. Nevertheless there are various heat sources that can, in principle, drive convection in these systems. Early work by Lin & Papaloizou (1980) considered very young (Class 0) protostellar disks that were still contracting towards the mid-plane. They demonstrated that such disks are inherently convectively unstable provided that the opacity is a sufficiently steep power law function of the temperature.²³ In the later stages of protoplanetary disks a potential heat source for convection could be dissipation of shocks launched by a high mass planet (Boley & Durisen, 2006; Lyra et al., 2016), or dissipation of zonal magnetic fields facilitated through ambipolar or Ohmic diffusion in mid-plane current sheets (Lesur et al., 2014;

²³An example of a sufficiently steep opacity law relevant to the primordial solar nebula is the melting of ice grains at temperatures of around 160 K.

B  thune et al., 2017; B  thune & Latter, 2020). Finally FU Orionis are a special subclass of protoplanetary disk that undergo outbursts. It is thought that these outbursts are similar to those in dwarf novae, with the gas within 1 AU of the star sufficiently hot that it is fully ionized, leading to MRI turbulence which ultimately dissipates heat around an optically thick mid-plane. Further away from the mid-plane the disk is cool and only partially ionized. Just like in dwarf novae systems this can set up an unstable entropy gradient that can drive convection (Hirose, 2015).

1.5 Dynamics of convection in disks

The papers we discussed in the previous section considered convective instability in various disk classes, but generally did not touch on the actual character of convective flows in disks. In this section we look at the problem from a slightly different angle, focusing on those works that investigated the dynamics of convection in disks more closely, in particular its relation to angular momentum transport and to the MRI. In the spirit of how the subject is presented in the results chapters of this thesis, we have divided this into a discussion of previous work on purely hydrodynamic convection in disks (Section 1.5.1), and a discussion of MHD convection in disks (Section 1.5.2).

1.5.1 Hydrodynamic convection in disks

As mentioned in Section 1.4.3, accretion may instigate convection. But convection might drive accretion itself. Motivated by considerations of the primordial solar nebula, Lin & Papaloizou (1980) constructed a simple hydrodynamic model of a cooling young protoplanetary disk that was still contracting towards the mid-plane. As we already mentioned earlier, they showed that such models are convectively unstable provided that the opacity increases sufficiently steeply with temperature. The heat source is initially provided by gravitational contraction. But, as Lin & Papaloizou (1980) pointed out, if convection is modeled, via a mixing length theory, with an effective viscosity, the process might *self-sustain*, i.e. convective eddies could extract

energy from the background orbital shear, and viscous dissipation of that energy might replace the initial gravitational contraction as a source of heat for maintaining convection.

In order to go beyond mixing length theory, Ruden et al. (1988) analysed linear axisymmetric modes in a thin, polytropic disk in the shearing box approximation and estimated that mixing of gas within convective eddies could result in values of $\alpha \sim 10^{-3} - 10^{-2}$. They cautioned, however, that axisymmetric convective cells could not by themselves exchange angular momentum: such an exchange would have to be facilitated either by non-axisymmetric modes, or by viscous dissipation of axisymmetric modes.

Following these early investigations, a debate ensued that centered not so much on the size of α but rather on its sign.²⁴ Dissipation of *non-linear axisymmetric* convective cells was investigated by Kley et al. (1993) who performed quasi-global simulations (spanning about 100 stellar radii) of axisymmetric disks with alpha viscosity heating and radiative cooling. Their aim was to measure the contribution of convection to the angular momentum (AM) flux, which of course would be in addition to the background AM flux facilitated by the explicit viscosity they included. They measured an inward flux of angular momentum, but warned that this might be due to the imposed axial symmetry and to their relatively high viscosity. Convective shearing waves were first investigated by Ryu & Goodman (1992), who concluded that linear *non-axisymmetric* perturbations would result in a net *inward* angular momentum flux at sufficiently large time. These results were questioned however by Lin et al. (1993), who examined analytically and numerically a set of *localized* linear non-axisymmetric disturbances in global geometry. They demonstrated that these modes could transport angular momentum outwards in some cases, and opined that the inward transport of angular momentum reported by Ryu & Goodman (1992) was an artifact of the shearing box approximation.

²⁴Interestingly the prevailing assumption among researchers at the time was that convection existed in protoplanetary disks, and thus discussion focused on whether it transported angular momentum inwards or outwards, rather than on whether a viable heat source that could sustain convection existed in the first place (Jim Stone; private communication).

Interest in hydro convection waned after local non-linear 3D compressible simulations showed that it resulted in *inward* rather than outward angular momentum transport. Using the finite-difference code ZEUS and rigid, isothermal vertical boundaries, Stone & Balbus (1996) (hereafter SB96) initialized inviscid, fully compressible, and vertically stratified shearing box simulations with a convectively unstable vertical temperature profile, and measured a time-averaged value of $\alpha \sim -4.2 \times 10^{-5}$. Inward angular momentum transport was also reported by Cabot (1996) who ran simulations similar to those of SB96 but included full radiative transfer and a relatively high explicit viscosity. Analytical arguments for the inward transport of angular momentum by convection were presented by SB96 which, crucially, assumed axisymmetry, especially in the pressure field. However, in a rarely cited paper Klahr et al. (1999) presented fully compressible, three-dimensional, global simulations including explicit viscosity and radiative transfer of hydrodynamic convection in disks that gave some indication that non-linear convection in disks actually assumes non-axisymmetric patterns.

Some fifteen years later, the claims of SB96 and of Cabot (1996) were called into question: fully local shearing box simulations of Boussinesq hydrodynamic convection in disks using the spectral code SNOOPY and employing explicit diffusion coefficients indicated that the sign of angular momentum transport due to vertical convection can be reversed provided the Rayleigh number (the ratio of buoyancy to kinematic viscosity and thermal diffusivity) is sufficiently large (Lesur & Ogilvie, 2010). It would appear then that vertical convection can possibly drive outward angular momentum transport after all.

1.5.2 MHD convection in disks

Although convection was often discussed in the literature in the context of models of dwarf novae disks (which should be sufficiently ionized to support the MRI, at least in the high state), these models were purely hydrodynamic in nature and modeled turbulent transport using an α viscosity (see Section 1.4.1). Only in the past few years has convection been investigated in MHD

simulations of disks. A key claim that has emerged from these studies is that the interplay between convection and MRI can enhance angular momentum transport. Bodo et al. (2012) performed 3D fully compressible, vertically stratified shearing box simulations with zero net magnetic flux (ZNF) in PLUTO with explicit thermal diffusivity and a perfect gas equation of state, and found that in regimes with low thermal diffusivity ($\chi \sim 10^{-3} - 10^{-4}$) convection resulted in a significant increase in the magnetic stress. A follow-up paper by Gressel (2013), however, countered that the enhancement in stress observed in Bodo et al. (2012) was in fact an artifact of the impenetrable stress-free boundary conditions employed in their simulations.

More extensive 3D MHD shearing box simulations using the finite-difference code ZEUS were later carried out to study the interplay between convection and the MRI in vertically stratified shearing boxes in more detail, with particular applications to the hydrogen-dominated disks found in dwarf novae (Hirose et al., 2014) and FU Orionis (Hirose, 2015). These simulations included radiative transfer through flux-limited diffusion (FLD), as well as opacity tables, a temperature- and density-dependent ratio of specific heats, and a non-ideal equation of state. Hirose et al. (2014) observed convective/radiative cycles on the upper branch of their simulations, and claimed that convection on the lower-end of the upper branch of the S-curve could enhance α by as much as an order of magnitude. Coleman et al. (2016) attempted to model dwarf novae outbursts in 1D disk models by incorporating the variation in α measured in the aforementioned simulations. Although their models were consistent with observed outburst and quiescent durations, they exhibited short bursts in magnitude during decay from outburst (known as reflares) that have not been observed in real dwarf novae. Using the data from Hirose et al. (2014) simulations, Coleman et al. (2017) investigated in more detail the effect of convection on the MRI dynamo, and found that convection in hydrogen-dominated disks could quench the toroidal field reversals observed in MRI simulations. Application to the helium-dominated disks in AM CVns exhibited a similar enhancement of α on the lower-end of the upper branch, but resulted in persistent rather than intermittent convection (Coleman et al., 2018). Note that many of these works used the sign

of the square of buoyancy frequency N_B^2 to determine whether convection was present in their simulations or not, a criterion that is necessary but not sufficient. We shall discuss this issue in more detail in Section 4.1.2.

The results of Hirose et al. (2014) – in particular intermittent convection and enhancement of α on the lower tip of the upper branch – were corroborated with 3D MHD flux-limited diffusion shearing box simulations carried out by Scepi et al. (2018a) using the finite-volume code PLUTO. Scepi et al. (2018a) also observed convection in hotter regions of the lower branch (which they referred to as the middle branch), but *without* an accompanying enhancement in α , thus calling into question the role of convection in enhancing α .

Furthermore, when Scepi et al. (2018b) investigated the effects of a net vertical magnetic field (NVF) configuration, they found that, in the absence of resistivity, this resulted in an α which was higher on the *lower* branch than in outburst. Scepi et al. (2018b) also found that with a NVF configuration *wind*-driven accretion dominated turbulent-driven accretion on the lower branch, and in a follow-up paper demonstrated that wind-driven accretion could reproduce dwarf nova light curves (Scepi et al., 2019). Thus Scepi et al. (2019) proposed wind-driven accretion as an alternative hypothesis to MRI/convection-driven turbulence as the mechanism for angular momentum transport in dwarf novae.

1.6 Outline of thesis

In Chapter 2 we discuss the methods that we employ throughout this thesis. We begin by presenting the governing equations (Section 2.1). All of the results in this thesis were obtained by considering a local region of an accretion disk, known as the shearing box formalism, and thus we review this formalism in Section 2.2. We then describe the numerical algorithms (Section 2.3), and astrophysical fluid dynamics codes that we use for our simulations (Section 2.4). Finally we define the various diagnostics that we use in our results chapters (Section 2.5).

In Chapter 3 we discuss hydrodynamic convection in disks. We begin by presenting our numerical set-up and equilibria in Section 3.1. In Section 3.2 we examine the linear theory of hydrodynamic convection in disks, including calculations of the linear growth rates. The non-linear phase is then investigated in Section 3.3 in unforced simulations, where we find that hydrodynamic convection can, in general, transport angular momentum outward. In Section 3.4 we discuss simulations in which convection is sustained using an artificial heat source, leading to large-scale oscillatory convective cells. We discuss our results and conclusions in Section 3.5.

In Chapter 4 we turn our attention to magnetohydrodynamic convection in disks. In Section 4.2 we discuss stratified MHD simulations without explicit cooling, with the aim of investigating the influence of the MRI on the vertical structure of the disk in the absence of convection. Next, in order to facilitate convection, we implement an explicit cooling term (Section 4.3), and find that the interplay between convection and the MRI can be manifested in various distinct ways. We discuss our results in Section 4.4, including how they compare to previous work and their application to dwarf novae. Finally, we summarize our key results in Section 4.5.

In Chapter 5 (which is largely independent of the first two), we discuss the stress-pressure relationship in disks. Pressure and stress are known to interact in a non-trivial way, certainly more complicated than the simple linear behavior posited by the alpha model, and this can have implications for thermal instability, which we discuss in Section 5.1. In Section 5.2 we provide a brief overview of our numerical set-up and cooling prescription. In Section 5.3 we examine the time-lag between stress and pressure in unstratified simulations by controlling the mean pressure using our cooling prescription. Finally, we conclude the thesis in Chapter 6, where we summarize our key results and discuss avenues for future work.

Chapter 2

Methods

In this section we will discuss the physical models and numerical methods used in later chapters. We start by presenting the governing equations and parameters in Section 2.1. These equations are solved in the local reference frame of an observer co-rotating at some fiducial radius in a disk. This is known as the shearing box approximation, which is presented in Section 2.2. In Section 2.3 we discuss the specific numerical algorithms used to solve the equations, in particular the technique of finite volume methods. We briefly present the codes that we use in Section 2.4. Finally, we discuss the various diagnostics for analyzing data from our simulations in Section 2.5.

2.1 Governing equations

We model the disk as a fluid, an approximation that is valid provided that the collision rate ω_c of gaseous particles in the disk is much greater than the orbital angular frequency Ω .¹ This assumption generally holds in dwarf novae, in protoplanetary discs, and in the outer regions of X-ray binaries and active galactic nuclei (Latter et al., 2017). Thus we work with the equations of magnetized gas dynamics (or of magnetohydrodynamics). These are given

¹An equivalent statement of the validity of the fluid approximation is that the Knudsen number, i.e. ratio of the mean free path λ of particles in the disk to the disk scale height H , is much less than unity, i.e. $\text{Kn}=\lambda/H \ll 1$.

by the continuity equation (corresponding to conservation of mass)

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (2.1)$$

the momentum equation (corresponding to conservation of momentum)

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho} \nabla P - \nabla \Phi + \nabla \cdot \mathbf{T} + \frac{1}{\mu_0 \rho} (\nabla \times \mathbf{B}) \times \mathbf{B}, \quad (2.2)$$

the thermal energy equation (corresponding to conservation of thermal energy)

$$\partial_t (\rho e) + \mathbf{u} \cdot \nabla (\rho e) = -\gamma \rho e \nabla \cdot \mathbf{u} + \mathbf{T} : \nabla \mathbf{u} + \kappa \nabla^2 T + \frac{\eta}{\mu_0} |\nabla \times \mathbf{B}|^2 + \Lambda, \quad (2.3)$$

and the induction equation governing the time evolution of the magnetic field

$$\partial_t \mathbf{B} = \nabla \times (\mathbf{u} \times \mathbf{B}) + \eta \nabla^2 \mathbf{B}. \quad (2.4)$$

These equations are closed with the caloric equation of state for a perfect gas $P = e(\gamma - 1)\rho$ where e is the specific internal energy. The adiabatic index (ratio of specific heats) is denoted by γ and is taken to be $\gamma = 5/3$ in all our simulations. The temperature T , when needed, is recovered using the thermal equation of state for a perfect gas $P = (\mathcal{R}/\mu)\rho T$. Here P is the thermal pressure of the fluid, ρ is mass density, \mathcal{R} is the gas constant, and μ is the mean molecular weight. In our simulations we set $\mu = 0.5$, which is appropriate for fully ionized hydrogen. Finally, in a few select simulations we keep the temperature of the gas constant, i.e. the equation of state is $P = c_s^2 \rho$ where the (isothermal) sound speed c_s is constant in both space and time. We refer to this as an *isothermal* equation state.

$\Phi(r, z)$ is the external gravitational potential due to a point source of mass M , i.e. $\Phi = -GM/\sqrt{r^2 + z^2}$ in cylindrical coordinates. In a non-inertial (i.e. rotating) frame Φ can be combined with the centrifugal potential to form an effective gravitational potential Φ_{eff} . Expansion of Φ_{eff} in a small parameter (specifically the ratio of disk thickness to cylindrical radius) results in an effective gravitational acceleration that depends only linearly on the radial and vertical coordinates. This simplification (together with appropriate boundary conditions) is known as the *shearing box approximation* (see Section 2.2).

Other terms in the equations include the viscous stress tensor $\mathbf{T} \equiv 2\rho\nu\mathbf{S}$. Here ν is the kinematic viscosity, and $\mathbf{S} \equiv (1/2)[\nabla\mathbf{u} + (\nabla\mathbf{u})^T] - (1/3)(\nabla\cdot\mathbf{u})\mathbf{I}$ is the traceless shear tensor, and we omit the bulk viscosity. The thermal conductivity is denoted by κ , but in our simulations we specify thermal diffusivity χ rather than κ . The former is related to the latter via $\chi = \kappa/(c_p\rho)$ where c_p is the specific heat capacity at constant pressure. The explicit magnetic diffusivity, where this is included, is denoted by η , and the permeability of free space is denoted by μ_0 . The magnetic diffusivity is related to the resistivity σ (not to be confused with the symbol we use for the growth rate later on) by $\sigma = 1/\mu_0\eta$, though we often refer to the magnetic diffusivity as a ‘resistivity’ later in this thesis. Note that where we employ explicit diffusion coefficients, for simplicity we treat these as being constant in both space and time. Thus the diffusion coefficients can be taken outside the derivatives in the governing equations.

In some of our hydrodynamic simulations we mock up the effects of radiative physics through a *thermal relaxation* term $\Lambda = \Lambda_{\text{relax}} \equiv -(\rho e - \rho_0 e_0)/\tau_{\text{relax}}$. This relaxes the thermal energy in each cell back to its initial state $e_0(z) \equiv e(z, t = 0)$ on a timescale given by τ_{relax} , where the initial state is a convectively unstable state that might be the result of radiative cooling balanced by internal or external heating. We refer to simulations run with $\Lambda_{\text{relax}} \neq 0$ as simulations of *forced* compressible convection. Conversely, simulations with $\Lambda_{\text{relax}} = 0$ are referred to as *unforced*. Since convection moves heat away from the mid-plane in our simulations, we choose the thermal relaxation timescale τ_{relax} to be the inverse of the linear growth rate of the convective instability (which we derive in Section 3.2).

In our MHD simulations heating is provided self-consistently through an internal heat source, thus we use a slightly simpler cooling prescription compared to that employed in our hydro simulations. We implement optically thin cooling through a linear cooling term (also known as *beta-cooling*) $\Lambda = \Lambda_c \equiv -\rho e/\tau_c$. This lowers the thermal energy in each cell on a cooling timescale given by τ_c . In fully turbulent MRI simulations this cooling is balanced by heating due to the dissipation of MRI turbulence (and resistive or Joule-heating, when explicit resistivity is included). In all but one of




Figure 2.1: Schematic of the shearing box, adapted from Mignone (2015). The inset shows a global representation of the shearing box embedded within a disk. The middle (bolded) box represents the active computational domain, while the two boxes on either side represent image boxes (used to determine the radial boundary conditions) that move at a constant speed $w = \pm q\Omega_0 L_x$ relative to the active domain.

the MHD simulations in which we employ explicit cooling, we use a height-dependent cooling prescription to mock up the effects of an optically thick mid-plane surrounded by an optically thin corona: this is facilitated by turning cooling on for $|z| > 0.75 H_0$. Note that in those regions where there is explicit cooling, the cooling timescale is constant both in space and time.

2.2 Shearing box formalism

2.2.1 Tidal expansion of effective gravitational potential

Equations 2.1-2.4 govern the motion and thermodynamics of a fluid element under the influence of an external gravitational potential Φ . A simplification

can be made to the momentum equation (Equation 2.2) by working in the shearing box approximation (Goldreich & Lynden-Bell, 1965; Hawley et al., 1995a; Umurhan & Regev, 2004; Latter & Papaloizou, 2017), which treats a local region of a disk as a Cartesian box centered at some fiducial radius $r = r_0$ and at the disk mid-plane plane $z = 0$, and orbiting with the angular frequency of the disk at that radius $\Omega_0 \equiv \Omega(r_0)$ (see Figure 2.1). A point in the box has Cartesian coordinates (x, y, z) which are related to the cylindrical coordinates (r, ϕ, z) through $x = r - r_0$, $y = r_0(\phi - \phi_0 - \Omega_0 t)$ and $z = z$.

To derive the shearing box equations, first consider a rotating (non-inertial) reference frame. In this frame terms due to the Coriolis and centrifugal forces are added to the right-hand side of the momentum equation

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = \dots - \nabla \Phi - 2\mathbf{\Omega} \times \mathbf{u} - \mathbf{\Omega} \times (\mathbf{\Omega} \times \mathbf{r}). \quad (2.5)$$

Here we have assumed that the angular frequency of the reference frame is equal to the local angular frequency Ω_0 of the disk at $r = r_0$, i.e. the angular frequency is given by $\mathbf{\Omega} = (0, 0, \Omega_0)^T$, where Ω_0 is constant. Thus we may rewrite the centrifugal term as the gradient of $\mathbf{\Omega} \times \mathbf{r}$ to obtain

$$\mathbf{\Omega} \times (\mathbf{\Omega} \times \mathbf{r}) = -\frac{1}{2} \nabla |\mathbf{\Omega} \times \mathbf{r}|^2, \quad (2.6)$$

which can then be combined together with the gradient of the external gravitational potential Φ to form a term corresponding to the gradient of an *effective* gravitational potential, i.e.

$$-\nabla \Phi + \frac{1}{2} \nabla |\mathbf{\Omega} \times \mathbf{r}|^2 = -\nabla \left[\Phi - \frac{1}{2} |\mathbf{\Omega} \times \mathbf{r}|^2 \right] \equiv -\nabla \Phi_{\text{eff}}. \quad (2.7)$$

Evaluating the cross product term $(1/2)|\mathbf{\Omega} \times \mathbf{r}|^2$, we find that the effective gravitational potential is given by

$$\Phi_{\text{eff}} = \Phi - \frac{1}{2} r^2 \Omega_0^2. \quad (2.8)$$

Expanding Φ_{eff} about some fiducial point ($r = r_0, z = 0$) in the disk to second order (and assuming that $x \equiv r - r_0 \ll 1$ and that $|z/r_0| \ll 1$) we obtain

$$\begin{aligned}
\Phi_{\text{eff}} = & \Phi(r_0, 0) + \left. \frac{\partial \Phi}{\partial r} \right|_{r_0, 0} x + \left. \frac{\partial \Phi}{\partial z} \right|_{r_0, 0} z \\
& + \frac{1}{2} \left. \frac{\partial^2 \Phi}{\partial r^2} \right|_{r_0, 0} x^2 + \frac{1}{2} \left. \frac{\partial^2 \Phi}{\partial z^2} \right|_{r_0, 0} z^2 + \left. \frac{\partial^2 \Phi}{\partial r \partial z} \right|_{r_0, 0} rz \\
& + \mathcal{O}(r_0^3, z^3) \\
& - \frac{1}{2} \Omega_0^2 (r_0^2 + 2r_0 x + x^2).
\end{aligned} \tag{2.9}$$

Note that in expanding Φ_{eff} we have implicitly assumed that $|z/r_0| \ll 1$, i.e. that the disk is *thin*. Furthermore, in Equation 2.9, the leading order gravitational and centrifugal terms $\Phi(r_0, 0) = -GM/r_0$ and $-(1/2)\Omega_0^2 r_0^2$, respectively, are constants and vanish on taking the derivative. Evaluating the remaining terms in Equation 2.9 we obtain the tidal expansion of the effective gravitational potential

$$\Phi_{\text{eff}} = -\Omega_0^2 S_0 x^2 + \frac{1}{2} \Omega_0^2 z^2, \tag{2.10}$$

where we have dropped the constant terms. Here $S_0 \equiv -rd\Omega/dr|_{r=r_0}$ is the *shear rate* evaluated at the point $r = r_0$.

The governing equations still involve gradient operators ∇ that are in cylindrical polar coordinates. Assuming that any perturbations of interest vary on length scales much smaller than the fiducial radius r_0 , we can adjust the size of the shearing box (and thus the extent of the shearing box coordinates x, y , and z) accordingly: therefore $x, y, z \ll r_0$. This allows us to drop subdominant cylindrical terms (which go as $x/r_0, y/r_0, z/r_0$) in the gradient operators. Thus $\nabla_{\text{cylindrical}} \sim \nabla_{\text{Cartesian}}$. Making this replacement and using the tidal expansion of the effective gravitational potential given by Equation 2.10, the momentum equation in the shearing box approximation becomes

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho} \nabla P + \mathbf{g}_{\text{eff}} - 2\mathbf{\Omega} \times \mathbf{u} + \nabla \cdot \mathbf{T} + \frac{1}{\mu_0 \rho} (\nabla \times \mathbf{B}) \times \mathbf{B}, \tag{2.11}$$

where all gradient operators in Equation 2.11 are now in *Cartesian* coordinates. Here the effective gravitational potential is embodied in the tidal

acceleration $\mathbf{g}_{\text{eff}} = 2q\Omega_0^2 x \hat{\mathbf{x}} - \Omega_0^2 z \hat{\mathbf{z}}$, where $q \equiv -d \ln \Omega / d \ln r$ is the (dimensionless) *shear parameter*. For Keplerian disks $q = 3/2$, a value we adopt throughout this paper.

To complete the shearing box approximation we require that the momentum equation with modified gravitational acceleration together with the usual equations for conservation of mass (Equation 2.1), conservation of energy (Equation 2.3), and the induction equation (Equation 2.4) are complemented by appropriate boundary conditions (see, for example, Hawley et al. (1995a)). We assume the shearing box (which constitutes the *active* domain, i.e. the domain in which the solution of the governing equations is updated in time) is accompanied by two *image* boxes at each radial boundary (see Figure 2.1). Due to the linear shear, these image boxes (whose purpose is purely to implement the radial boundary conditions) appear to slide past the active domain relative to an observer located at the center of the active domain. A fluid element leaving the active domain at the left-hand radial boundary at the azimuthal position y will reappear at the right-hand radial boundary at the azimuthal position of the right-hand image box $y - q\Omega L_x t$. To quantify this, let the radial, azimuthal, and vertical boundaries be located at $x = \pm L_x/2$, $y = \pm L_y$, and $z = \pm L_z$, respectively. For any flow variable f (except the y -component of the velocity u_y which we shall discuss shortly) the boundary conditions are given by

$$f(x, y, z) = f(x + L_x, y - q\Omega L_x t, z) \quad (x \text{ boundary}), \quad (2.12)$$

$$f(x, y, z) = f(x, y + L_y, z) \quad (y \text{ boundary}), \quad (2.13)$$

at the x - and z -boundaries, respectively. For the y -component of the velocity u_y we use periodic boundary conditions in the y -direction (cf. Equation 2.13), but the radial boundary condition is given by

$$u_y(x, y, z) = u_y(x + L_x, y - q\Omega L_x t, z) + q\Omega L_x \quad (x \text{ boundary}). \quad (2.14)$$

Thus in the shearing box approximation we are restricted to using so called shear-periodic boundary conditions in the x -direction and periodic boundary

conditions in the y -direction. (We are free to choose whatever boundary conditions we wish in the z -direction.)

2.2.2 Conservation of energy in the shearing box

By dotting the momentum equation (Equation 2.11) with \mathbf{u} , combining the result with the continuity equation (Equation 2.1), thermal energy equation (Equation 2.3), and induction equation (Equation 2.4), and integrating over the volume of the box we can synthesize an equation for the volume-integrated total energy density E in the shearing box

$$\frac{\partial}{\partial t} \int_V E dV = - \int_V \nabla \cdot \mathcal{F} dV + \int_V \mathbf{u} \cdot (\rho \mathbf{g}_{\text{eff}}) dV, \quad (2.15)$$

where the total energy density is defined as

$$E \equiv \frac{1}{2} \rho u^2 + \frac{1}{2\mu_0} B^2 + \rho e, \quad (2.16)$$

and the flux of total energy density²

$$\mathcal{F} \equiv \left[\mathbf{u}(E + P_t) - \frac{1}{\mu_0} (\mathbf{u} \cdot \mathbf{B}) \mathbf{B} \right], \quad (2.17)$$

where the total pressure is given by $P_t = P + B^2/2\mu_0$. Using the divergence theorem to convert the first term on the right-hand-side into a flux through the boundaries we obtain

$$\frac{\partial}{\partial t} \int_V E dV = - \int_S \mathcal{F} \cdot d\mathbf{S} + \int_V \mathbf{u} \cdot (\rho \mathbf{g}_{\text{eff}}) dV, \quad (2.18)$$

where the surface integral is over the six faces of the (Cartesian) shearing box.

To evaluate the surface integral, let us adopt periodic boundary conditions at both the y - and z -directions. Thus the integrated flux through the y -boundary oriented along the positive y -axis is equal and opposite to the

²For simplicity we have dropped diffusive terms (which can be incorporated as part of the total flux) as well as the cooling term Λ (which, had we kept it, would appear as a source term on the right-hand-side of Equation 2.15.)

integrated flux through the opposite y -boundary, and likewise for the integrated flux through the z -boundaries. All surface integrals except those over the x -faces vanish, and the flux term reduces to

$$\int_S \mathcal{F} \cdot d\mathbf{S} = \int_X \mathcal{F}_x dydz, \quad (2.19)$$

where \mathcal{F}_x is the radial component of the flux vector defined in Equation 2.17 evaluated at $x = L_x/2$ and $x = -L_x/2$ at the right- and left-hand boundaries, respectively. Applying the x -boundary conditions and substituting $\delta u_y = q\Omega_0 x + u_y$ we obtain

$$\int_X \mathcal{F}_x dydz = -q\Omega_0 L_x \int_X \left[\rho u_x \delta u_y - \frac{1}{\mu_0} B_x B_y \right] dydz. \quad (2.20)$$

Finally, substituting Equation 2.20 into Equation 2.18 we obtain the following equation for the time-evolution of the volume-integrated total energy density

$$\frac{\partial}{\partial t} \int_V E dV = q\Omega_0 L_x \int_X \left[\rho u_x \delta u_y - \frac{1}{\mu_0} B_x B_y \right] dydz + \int_V \mathbf{u} \cdot (\rho \mathbf{g}_{\text{eff}}) dV. \quad (2.21)$$

Thus total energy is *not* conserved in the shearing box: flux of angular momentum through the x -boundaries due to the Reynolds and magnetic stresses act as sources for the total energy.

Note that in our code (see Section 2.4) the total energy density is defined so as to omit contributions from the gravitational energy density $\rho\Phi_{\text{eff}}$. In this formulation, work done by the tidal expansion of the gravitational potential also acts as a source term in the total energy equation (Mignone et al., 2012). It is possible, however, to avoid this by redefining the total energy density such that it includes the contribution due to the effective gravity, i.e. $\tilde{E} \equiv E + \rho\Phi_{\text{eff}}$. By rewriting the work done by gravity in terms of the effective gravitational potential

$$\mathbf{u} \cdot (\rho \mathbf{g}_{\text{eff}}) = -\rho \mathbf{u} \cdot \nabla \Phi_{\text{eff}}, \quad (2.22)$$

then using the identity

$$\rho \mathbf{u} \cdot \nabla \Phi_{\text{eff}} = \nabla \cdot (\rho \mathbf{u} \Phi_{\text{eff}}) - \Phi_{\text{eff}} \nabla \cdot (\rho \mathbf{u}), \quad (2.23)$$

and, finally, using the continuity equation to rewrite the second term on the right-hand-side as $-\Phi_{\text{eff}}\nabla \cdot (\rho\mathbf{u}) = \partial_t(\rho\Phi_{\text{eff}})$ (note that Φ_{eff} is time-independent) we obtain the following equation for the time-evolution of the total energy density \tilde{E}

$$\frac{\partial}{\partial t} \int_V \tilde{E} dV = q\Omega_0 L_x \int_X \left[\rho u_x \delta u_y - \frac{1}{\mu_0} B_x B_y \right] dy dz, \quad (2.24)$$

where the total energy density $\tilde{E} \equiv E + \rho\Phi_{\text{eff}}$ now includes contributions due to the effective gravitational potential and the gravitational energy no longer appears as a source term in the equation.

2.3 Numerical Algorithms

2.3.1 Finite-volume methods

In order to solve the compressible MHD equations in the shearing box approximation numerically, we employ finite-volume methods. The basic idea behind this algorithm (strictly speaking a set of algorithms) is to discretize the governing equations in conservative form (see Sections 2.3.2 and 2.3.3) such that we evolve volume-averaged quantities (integral equations) rather than cell-centered quantities (i.e. partial differential equations, as in finite-difference methods). Details can be found in the monographs by Laney (1998); LeVeque (2002) and Toro (2013). Therefore each discrete point corresponds to a volume-average over a cell.

There are several advantages of using finite-volume methods (FVM) over finite difference methods (FDM). As we shall demonstrate shortly, one advantage is that no truncation errors are introduced in discretizing the governing equations. Second, because FVM evolve integrals rather than PDEs, they overcome the problem of having to calculate the derivative of a discontinuous quantity, and therefore they are much more suitable for capturing discrete jumps in fluid variables (i.e. *shocks*) compared to FDM. Finally because FVM inherently work with the conservative form of the equations, conservation of mass, momentum, and total energy is automatically guaranteed




Figure 2.2: Top: schematic of the numerical grid in one dimension. Bottom: schematic of the reconstruct-solve-average strategy. Note that some RSA methods, including that employed by our code (see Section 2.4), perform reconstruction in primitive variables $\mathbf{V} = (\rho, \mathbf{u}, P, \mathbf{B})^T$ rather than conservative variables $\mathbf{U} = (\rho, \rho \mathbf{u}, E, \mathbf{B})^T$.

unless there are explicit source terms (note, however, that conservative finite difference methods also exist, see, for example, Hirose et al. (2006)).

Most finite-volume methods employ a three-step reconstruct-solve-average (RSA) strategy to update the solution array in each cell (see, for example, (Mignone et al., 2007) and Figure 2.2). First (step 1), cell-centered volume-averages are interpolated to obtain the value of each fluid variable at each cell interface. This results in two discrete states of initial data on either side of each cell interface, known as a *Riemann problem*. The total flux through each interface is then obtained by solving this Riemann problem (step 2). Finally (step 3), the fluxes are used together with a time-stepping algorithm to update the solution in each cell. We will discuss each of these three steps in greater detail in Sections 2.3.4, 2.3.5, and 2.3.6, respectively.

2.3.2 Conservative form of the equations

Finite volume-methods solve the governing equations not in the form given by Equations 2.1-2.4 (known as *primitive form*, where the primitive variables are given by $\mathbf{V} = (\rho, \mathbf{u}, P, \mathbf{B})^T$), but in *conservative form* in which the equations are rewritten as a set of conservation laws given by

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U}), \quad (2.25)$$

where $\mathbf{U} = (\rho, \rho \mathbf{u}, E, \mathbf{B})^T$ is a vector of conservative variables (i.e. mass density, momentum density, total energy density, and magnetic field, respectively). Here $\mathbf{F}(\mathbf{U})$ is an 8×1 matrix, the rows of which correspond to the fluxes of each conservative variable. Written out in full

$$\mathbf{F}(\mathbf{U}) = \begin{pmatrix} \rho \mathbf{u} \\ \rho \mathbf{u} \mathbf{u} - \frac{1}{\mu_0} \mathbf{B} \mathbf{B} + \mathbf{1} P_t - \mathbf{T} \\ (E + P_t) \mathbf{u} - \frac{1}{\mu_0} \mathbf{B} (\mathbf{u} \cdot \mathbf{B}) - \mathbf{u} \cdot \mathbf{T} - \kappa \nabla T + \frac{\eta}{\mu_0} (\nabla \times \mathbf{B}) \times \mathbf{B} \\ \mathbf{u} \mathbf{B} - \mathbf{B} \mathbf{u} - \eta \nabla \mathbf{B} \end{pmatrix} \quad (2.26)$$

The terms in black in Equation 2.26 are known as *hyperbolic* terms, and correspond to the transport (or advection) of mass, momentum, and total energy flux. These terms are associated with wave propagation at finite speed. The terms in red in Equation 2.26 are known as *parabolic* terms, and correspond to diffusive processes, such as viscosity, thermal diffusivity, and resistivity. Thus the equations of MHD are of mixed hyperbolic/parabolic type.

Finally source terms are embodied in the vector $\mathbf{S}(\mathbf{U})$ which is given by

$$\mathbf{S}(\mathbf{U}) = \begin{pmatrix} 0 \\ \rho \mathbf{g}_{\text{eff}} - 2\Omega_0 \hat{\mathbf{z}} \times \rho \mathbf{u} \\ \rho \mathbf{u} \cdot \mathbf{g}_{\text{eff}} + \Lambda \\ 0 \end{pmatrix}. \quad (2.27)$$

These comprise terms due to external body forces (such as the effective gravitational force), the Coriolis force (due to rotation), and radiative processes (such as the cooling term Λ).

2.3.3 Discretization of governing equations

For simplicity, consider the set of conservation laws given in Equation 2.25 in one dimension and without source terms, i.e.

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} = 0. \quad (2.28)$$

As mentioned in Section 2.3.1, finite volume methods evolve volume-averages of fluid quantities over discrete domains known as cells. In 1D, a cell is defined as occupying the spatial extent $x \in \{x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}\}$ with x_i being the cell-center and $x_{i\pm\frac{1}{2}}$ the right- and left-hand cell-interfaces, respectively. Likewise we discretize time into distinct time-levels with t^n denoting the n th time-level, t^{n+1} denoting the $(n+1)$ th time-level, and so on. The time interval between two subsequent time-levels is denoted $\Delta t \equiv t^{n+1} - t^n$ and defines a single time-step.

Integrating the differential form of the equations over space and time we obtain

$$\int_{t^n}^{t^{n+1}} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \left(\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} \right) dx dt = 0, \quad (2.29)$$

where the time-integral is taken over a time-interval t^n to t^{n+1} and the spatial integral is taken over the width of the cell. Integrating over space

$$\int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \left(\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} \right) dx = \frac{\partial}{\partial t} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \mathbf{U} dx + \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \frac{\partial \mathbf{F}}{\partial x} dx \quad (2.30)$$

$$= \Delta x \frac{\partial \langle \mathbf{U} \rangle_i}{\partial t} + \left(\mathbf{F}_{i+\frac{1}{2}} - \mathbf{F}_{i-\frac{1}{2}} \right), \quad (2.31)$$

where $\langle \mathbf{U} \rangle_i$ is the spatial-average of the state vector \mathbf{U} over the cell centered at x_i

$$\langle \mathbf{U} \rangle_i \equiv \frac{1}{\Delta x} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \mathbf{U}(x, t) dx. \quad (2.32)$$

Finally, integrating Equation 2.31 over time we obtain

$$\int_{t^n}^{t^{n+1}} \left[\Delta x \frac{\partial \langle \mathbf{U} \rangle_i}{\partial t} + \left(\mathbf{F}_{i+\frac{1}{2}} - \mathbf{F}_{i-\frac{1}{2}} \right) \right] dt = \quad (2.33)$$

$$\Delta x \left(\langle \mathbf{U} \rangle_i^{n+1} - \langle \mathbf{U} \rangle_i^n \right) + \Delta t \left(\tilde{\mathbf{F}}_{i+\frac{1}{2}}^{n+\frac{1}{2}} - \tilde{\mathbf{F}}_{i-\frac{1}{2}}^{n+\frac{1}{2}} \right), \quad (2.34)$$

where $\langle \mathbf{U} \rangle_i^n \equiv \langle \mathbf{U} \rangle_i(t^n)$, and $\tilde{\mathbf{F}}_{i\pm\frac{1}{2}}^{n+\frac{1}{2}}$ is the time-averaged flux vector at the left (right) cell-interfaces

$$\tilde{\mathbf{F}}_{i\pm\frac{1}{2}}^{n+\frac{1}{2}} \equiv \frac{1}{\Delta t} \int_{t^n}^{t^{n+1}} \mathbf{F}_{i\pm\frac{1}{2}} dt, \quad (2.35)$$

where $\mathbf{F}_{i\pm\frac{1}{2}} \equiv \mathbf{F}_{i\pm\frac{1}{2}}(\mathbf{U}(x_{i\pm\frac{1}{2}}))$.

Finally the volume-averaged solution vector at the cell center x_i at the $(n+1)$ th time-step is given in terms of the solution at the n th time-step and the fluxes through the right- and left-interfaces by

$$\langle \mathbf{U} \rangle_i^{n+1} = \langle \mathbf{U} \rangle_i^n - \frac{\Delta t}{\Delta x} \left(\tilde{\mathbf{F}}_{i+\frac{1}{2}}^{n+\frac{1}{2}} - \tilde{\mathbf{F}}_{i-\frac{1}{2}}^{n+\frac{1}{2}} \right). \quad (2.36)$$

An important advantage of finite-volume methods over finite-difference methods is that no truncation errors have been introduced, i.e. Equation 2.36 is an *exact* relation (though, as we shall see, an approximation is usually introduced in evaluating the fluxes). As can be seen from Equation 2.36, an update of the solution in the i th cell (time-stepping) requires knowledge of the fluxes at the left and right cell-interfaces (solving a Riemann problem), which in turn require knowledge of the solution at the current step at those interfaces (reconstruction). The methods for obtaining these quantities are discussed, in reverse order, in the following three sections.

2.3.4 Reconstruction and interpolation methods

The first step of the reconstruct-solve-average strategy involves converting the cell-centered solution $\langle \mathbf{U} \rangle_i^n$ to a solution at the left- and right-interfaces given by $\mathbf{U}_{i-\frac{1}{2}}^n$ and $\mathbf{U}_{i+\frac{1}{2}}^n$, respectively. Various interpolation methods can be used to achieve this. For simplicity we restrict our discussion to one-dimension.

The simplest and most computationally inexpensive case is that of *flat* reconstruction (1st order accurate in space) in which the solution is simply taken to be piecewise constant inside each cell

$$\mathbf{U}^n(x) = \langle \mathbf{U} \rangle_i^n \quad (2.37)$$

A more sophisticated approach which is still fairly computationally inexpensive (and which we frequently employ in our MHD simulations) approximates the solution as being piecewise linear inside each cell. This is known as *linear* reconstruction. Here the slope of the piecewise linear function is computed using the gradient of the function at the cell center. The interpolated solution is given by

$$\mathbf{U}^{n,(k)}(x) = \langle \mathbf{U} \rangle_i^n + \Delta \mathbf{U}_i^{n,(k)} (x - x_i), \quad (2.38)$$

where the slope of the linear equation is denoted by $\Delta \mathbf{U}_i^{n,(k)}$ and $k \in \{F, B\}$. This can be approximated either as a forward ($k = F$) or backward difference ($k = B$), i.e.

$$\Delta \mathbf{U}_i^{n,(F)} \equiv \frac{\langle \mathbf{U} \rangle_{i+1}^n - \langle \mathbf{U} \rangle_i^n}{\Delta x}, \quad \Delta \mathbf{U}_i^{n,(B)} \equiv \frac{\langle \mathbf{U} \rangle_i^n - \langle \mathbf{U} \rangle_{i-1}^n}{\Delta x}. \quad (2.39)$$

A major drawback of this scheme, however, is that if the solution varies very rapidly between cells, such as at a discontinuity, the straight line given by Equation 2.38 will over- or underestimate the solution at the interfaces. This introduces new maxima or minima which result in unphysical oscillations (manifested as a Gibbs-type overshoot at the discontinuity). To rectify this, linear reconstruction methods replace the forward and backward differences given by Equation 2.39 by a non-linear function of the two known as a *slope limiter* $\phi(r)$

$$\mathbf{U}^n(x) = \langle \mathbf{U} \rangle_i^n + \phi(r) \frac{x - x_i}{\Delta x}, \quad (2.40)$$

where $r \equiv \Delta \mathbf{U}_i^{n,(B)} / \Delta \mathbf{U}_i^{n,(F)}$. Although we omit the details here for brevity, a detailed discussion of many different slope limiters can be found in Mignone (2014).

Finally a third reconstruction scheme is the second-order accurate *weighted-essentially-non-oscillatory* (WENO3) scheme (Yamaleev & Carpenter, 2009; Mignone, 2014). This is more computationally expensive than linear reconstruction, and therefore we restrict its use to our hydrodynamic simulations. However WENO3 has higher order spatial accuracy in general than slope-limited linear reconstruction methods, and it retains its accuracy near extrema whereas slope-limiters reduce the accuracy at extrema to first-order (Harten et al., 1986; Shu, 1997). In this scheme reconstruction is carried out by assigning weights to a linear superposition of the right- and left-linearly reconstructed values (i.e. to $\mathbf{U}^{n,(F)}(x)$ and $\mathbf{U}^{n,(B)}(x)$, respectively). Explicitly, the reconstructed solution is given by

$$\mathbf{U}^n(x) = \omega_{i,0}\mathbf{U}^{n,(F)}(x) + \omega_{i,1}\mathbf{U}^{n,(B)}(x). \quad (2.41)$$

The weights $\omega_{i,k}$ (with $k \in \{0, 1\}$) are non-linear functions of the forward and backward differences (see Equation 2.39) that measure the smoothness of the numerical solution on the right- and left-hand stencils $\{i, i+1\}$ and $\{i-1, i\}$, respectively (see Mignone (2014) for further details).

2.3.5 Riemann solvers

Following the interpolation step, the solution is known at either side of each cell-interface. For example, at the right-hand interface of the cell centered at x_i this results in a set of initial data given by

$$\langle \mathbf{U} \rangle_{i+\frac{1}{2}}^n = \begin{cases} \langle \mathbf{U} \rangle_L^n, & (x < x_{i+\frac{1}{2}}), \\ \langle \mathbf{U} \rangle_R^n, & (x > x_{i+\frac{1}{2}}), \end{cases} \quad (2.42)$$

where $\langle \mathbf{U} \rangle_L$ is the solution at the interface $x_{i+\frac{1}{2}}$ obtained by interpolating from the cell center at x_i and $\langle \mathbf{U} \rangle_R$ is the solution at the same interface obtained by interpolating from the neighboring cell center at x_{i+1} . This set of discontinuous initial data together with the conservation laws (Equation 2.28) constitutes a *Riemann problem*. Specific algorithms for solving the Riemann problem are known as *Riemann solvers*.

The solution to the Riemann problem formulated in the previous paragraph consists (in MHD) of seven waves separating eight constant states. To see this it is best to rewrite the set of governing equations (Equation 2.28) in quasi-linear form as

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{A} \cdot \frac{\partial \mathbf{U}}{\partial x} = 0, \quad (2.43)$$

where $\mathbf{A} \equiv \partial \mathbf{F} / \partial \mathbf{U}$ is a Jacobian matrix. From the theory of partial differential equations, we know that the solution is constant along integral curves in the space of independent variables (t, x) , known as characteristics (see, for example, Riley et al. (2006)). Information (i.e. initial data) propagates along characteristics, and the characteristics separate regions within which the solution is constant.

Physically the characteristics correspond to various waves separating constant states. Information about the waves can be determined by calculating the left- and right-eigenvectors of \mathbf{A} (denoted $\mathbf{L}^{(k)}$ and $\mathbf{R}^{(k)}$, respectively), and the associated eigenvalues $\lambda^{(k)}$. In particular the speed at which information propagates along characteristics is given by the *eigenvalues* of \mathbf{A} , while the *eigenvectors* of \mathbf{A} indicate how the solution changes across characteristics.³

We illustrate the solution (in MHD) in the space of independent variables (t, x) in Figure 2.3, which shows the seven wave families, separating eight constant states (known as a Riemann fan). Proceeding from the outside-in, the outermost wave family consists of fast magnetosonic waves, followed by Alfvén waves, and finally slow magnetosonic waves. The magnetosonic waves are manifested as two types of non-linear wave: shocks or rarefactions. The middle wave is always a linear wave known as a contact wave (or entropy wave), across which only the density is discontinuous.

The two regions on each side of any given characteristic are related through a non-linear set of algebraic equations known as jump conditions.

³Strictly speaking the eigenvalues of \mathbf{A} are equivalent to the wave speeds only for a *linear* system of equations (Toro, 2019). This is no longer true for a non-linear system of equations such as the HD or MHD equations, although we omit the distinction here for brevity.




Figure 2.3: Solution of the Riemann problem in MHD at the $x_{i+\frac{1}{2}}$ interface. The solution consists of seven waves, represented by the colored straight-lines (characteristics), separating eight constant states. The left-hand-most and right-hand-most states are given by the initial data \mathbf{U}_L and \mathbf{U}_R , respectively.

These relate jumps in the fluxes and in the solution across the discontinuity, and the speed of the discontinuity. In particular the jump condition across the k th discontinuity is given by

$$\lambda^{(k)}(\mathbf{U}_{k+1} - \mathbf{U}_k) = \mathbf{F}(\mathbf{U}_{k+1}) - \mathbf{F}(\mathbf{U}_k) \quad (2.44)$$

where in MHD $k \in \{1, 7\}$ and $\mathbf{U}_1 \equiv \mathbf{U}_L$ and $\mathbf{U}_8 \equiv \mathbf{U}_R$.

An *exact* Riemann solver would involve solving the set of jump conditions given by Equation 2.44 at each cell interface, which in turn determines the solution, and therefore the flux, at that interface. This is generally prohibitively computationally expensive. Therefore various techniques are used to approximate the total flux $\mathbf{F}_{i+\frac{1}{2}}$ through the interface. Below we discuss four approximate Riemann solvers that we use in our simulations: the first three (Lax-Friedrichs-Rusanov, HLL, and HLLD) approximate the flux by averaging over the Riemann fan, while the fourth (Roe) approximates the flux by linearizing the jump conditions about some average state.

Lax-Friedrichs-Rusanov solver

The simplest (and most diffusive) approximate Riemann solver estimates the flux by taking the arithmetic mean of the right- and left-hand fluxes and states. No jump conditions are used, and the solver makes no use of the intermediate state(s) except the maximum wave speed λ_{\max} . This is known as the *Lax-Friedrichs-Rusanov* (LFR) solver, and the flux is given by

$$\mathbf{F}_{i+\frac{1}{2}}^{\text{LFR}} = \frac{1}{2} (\mathbf{F}(\mathbf{U}_R) + \mathbf{F}(\mathbf{U}_L)) - \frac{1}{2} |\lambda_{\max}| (\mathbf{U}_R - \mathbf{U}_L) \quad (2.45)$$

See, for example, Tóth & Odstrčil (1996).

HLL solver

An improvement over the Lax-Friedrich-Rusanov solver is given by the *Harten-Lax-van Leer* (HLL) solver, which is an approximate Riemann solver in which the middle (contact) discontinuity is omitted (Harten et al., 1983; Toro, 2013). Thus the solution in hydrodynamics is approximated by a left-propagating wave with speed λ_L and a right-propagating wave with speed λ_R , separating three regions: the known left- and right-hand states \mathbf{U}_L and \mathbf{U}_R , respectively, and an intermediate region with unknown state \mathbf{U}^* . The flux is given by

$$\mathbf{F}_{i+\frac{1}{2}}^{\text{HLL}} = \begin{cases} \mathbf{F}_L & (\lambda_L > 0) \\ \mathbf{F}^* & (\lambda_L \leq 0 \leq \lambda_R) \\ \mathbf{F}_R & (\lambda_R < 0) \end{cases} \quad (2.46)$$

where $\mathbf{F}_L \equiv \mathbf{F}(\mathbf{U}_L)$ and $\mathbf{F}_R \equiv \mathbf{F}(\mathbf{U}_R)$. The intermediate state flux \mathbf{F}^* is obtained by using the jump conditions and is given by

$$\mathbf{F}^* = \frac{\lambda_R \mathbf{F}_L - \lambda_L \mathbf{F}_R + \lambda_R \lambda_L (\mathbf{U}_R - \mathbf{U}_L)}{\lambda_R - \lambda_L}. \quad (2.47)$$

Note that various estimates are possible for the left- and right-wavespeeds λ_L and λ_R , respectively. A commonly used estimate is that given byinfeldt (1988), in which case the solver is referred to as the HLLE solver. Here the speeds of the left and right non-linear waves are each estimated using

density-weighted averages of the left and right advection and sound speeds (the reader may wish to refer to Equations 10.39-10.41 of Toro (2013) for the explicit expressions).

The key simplifying feature of the HLL solver is also its major drawback, i.e. the omission of the middle wave. Unsurprisingly this results in poor numerical resolution of both contact waves (and therefore of jumps in density and entropy) as well as shear waves (associated with jumps in the tangential components of the velocity).

HLLC solver

The HLLC solver is a refinement of the HLL solver which restores the middle (contact) discontinuity (the ‘C’ stands for contact) (Toro et al., 1994; Toro, 2019). In hydrodynamics the Riemann fan is now approximated by three waves separating four states. The flux is given by

$$\mathbf{F}_{i+\frac{1}{2}}^{\text{HLLC}} = \begin{cases} \mathbf{F}_L & (\lambda_L \geq 0) \\ \mathbf{F}_L^* = \mathbf{F}_L + \lambda_L(\mathbf{U}_{*L} - \mathbf{U}_L) & (\lambda_L \leq 0 \leq \lambda_*) \\ \mathbf{F}_R^* = \mathbf{F}_R + \lambda_R(\mathbf{U}_{*R} - \mathbf{U}_R) & (\lambda_* \leq 0 \leq \lambda_R) \\ \mathbf{F}_R & (\lambda_R \leq 0) \end{cases} \quad (2.48)$$

where the left- and right-intermediate state solutions \mathbf{U}_{*L} and \mathbf{U}_{*R} are given by Equation 10.33 in Toro (2013). A generalization of the HLLC solver to MHD is known as the HLLD solver. Here the solution consists of five waves separating six constant states. Further details may be found in Miyoshi & Kusano (2005) and Li (2005).

Roe solver

Finally we discuss the Riemann solver of Roe (Roe (1981); Toro (2013); see Cargo & Gallice (1997) for the MHD implementation). This is different in flavor to the previous HLL-family of Riemann solvers, which involved solving an exact Riemann problem (given by Equations 2.42-2.43) approximately (the approximation being introduced by averaging over intermediate states).

The Roe solver, on the other hand, solves an *approximate* Riemann problem *exactly*. Explicitly, Roe's method solves the Riemann problem given by

$$\frac{\partial \mathbf{U}}{\partial t} + \tilde{\mathbf{A}} \cdot \frac{\partial \mathbf{U}}{\partial x} = 0, \quad (2.49)$$

$$\mathbf{U} = \begin{cases} \mathbf{U}_L, & (x < x_{i+\frac{1}{2}}), \\ \mathbf{U}_R, & (x \geq x_{i+\frac{1}{2}}), \end{cases} \quad (2.50)$$

where the Jacobian matrix $\mathbf{A}(\mathbf{U})$, whose elements are functions of the fluid variables, has been replaced with a *constant* Jacobian matrix $\tilde{\mathbf{A}} = \tilde{\mathbf{A}}(\mathbf{U}_L, \mathbf{U}_R)$ known as the Roe matrix. Various methods exist for constructing the Roe matrix subject to the constraints that it satisfy hyperbolicity (has real eigenvalues), self-consistency [i.e. that $\tilde{\mathbf{A}}(\mathbf{U}, \mathbf{U}) = \mathbf{A}(\mathbf{U})$], and the jump conditions (Roe & Pike, 1984; Toro, 2013).

Ultimately the approximation of replacing the Jacobian with a matrix containing constant elements reduces the problem to one of solving a *linear* system of conservation laws, for which it can be shown that the flux through the interface at $x_{i+\frac{1}{2}}$ is given by

$$\mathbf{F}_{i+\frac{1}{2}}^{\text{Roe}} = \frac{1}{2} (\mathbf{F}_{i+1} + \mathbf{F}_i) - \frac{1}{2} \sum_k |\tilde{\lambda}^{(k)}| \tilde{\mathbf{L}}^{(k)} \cdot (\mathbf{U}_{i+1} - \mathbf{U}_i) \tilde{\mathbf{R}}^{(k)}, \quad (2.51)$$

where $\tilde{\lambda}^{(k)}$ is the eigenvalue of $\tilde{\mathbf{A}}$ associated with the k th discontinuity in the Riemann fan, and $\tilde{\mathbf{L}}^{(k)}$ and $\tilde{\mathbf{R}}^{(k)}$ the corresponding left- and right-eigenvectors, respectively (see, for example, Toro (2013)).

2.3.6 Time-stepping algorithms

The final stage of the reconstruct-solve-average strategy involves updating the solution from that at the n th time-level to the $(n+1)$ th time-level. Generalizing Equation 2.36 to include source terms (see Equation 2.27) and working in one spatial dimension for simplicity we have

$$\langle \mathbf{U} \rangle_i^{n+1} = \langle \mathbf{U} \rangle_i^n + \Delta t \mathbf{H}_i^n, \quad (2.52)$$

where

$$\mathbf{H}_i^n \equiv -\frac{1}{\Delta x} \left(\tilde{\mathbf{F}}_{i+\frac{1}{2}}^{n+\frac{1}{2}} - \tilde{\mathbf{F}}_{i-\frac{1}{2}}^{n+\frac{1}{2}} \right) + \langle \mathbf{S} \rangle_i^n. \quad (2.53)$$

There are various time-stepping algorithms that are available in the codes we use, but the scheme we employ in most of our simulations is the Runge-Kutta algorithm (see, for example, Press et al. (1988)). This introduces an intermediate step to the update process (known a predictor step) given by

$$\langle \mathbf{U} \rangle_i^* = \langle \mathbf{U} \rangle_i^n + \Delta t \mathbf{H}_i^n. \quad (2.54)$$

The full update is then obtained by taking the arithmetic mean of the solution at the n th time-level $\langle \mathbf{U} \rangle_i^n$ and the solution at the intermediate time-level $\langle \mathbf{U} \rangle_i^*$, resulting in

$$\langle \mathbf{U} \rangle_i^{n+1} = \frac{1}{2} (\langle \mathbf{U} \rangle_i^n + \langle \mathbf{U} \rangle_i^* + \Delta t \mathbf{H}_i^*), \quad (2.55)$$

where $\mathbf{H}_i^* \equiv \mathbf{H}(\mathbf{U}_i^*)$ is evaluated using the solution at the intermediate time-step. The global error after N steps is second-order accurate, thus the time-stepping algorithm presented in Equations 2.54 and 2.55 is known as the 2nd-order Runge-Kutta algorithm. We also employ a third-order accurate Runge-Kutta algorithm, which involves an extra predictor step, but we will not discuss it here for brevity (see, for example, Mignone et al. (2007)).

In order to determine the time-step Δt itself (and ensure numerical stability) we require that the fastest wave cannot propagate over a distance greater than the width of a cell between the n th and $(n+1)$ th time-level (known as the Courant-Friedrichs-Lewy (CFL) condition; Courant et al. (1928)). This is formally given by

$$\Delta t = C_a \min \left(\frac{\Delta x}{|\lambda_{\max}|} \right), \quad (2.56)$$

where the minimum is taken over all active cells in the computational domain, λ_{\max} is the maximum signal (or wave) speed in the domain, and C_a is the Courant number, which (by rearranging Equation 2.56) can be interpreted as the ratio of the fastest wave speed in the computational domain to the

grid speed $\Delta x/\Delta t$. We take to be $C_a = 0.3$ in all our simulations. Note that if explicit diffusion coefficients or cooling prescriptions are used these will introduce their own time-step constraints.

2.4 Codes

Most of our simulations are run with the conservative, finite-volume code PLUTO (Mignone et al., 2007). Unless stated otherwise, for our hydrodynamic simulations we employ the Roe Riemann-solver (see Section 2.3.5), 3rd order in space WENO interpolation (Section 2.3.4), and the 3rd order in time Runge-Kutta algorithm. Other configurations are explored in Section 3.3.3. For our MHD simulations we use the HLLD Riemann solver (see Section 2.3.5), 2nd order in space linear interpolation, and the 2nd order in time Runge-Kutta algorithm (Section 2.3.6). Note that due to the code’s conservative form kinetic energy is not lost to the grid but converted directly into thermal energy. Ghost zones are used to implement the boundary conditions. All of our simulations are parallelized using the message passing interface (MPI) library, thus significantly reducing computational time.

To allow for longer time-steps, we employ the FARGO (fast advection in rotating gaseous objects) scheme in all our simulations (Mignone et al., 2012). Thin accretion disks are supersonic in the azimuthal direction. This supersonic motion in turn significantly reduces the time-step when the orbital speed is greater than any other wave speeds in the domain, because the flow is required to satisfy the CFL condition for numerical stability (cf. Equation 2.56). In order to work around this restriction, the FARGO scheme decomposes the total velocity into an azimuthally averaged background term (corresponding to the supersonic azimuthal flow) and a residual term (corresponding to fluctuations). Only the residual term is required to satisfy the CFL condition, leading to longer time-steps. The FARGO scheme works provided there is a supersonic dominant background orbital motion that is aligned with the grid (as is the case in the shearing box approximation). In addition, the scheme conserves total angular momentum and energy to

machine precision, and is known to be less dissipative than the standard numerical implementation (Mignone et al., 2012).

When explicit viscosity ν , thermal diffusivity χ , or resistivity η are included, we further reduce the computational time by using the Super-Time-Stepping (STS) scheme (Alexiades et al., 1996). When this scheme is used, parabolic (i.e. diffusive) terms are solved separately from hyperbolic (i.e. advective) terms using operator splitting. The solution vector for the parabolic terms is solved over a ‘super time-step’ ΔT , which in turn consists of N unequal sub-steps Δt . The CFL condition is enforced only at the end of this super time-step rather than after each sub-step. The super time-step itself depends on the number of sub-steps, the diffusion coefficient, the grid-size, and a free parameter ν (not to be confused with the kinematic viscosity) that depends on the minimum and maximum wavespeeds and which is adjusted to allow Δt to be as large as possible while guaranteeing stability (see Equation 2.10 of (Alexiades et al., 1996)). In the limit that the free parameter ν tends to zero, a single super-time-step consisting of N sub-steps advances the solution in time by an increment that is N times longer than that achieved in the same number of steps using an explicit scheme. The greatest speed-up is obtained in problems where the diffusion or resolution are sufficiently high that the parabolic time-step falls below the advection time-step (Mignone et al., 2007). Another example in which speed-up can be obtained using STS is in problems involving anisotropic thermal conductivity (Vaidya et al., 2017).

We use the built-in shearing box module in PLUTO (Mignone et al., 2012). As explained in Section 2.3.2, rather than solving Equations 2.1-2.4 (primitive form), PLUTO solves the governing equations in *conservative form*, evolving the total energy equation rather than the thermal energy equation, where the total energy density of the fluid is given by $E = (1/2)\rho u^2 + (1/2\mu_0)B^2 + \rho e$ (kinetic + magnetic + internal). In PLUTO the cooling term Λ is not implemented directly in the total energy equation. Instead, it is included on the right-hand-side of the thermal energy equation, which is then integrated (in time) analytically.

Finally, we have verified the results of our fiducial hydrodynamic simu-

lation (presented in Section 3.3.2) with the conservative finite-volume code ATHENA (Stone et al., 2008).

2.5 Diagnostics

Averaged quantities

We follow the time-evolution of various volume-averaged quantities (e.g. kinetic energy density, thermal energy density, magnetic energy density, Reynolds and Maxwell stress, α , mass density, thermal pressure, and temperature). For a quantity X the volume-average of that quantity is denoted $\langle X \rangle$ and is defined as

$$\langle X \rangle(t) \equiv \frac{1}{V} \int_V X(x, y, z, t) dV \quad (2.57)$$

where V is the volume of the box.

We are also interested in averaging certain quantities (e.g. the Reynolds stress) over time. The temporal average of a quantity X is denoted $\langle X \rangle_t$ and is defined as

$$\langle X \rangle_t(x, y, z) \equiv \frac{1}{\Delta t} \int_{t_i}^{t_f} X(x, y, z, t) dt, \quad (2.58)$$

where we integrate from some initial time t_i to some final time t_f and $\Delta t \equiv t_f - t_i$.

If we are interested only in the vertical structure of a quantity X then we average over the x - and y -directions, only. The horizontal average of that quantity is denoted $\langle X \rangle_{xy}$ and is defined as

$$\langle X \rangle_{xy}(z, t) \equiv \frac{1}{A} \int_A X(x, y, z, t) dA, \quad (2.59)$$

where A is the horizontal area of the box. Horizontal averages over different coordinate directions (e.g. over the y - and z -directions) are defined in a similar manner.

We track the vertical profiles of horizontally averaged pressure, density, temperature, Reynolds stress, magnetic stress, and plasma beta. Of special

interest is the buoyancy frequency, which is calculated from the pressure and density data by finite differencing the formula

$$\left\langle \frac{N_B^2}{\Omega^2} \right\rangle_{xy} = \left\langle z \left[\frac{1}{\gamma} \frac{d \ln P}{dz} - \frac{d \ln \rho}{dz} \right] \right\rangle_{xy}. \quad (2.60)$$

Note we calculate horizontal (and time) averages after calculating N_B^2 . To check the effects of the order of operations (i.e. averaging N_B^2 compared to averaging P and ρ and then calculating N_B^2) we have also calculated the buoyancy frequency by finite differencing the formula

$$\left\langle \frac{N_B^2}{\Omega^2} \right\rangle_{xy} = z \left[\frac{1}{\gamma} \frac{d \ln \langle P \rangle_{xy}}{dz} - \frac{d \ln \langle \rho \rangle_{xy}}{dz} \right]. \quad (2.61)$$

We measured the percentage difference between Equations 2.60 and 2.61 in our fiducial stratified MHD simulation (see Section 4.2) and found it to be less than 1% within $z = \pm 2H_0$.

We also calculate the (horizontally averaged) vertical profiles of mass and heat flux. We define the mass flux as

$$F_{\text{mass}} = \langle \rho u_z \rangle_{xy}, \quad (2.62)$$

and the heat flux as

$$F_{\text{heat}} = \langle \rho u_z T \rangle_{xy}. \quad (2.63)$$

Reynolds and magnetic stresses, and alpha

In accretion discs, the radial transport of angular momentum is related to the xy -component of the total stress

$$\Pi_{xy} \equiv R_{xy} + M_{xy}. \quad (2.64)$$

$\Pi_{xy} > 0$ corresponds to positive (i.e. radially outward) angular momentum transport, while $\Pi_{xy} < 0$ corresponds to negative (i.e. radially inward) angular momentum transport. In Equation 2.64, R_{xy} is the Reynolds stress, defined as

$$R_{xy} \equiv \rho u_x \delta u_y \quad (2.65)$$

where $\delta u_y \equiv u_y + q\Omega x$ is the fluctuating part of the y -component of the total velocity u_y . The magnetic stress M_{xy} is defined as

$$M_{xy} \equiv -B_x B_y, \quad (2.66)$$

where B_x and B_y are the magnetic field components in the x - and y -directions, respectively. Note that here (and for the remainder of this thesis) we have chosen our units such that the constant μ_0 disappears.

Finally, the total stress is related to the classic dimensionless parameter α by

$$\alpha \equiv \frac{\langle \Pi_{xy} \rangle}{q \langle P \rangle}, \quad (2.67)$$

where $q \equiv -d \ln \Omega / d \ln r|_{r=r_0}$ is the dimensionless shear parameter and $q = 3/2$ for a Keplerian disk. It is customary for many authors to drop the factor q in Equation 2.67. To compare our results more easily with the literature we too drop the factor of q in our measurements of α in our MHD simulations. In our hydrodynamic simulations, however, the factor of q has been retained (this is because we retained this factor when we published our hydrodynamic results, see Held & Latter (2018)).

Energy densities

The kinetic energy density of a fluid element is defined as

$$E_{\text{kin}} \equiv \frac{1}{2} \rho u^2, \quad (2.68)$$

where u is the magnitude of the *total* velocity of a fluid element. Often we will plot the vertical kinetic energy density $E_{\text{kin},z} = \frac{1}{2} \rho u_z^2$, in which case u in the above equation is replaced by u_z .

The magnetic energy density (equivalently the magnetic *pressure*) of a fluid element is defined as

$$E_{\text{mag}} \equiv \frac{1}{2} B^2, \quad (2.69)$$

where B is the magnitude of the *total* magnetic field of a fluid element.

The thermal energy density of a fluid element is defined as

$$E_{\text{th}} \equiv \rho e = \frac{P}{\gamma - 1}. \quad (2.70)$$

Finally, the total energy density is given by

$$E_{\text{total}} = \frac{1}{2}\rho u^2 + \frac{1}{2}B^2 + \rho\Phi + \rho e, \quad (2.71)$$

where the terms on the right-hand side correspond to the kinetic, magnetic, gravitational potential, and thermal energy densities, respectively. Here $\Phi = \frac{1}{2}\Omega_0^2 z^2 - \frac{3}{2}\Omega_0^2 x^2$ is the effective gravitational potential in the shearing box approximation for a Keplerian disk.

Because we employ open boundaries in the vertical direction, energy is lost through advection of fluid across the vertical boundaries. It is important to determine the extent to which this energy loss (which we refer to as *box-cooling*) influences the total energy budget. The integrated flux of total energy across the vertical boundaries is given by

$$\mathcal{F}_z \equiv \frac{1}{V} \left[\iint (u_z E_{\text{tot}} + u_z P_t - B_z(\mathbf{u} \cdot \mathbf{B})) dx dy \right]_{z=-L_z/2}^{z=+L_z/2}, \quad (2.72)$$

where the $P_t = P + B^2/2$, and the $u_z P$ term corresponds to the contribution due to the ram pressure.

The integrated flux of total energy across the vertical boundaries (normalized by the volume-averaged thermal energy density) can be used to estimate the rate at which energy is lost through the vertical boundaries, known as the wind cooling timescale τ_w (see Riols & Latter (2018)), and this is defined by

$$\frac{1}{\tau_w(t)} \equiv \frac{\mathcal{F}_z}{\langle E_{\text{th}} \rangle}. \quad (2.73)$$

2D Power Spectra

In order to distinguish convection from the MRI in a quantitative manner we use the 2D power spectrum of the specific vertical kinetic energy. To calculate this, we first extract the (y -averaged) vertical component of the velocity $w \equiv \langle u_z \rangle_y$ from $z = 0$ and $z = +H_0$ (as most of the activity in the disk lies within a scale height of the mid-plane). Because we use open

boundary conditions (in z) and shear periodic boundary conditions (in x) this data is *not* periodic in either direction. To make it periodic we reflect the extracted data in x and z to create a doubly-periodic array spanning $z = -H_0$ to $z = +H_0$ and $x = -4H_0$ to $x = +4H_0$. The kl th component of the 2D discrete Fourier transform of this data is then defined as

$$\widehat{w}_{kl} = \sum_{n=0}^{\tilde{N}_x-1} \sum_{m=0}^{\tilde{N}_z-1} w_{mn} \exp \left\{ -2\pi i \left(\frac{mk}{\tilde{N}_x} + \frac{nl}{\tilde{N}_z} \right) \right\}, \quad (2.74)$$

where $k = 0, \dots, \tilde{N}_x - 1$, $l = 0, \dots, \tilde{N}_z - 1$, and w_{mn} denotes the m th component of the 2D array $w \equiv \langle u_z \rangle_y$. Note that \tilde{N}_x (and \tilde{N}_z) are the number of radial (vertical) cells in the extended periodic partition of the disk; in general they are *not* equal to the total number of radial (vertical) cells in the domain N_x (and N_z).

The power in the specific vertical kinetic energy can then be obtained from

$$\hat{E}_{\text{kin},z} = \frac{1}{2} |\widehat{w}|^2, \quad (2.75)$$

where $|\widehat{w}|^2 \equiv \widehat{w}\widehat{w}^*$. Finally we plot Equation 2.75 in the (k_x, k_z) -plane to obtain the 2D power spectrum of the specific vertical kinetic energy, where k_x and k_z are the radial and vertical wavenumber, respectively.

Diagnostics for detecting convection

In general it is difficult to detect convection against the backdrop of MRI turbulence. To aid us in determining whether convection is present in our simulations or not, we employ various diagnostics. Our first diagnostic is the sign of the buoyancy frequency N_B^2 (Equation 2.60), though as discussed in Section 4.1.2 this is a necessary but not a sufficient criterion in turbulent flows. Thus where $N_B^2 > 0$ we can definitely rule out the presence of convection, but where $N_B^2 < 0$ further diagnostics are needed to determine whether convection is present or not.

As a second diagnostic, we employ visual inspection of the flow field in the xz -plane (i.e. that the vertical velocity u_z exhibit hot updrafts and cool downdrafts). To try to quantify any vertical ‘structure’ that we pick up

visually, we use a third diagnostic, namely 2D power spectra (cf. Section 2.5). From linear theory, convective cells are arranged such as to minimize radial fluid motion (i.e. $k_x/k_z \gg 1$), whereas the MRI tends to *maximize* radial fluid motions (i.e. $k_x/k_z \ll 1$). Thus each instability conveniently inhabits different regions of wavenumber space.

Finally we also use the horizontal- and time-averaged values of the vertical heat and mass flux. In particular, we have found in our hydrodynamic simulations that convection exhibits not just a heat flux *away* from the mid-plane (as does, to a lesser extent, the MRI) but - surprisingly - a small mass flux *towards* the mid-plane (see Figure 3.12 of Chapter 3). Thus convection is rearranging the vertical disk structure (by changing the temperature and density profiles) so that the disk is closer to a state of marginal stability. A caveat, however, is that the inward mass flux appears to decrease over very long timescales, and is thus probably a slow transient effect. Nevertheless over the timescales of interest (c. 100 orbits) we have found it to be a reliable diagnostic for convection.⁴

⁴The unforced hydro simulations in question did not include a mass source term, thus we are confident that the inward mass flux observed when convection is present is not an artifact of our mass source term.

Chapter 3

Hydrodynamic convection in disks

Our aim in this chapter is to explore the competing results and claims concerning hydrodynamic convection in accretion discs, in particular to determine the sign and magnitude of angular momentum transport. We also isolate and characterize other generic features of convection that should be shared by multiple disk classes (dwarf novae, protoplanetary disks, etc) and by different driving mechanisms (MRI turbulence, spiral shock heating, etc). We do so both analytically and through numerical simulations, working in the fully compressible, vertically stratified shearing box approximation. We restrict ourselves to an idealized hydrodynamic set-up, omitting magnetic fields and complicated opacity transitions, and in so doing ensure our results are as general as possible.

The structure of this chapter is as follows: first, in Section 3.1 we provide a brief overview of the numerical parameters and set-up. In Section 3.2 we investigate the linear behavior of the convective instability, employing both analytical WKBJ and semi-analytical spectral methods to calculate the growth rates and the eigenfunctions. In Section 3.3 we explore the non-linear regime through *unforced* simulations, in which the convection is not sustained and is permitted to decay after non-linear saturation. Because this unforced convection is a transient phenomenon which might depend on the initial con-

ditions, in Section 3.4 we explore the non-linear regime through simulations of *forced* convection, in which the internal energy is relaxed to its initial, convectively unstable, state to mimic, possibly, the action of background MRI turbulent dissipation and optically thick radiative cooling. Finally, Section 3.5 summarizes and discusses the results.

3.1 Numerical set-up

3.1.1 Initial conditions and units

Time units are selected so that $\Omega_0 = 1$. From now the subscript on the angular frequency is dropped if it appears. The length unit is chosen so that the initial mid-plane isothermal sound speed $c_{s0} = 1$, which in turn defines a constant reference scale height $H \equiv c_{s0}/\Omega = 1$. Note, however, that the sound speed is generally a function of both space and time.

We use two different convectively unstable profiles to initialize our simulations. Nearly all employ an equilibrium exhibiting a Gaussian temperature profile

$$T = T_0 \exp \left[-\frac{z^2}{\beta H^2} \right], \quad (3.1)$$

where T_0 is mid-plane temperature and β is a dimensionless tuning parameter. See Figure 3.2 for all associated thermal profiles. For comparison we also use the profile introduced by SB96 in which the temperature follows a power law

$$T = T_0 - Az^p, \quad (3.2)$$

where A and p are parameters, with $p = 3/2$ usually (see Figure 3.1). Further details of both profiles are given later in this section. The equilibria are convectively unstable within a confined region (of size L_c) about the disk mid-plane, and convectively stable outside this region. We have checked that these vertical profiles satisfy vertical hydrostatic equilibrium in our numerical set-up by running a series of simulations without perturbations; these show that after 55 orbits the initial profiles are unchanged with velocity fluctuation

amplitudes typically less than $10^{-3} c_{s0}$ at the vertical boundaries and less than $10^{-7} c_{s0}$ at the mid-plane.

The background velocity is given by $\mathbf{u} = -(3/2)\Omega_0 x \mathbf{e}_y$. At initialization we usually perturb all the velocity components with random noise exhibiting a flat power spectrum. The perturbations $\delta\mathbf{u}$ have maximum relative amplitude of about $5 \times 10^{-5} c_{s0}$ and can be either positive or negative. In order to investigate specifically the nature of linear axisymmetric convective modes, we initialize several PLUTO simulations with linear axisymmetric modes calculated semi-analytically rather than with random noise (see Section 3.2.5).

Convectively unstable vertical disk profiles

We describe the two convectively unstable vertical profiles that are used to initialize our simulations. The reader should note that these profiles may or may not correspond to those in *real* astrophysical disks, which will be determined by several sources of heating and cooling, none considered in this chapter. Here we simply present convectively unstable disk profiles that satisfy hydrostatic equilibrium, are convectively unstable, and can conveniently initialize simulations.

Stone and Balbus (1996) profile The SB96 profile employs a power law profile in temperature, see Eq. 3.2. For $p = 3/2$, the density can be obtained analytically

$$\rho = \rho_0 (1 - s^3)^{-(1+g/3)} (1 - s)^g \exp \left\{ 2g \left[s - \frac{1}{\sqrt{3}} \tan^{-1} \left(\frac{\sqrt{3}s}{s+2} \right) \right] \right\}$$

where ρ_0 is the mid-plane density, $s = (z/H)^{1/2} f^{1/3}$, and $g = f^{-4/3}$, in which $f = H^{3/2} A / T_0$, and $H = c_s / \Omega$ is the mid-plane scale height. The pressure is obtained from the ideal gas equation of state. Note that for $|z| > (T_0/A)^{1/p}$ we must have vacuum, which ties the maximum numerical domain to the ratio T_0/A .




Figure 3.1: Vertical disk structure of Stone and Balbus (1996) in a disk of height $2H$ about the mid-plane. *Left*: vertical profiles for density ρ , pressure P , and temperature T . *Right*: vertical profile of the buoyancy frequency N_B^2/Ω^2 . For clarity, only half of the vertical domain is shown. The profile parameters are $\{T_0 = 1.0, \rho_0 = 1.0, g = 5.0\}$ and the adiabatic index is $\gamma = 5/3$.

The buoyancy frequency for $p = 3/2$ is given by

$$\frac{N_B^2}{\Omega^2} = \frac{1}{1 - \left(\frac{z}{H}\right)^{3/2} f} \left[\left(1 - \frac{1}{\gamma}\right) \left(\frac{z}{H}\right)^2 - \frac{3}{2} \left(\frac{z}{H}\right)^{3/2} f \right] \quad (3.3)$$

which corresponds to a profile that is negative (convectively unstable) within some region $|L_c| < 0$ about the mid-plane and positive (convectively unstable) outside of this region. The width of the convectively unstable region is given by

$$|L_c| = \left[\frac{3}{2} \left(1 - \frac{1}{\gamma}\right)^{-1} H^2 \frac{A}{T_0} \right]^2. \quad (3.4)$$

Although convenient within a limited choice of parameters, the SB96 profile suffers from the drawback that the width of the convectively unstable region is sensitive to the size of the box through the ratio T_0/A . Increasing the vertical box size necessarily decreases the size of the convectively unstable region.




Figure 3.2: Vertical disk structure for Gaussian temperature profile in a disk of height $2H$ about the mid-plane. *Left*: vertical profiles for density ρ , pressure P , and temperature T . *Right*: vertical profile of the buoyancy frequency N_B^2 / Ω^2 . For clarity, only half of the vertical domain is shown. The profile parameters are $\{T_0 = 1.0, \rho_0 = 1.0, \beta = 3.0\}$ and the adiabatic index is $\gamma = 5/3$.

Gaussian temperature profile The drawbacks of the SB96 profile motivated us to search for a more convenient unstable profile, one that would leave the size and depth of the convectively unstable region independent of the vertical extent of the box. Setting the temperature to a Gaussian, cf. Eq. 3.1, provided such a profile.

The associated density is

$$\rho = \rho_0 e^{z^2/\beta H^2} \exp\left(-\frac{\beta H^2 \Omega^2 \mu}{2T_0 \mathcal{R}} e^{z^2/\beta H^2}\right), \quad (3.5)$$

where ρ_0 is mid-plane density and H the mid-plane scale height. Note that the factor $-\frac{\beta H^2 \Omega^2 \mu}{2T_0 \mathcal{R}}$ is dimensionless.¹ Pressure is obtained from the ideal gas equation of state, as above.

¹In this subsection, for completeness, the equations are dimensional. In our codes, however we work with dimensionless variables, in which the scales are chosen such that the factor \mathcal{R}/μ disappears. In addition the (dimensionless) scale height and angular frequency are set to one such that the factor $-\frac{\beta H^2 \Omega^2 \mu}{2T_0 \mathcal{R}}$ reduces to $-\beta/2T_0$ where T_0 is now the dimensionless mid-plane temperature.




Figure 3.3: Horizontally y -averaged vertical profiles of temperature (green), density (red) and pressure (blue) in a PLUTO simulation initialized without perturbations to test hydrostatic equilibrium. The solid curves are the profiles at initialization; the dots are data from the simulation after 55 orbital time-scales.

The buoyancy frequency is given by

$$\frac{N_B^2}{\Omega^2} = \frac{2}{\beta H^2} z^2 \left[\left(1 - \frac{1}{\gamma} \right) \frac{\beta H^2 \Omega^2 \mu}{2T_0 \mathcal{R}} e^{z^2/\beta H^2} - 1 \right]. \quad (3.6)$$

The boundary of the convectively unstable region is hence given by

$$|L_c| \equiv \sqrt{\beta H^2 \ln \left[\left(1 - \frac{1}{\gamma} \right)^{-1} \frac{2T_0 \mathcal{R}}{\beta \mu H^2 \Omega^2} \right]}. \quad (3.7)$$

Hydrostatic equilibrium tests

Here we briefly present checks that the convectively unstable vertical profile used to initialize our simulations (see Section 3.1.1) as well as the vertical boundary conditions both satisfy hydrostatic equilibrium in the vertical direction.

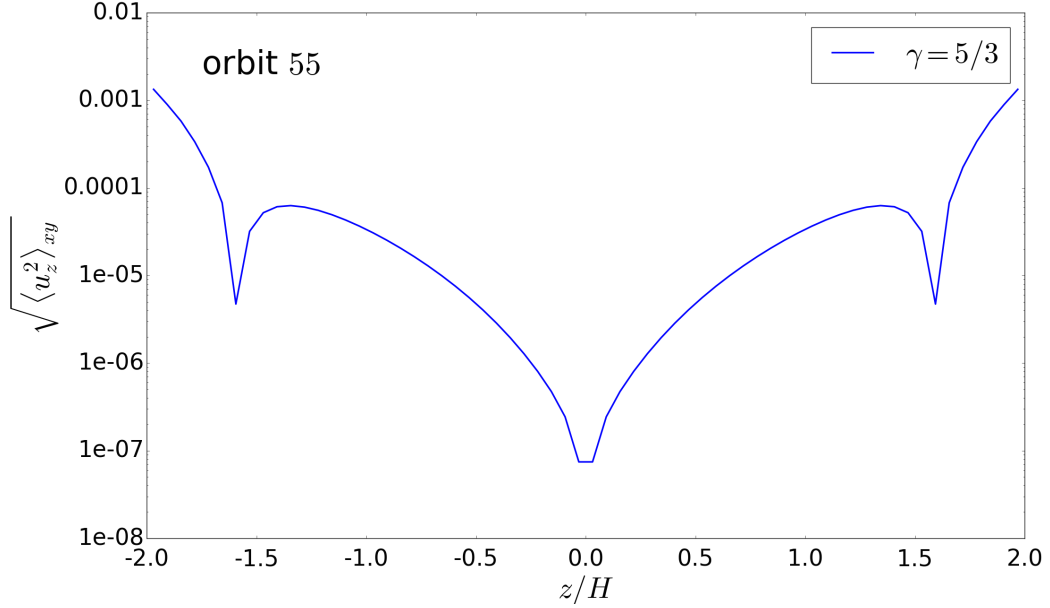


Figure 3.4: Vertical profile of root-mean-square of horizontally-averaged vertical velocity component after 55 orbits in a PLUTO simulation initialized without perturbations.

We initialize a PLUTO simulation at a resolution of $64 \times 64 \times 64$ in a box of size $4H \times 4H \times 4H$ with the vertical density profile given by Equation 3.5. The parameters characterizing our profile are given by $\{T_0 = 1.0, \rho_0 = 1.0$ and $\beta = 3.0\}$. The ratio of specific heats is taken to be $\gamma = 5/3$. This profile is convectively unstable within a region $|z| < L_c \sim 1.2H$ about the mid-plane, but should satisfy hydrostatic equilibrium throughout the entire vertical domain. In the vertical direction we use outflow boundary conditions on the velocity, while the ghost cells are kept in hydrostatic equilibrium.

In Figure 3.3 we plot the horizontally averaged vertical profiles of temperature, density and pressure after 55 orbits. The profiles taken from the simulations data after 55 orbits are almost indistinguishable from those at initialization. In Figure 3.4 we plot the vertical profile of the root-mean-square vertical component of the velocity at orbit 55. The fluctuations are typically less than 10^{-3} even after 55 orbits indicating that there is very little excitation due to the vertical boundaries and that the code is maintaining

hydrostatic equilibrium very well.

3.1.2 Box size and resolution

We measure box size in units of scale height H , defined above. We employ resolutions of 64^3 , 128^3 , 256^3 and 512^3 in boxes of size $4H \times 4H \times 4H$, which correspond to resolutions of 16, 32, 64 and 128 grid cells per scale height H in all directions. However, in simulations in which we force convection we employ a ‘wide-box’ ($6H \times 6H \times 4H$) with a resolution of $256 \times 256 \times 256$ which corresponds to about 43 grid cells per H in the horizontal directions, and 64 grid cells per H in the vertical direction.

3.1.3 Boundary conditions

We use shear-periodic boundary conditions in the x -direction (see Hawley et al. (1995a)) and periodic boundary conditions in the y -direction. In the vertical direction, we keep the ghost zones associated with the thermal variables in isothermal hydrostatic equilibrium (the temperature of the ghost zones being kept equal to the temperature of the vertical boundaries at initialization) in the manner described in Zingale et al. (2002). For the velocity components we use mostly standard *outflow* boundary conditions in the vertical direction, whereby the vertical gradients of all velocity components are zero, and variables in the ghost zones are set equal to those in the active cells bordering the ghost zones. In a handful of our simulations we also use free-slip and periodic boundary conditions to test the robustness of our results.

3.1.4 Important parameters and instability criteria

The onset of thermal convection is controlled by the local *buoyancy frequency*, defined as

$$N_B^2 = z \left[\frac{1}{\gamma} \frac{\partial \ln P}{\partial z} - \frac{\partial \ln \rho}{\partial z} \right] \Omega_0^2. \quad (3.8)$$

But convection is opposed by the effects of viscosity and thermal diffusivity, with the ratio of destabilizing and stabilizing processes quantified by the

Rayleigh number

$$\text{Ra} \equiv \frac{|N_B^2| H^4}{\nu \chi}, \quad (3.9)$$

On the other hand, the ratio of buoyancy to shear is expressed through the *Richardson number*

$$\text{Ri} \equiv \frac{|N_B^2|}{q^2 \Omega_0^2}. \quad (3.10)$$

But for convection perpendicular to the plane of an accretion disk, the buoyancy force is perpendicular to the direction of the background shear. Thus the Richardson number in our set-up is more an expression for the scaled intensity of N_B^2 . Finally, the *Prandtl number* is defined as the ratio of kinematic viscosity to thermal diffusivity

$$\text{Pr} \equiv \frac{\nu}{\chi}. \quad (3.11)$$

When the explicit diffusion coefficients are neglected ($\nu, \chi = 0$), convective instability occurs when the buoyancy frequency is negative (i.e. in regions where $N_B^2 < 0$). When explicit viscosity and thermal diffusivity are included, however, convective instability requires both that $N_B^2 < 0$, *and* that the Rayleigh number exceed some critical value Ra_c . Though in real astrophysical disks the microscopic viscosity is negligible, the inclusion of magnetic fields complicates the stability criterion, as does the presence of pre-existing turbulence (as might be supplied by the MRI) which itself may diffuse momentum and heat. In the latter case it may be possible to define a turbulent ‘Rayleigh number’, which must be sufficiently large so that convection resists the disordered background flow. In any case, the sign of N_B^2 is certainly insufficient to assign convection to MRI-turbulent flows, as is often done in recent work (Hirose et al., 2014; Scepi et al., 2018a).

3.2 Linear theory

In this section we investigate the linear phase of the axisymmetric convective instability both semi-analytically by solving a 1D boundary value / eigenvalue problem using spectral methods, and also analytically using WKBJ methods.

Our aim is to determine the eigenvalues (growth rates) and eigenfunctions (density, velocity, and pressure perturbations) of the axisymmetric convective instability in the shearing box as a function of radial wavenumber k_x . The convectively unstable background vertical profile used in all calculations is shown in Figure 3.2, and the profile is discussed in detail in Section 3.1.1. All calculations were carried out with profile parameters $T_0 = 1.0, \rho_0 = 1.0, \beta = 3.0$ and an adiabatic index of $\gamma = 5/3$. For our chosen background equilibrium, this corresponds to a Richardson number of $\text{Ri} \sim 0.05$.

We proceed as follows: first, we produce linearized equations for the perturbed variables; second, these are Fourier analyzed so that $X'(x, z, t) = e^{\sigma t} e^{ik_x x} X'(z)$ for each of the perturbed fluid variables ρ', \mathbf{u}' and P' . Here σ can be complex (if σ is real and positive it is referred to as a *growth rate*). The linearized equations are given by

$$\sigma \rho' = -ik_x \rho_0 u'_x - \partial_z(\rho_0 u'_z), \quad (3.12)$$

$$\sigma u'_x = 2\Omega u'_y - \frac{ik_x}{\rho_0} P', \quad (3.13)$$

$$\sigma u'_y = (q - 2)\Omega u'_x, \quad (3.14)$$

$$\sigma u'_z = -\frac{1}{\rho_0} \partial_z P' + \frac{\rho'}{\rho_0^2} \partial_z P_0, \quad (3.15)$$

$$\sigma P' = -u'_z \partial_z P_0 - \gamma P_0 (ik_x u'_x + \partial_z u'_z). \quad (3.16)$$

Background fluid quantities are denoted $X_0 = X_0(z)$. For equilibrium ρ_0 and P_0 see Section 3.1.1.

3.2.1 Spectral methods

We discretize the fluid variables on a vertical grid containing N_z points by expanding each eigenfunction as a linear superposition of Whittaker cardinal (sinc) functions (see below for more detail). These are appropriate as all perturbations should decay to zero far from the convectively unstable region. This choice also saves us from explicitly imposing boundary conditions. The discretized equations are next gathered up into a single algebraic eigenvalue

problem (Boyd, 2001). Finally, we solve this matrix equation numerically using the QR algorithm to obtain the growth rates σ and eigenfunctions for a given radial wavenumber k_x . The methodology is discussed in detail below. The reader interested only in the solutions should skip directly to Section 3.2.3.

Whittaker cardinal basis

Consider the one-dimensional eigenvalue problem

$$\mathcal{L}G(z) = \lambda G(z) \quad (3.17)$$

where \mathcal{L} is a differential operator with corresponding eigenvalue λ and eigenfunction $G(z)$. To solve Equation 3.17 using *spectral methods*, we approximate the eigenfunction $G(z)$ as a truncated series of $N + 1$ terms of the form

$$G(z) \approx G_N(z) = \sum_{j=-N/2}^{N/2} a_j C_j(z). \quad (3.18)$$

Note that the truncated series $G_N(z)$ given by Equation 3.18 has not yet been discretized on a grid, i.e. it is a *continuous* function of the independent variable z . The form of the basis functions $C_j(z)$ depends on the type of problem, specifically on how the eigenfunctions behave at the boundaries. For our problem a set of appropriate basis functions are the so-called *Whittaker cardinal functions* defined as

$$C_j(z) \equiv \text{sinc} \left(\frac{z - jh}{h} \right), \quad (3.19)$$

where h is the grid spacing (i.e. the distance between neighboring interpolation points), and the normalized sinc function is defined as

$$\text{sinc}(z) \equiv \frac{\sin(\pi z)}{\pi z}. \quad (3.20)$$

The motivation for choosing Whittaker cardinal functions as our basis is because these enable spectral methods to be employed in calculations on

a doubly infinite domain, i.e. $z \in [-\infty, \infty]$, provided that the eigenfunctions decay exponentially at infinity. The eigenfunctions in our problem correspond physically to perturbations to the density, pressure and the three velocity components: in other words, they describe the structure of the convective cells ‘sitting’ on top of the background equilibrium. Our domain is, of course, not infinite, but bounded in the vertical direction. Nevertheless, outside the convectively unstable region there is no buoyancy force for the perturbations to feed off. The fluid is convectively stable in these regions so we expect the perturbations to decay rapidly outside the convectively unstable region. Thus as long as we locate the vertical boundaries a sufficient distance from the boundaries of the convectively unstable region, the perturbations should decay to zero before reaching the edge of the vertical domain.² Therefore Whittaker cardinal functions should be suitable as a basis to describe convection vertical to the plane of the disk. The vertical structure and physical interpretation of the eigenfunctions are discussed in greater detail in Section 3.2.4, where we find that all the eigenfunctions do indeed decay to zero near the vertical boundaries, as required.

Note from Equation 3.19 that the Whittaker cardinal functions have the property that $C_j(z_i = ih) = \delta_{ij}$. In other words, if we discretize the independent variable z on a grid such that $z_i \equiv ih$, $i = -N/2, \dots, N/2$ with grid spacing h , the Whittaker cardinal functions will have non-zero values only at the grid points (where they are equal to unity). Evaluating the truncated expansion given by Equation 3.18 we obtain

$$G_N(z_i) = \sum_{j=-N/2}^{N/2} a_j C_j(z_i) = \sum_{j=-N/2}^{N/2} a_j \delta_{ij} = a_i, \quad (3.21)$$

thus we see that the coefficient a_i in the i th term of the truncated expansion of the eigenfunction $G(z)$ is identical to the value of the eigenfunction evaluated at the i th grid point.

²We have confirmed this by running the eigensolver in a box of size $\pm H$ about the mid-plane, so that the vertical boundaries are well within the boundaries of the convectively unstable region. The resulting eigenfunctions have a sawtooth pattern, which is clearly unphysical.

Now consider again Equation 3.17 but evaluated at the point $z = z_i$

$$(\mathcal{L}G)_i = \lambda G_i, \quad (3.22)$$

where $G_i \equiv G(z_i) = a_i$, and

$$(\mathcal{L}G)_i \equiv \mathcal{L}G(z)|_{z=z_i} \approx \sum_{j=-N/2}^{N/2} a_j [\mathcal{L}C_j(z)]_{z=z_i}, \quad (3.23)$$

where the approximation arises from truncation of the series. Thus, defining $\mathcal{L}_{ij} \equiv [\mathcal{L}C_j(z)]_{z=z_i}$, we may write the eigenvalue problem given by Equation 3.17 as

$$\sum_{j=-N/2}^{N/2} \mathcal{L}_{ij} a_j = \lambda a_i, \quad i = -\frac{N}{2}, \dots, \frac{N}{2}. \quad (3.24)$$

Equation 3.24 corresponds to a system of $N + 1$ equations for the coefficients a_i , which, the reader will recall, are equivalent to the values of the eigenfunction $G(z)$ at the grid points z_i (see Equation 3.21). Expressing Equation 3.24 in matrix form we have

$$\underline{\underline{\mathcal{L}}} \mathbf{G}^{(k)} = \lambda^{(k)} \mathbf{G}^{(k)} \quad (3.25)$$

where the superscript k refers to the k th eigenvector (or eigenvalue) of the matrix $\underline{\underline{\mathcal{L}}}$.

We may also write out Equation 3.24 in full matrix form as

$$\begin{pmatrix} [\mathcal{L}C_{-\frac{N}{2}}(z)]_{z=-\frac{N}{2}} & \cdots & [\mathcal{L}C_{-\frac{N}{2}}(z)]_{z=\frac{N}{2}} \\ \vdots & \ddots & \vdots \\ [\mathcal{L}C_{\frac{N}{2}}(z)]_{z=-\frac{N}{2}} & \cdots & [\mathcal{L}C_{\frac{N}{2}}(z)]_{z=\frac{N}{2}} \end{pmatrix} \begin{pmatrix} a_{-\frac{N}{2}}^{(k)} \\ \vdots \\ a_{\frac{N}{2}}^{(k)} \end{pmatrix} = \lambda^{(k)} \begin{pmatrix} a_{-\frac{N}{2}}^{(k)} \\ \vdots \\ a_{\frac{N}{2}}^{(k)} \end{pmatrix}. \quad (3.26)$$

To summarize, we may employ spectral methods with a Whittaker cardinal basis to solve eigenvalue problems involving a linear differential operator. Beginning with a differential equation of the form $\mathcal{L}G(z) = \lambda G(z)$, we first express the eigenfunctions $G(z)$ as a truncated series of $N + 1$ terms, with

Whittaker cardinal functions acting as the basis functions. By discretizing the differential equation $\mathcal{L}G(z) = \lambda G(z)$, we then obtain a system of $N + 1$ equations given in matrix form by $\underline{\underline{\mathcal{L}}}\mathbf{G}^{(k)} = \lambda^{(k)}\mathbf{G}^{(k)}$. Here $\underline{\underline{\mathcal{L}}}$ is an $(N + 1) \times (N + 1)$ matrix whose elements correspond to derivatives of Whittaker cardinal functions evaluated at the $N + 1$ vertical grid points, and $\mathbf{G}^{(k)}$ denotes a vector containing $N + 1$ elements, each of which corresponds to the k th eigenfunction $G^{(k)}$ evaluated at one of the $N + 1$ grid points.

To see how the aforementioned method applies to our specific problem, let us gather up Equations 3.12-3.16 into matrix form to obtain

$$\sigma \begin{pmatrix} \rho' \\ u'_x \\ u'_y \\ u'_z \\ P' \end{pmatrix} = \begin{pmatrix} 0 & -ik_x\rho_0 & 0 & -\partial_z\rho_0 - \rho_0\partial_z & 0 \\ 0 & 0 & 2\Omega & 0 & -\frac{ik_x}{\rho_0} \\ 0 & -\Omega/2 & 0 & 0 & 0 \\ \frac{1}{\rho_0^2}(\partial_z P_0) & 0 & 0 & 0 & -\frac{1}{\rho_0}\partial_z \\ 0 & -\gamma P_0 ik_x & 0 & -\partial_z P_0 - \gamma P_0 \partial_z & 0 \end{pmatrix} \begin{pmatrix} \rho' \\ u'_x \\ u'_y \\ u'_z \\ P' \end{pmatrix}. \quad (3.27)$$

Now, to use spectral methods to solve the system of equations given by Equation 3.27, we approximate each fluid variable $X'(z)$ as a truncated series with Whittaker cardinal functions as the basis

$$X'(z) \approx X'_N(z) = \sum_{j=-N/2}^{N/2} a_j C_j(z), \quad (3.28)$$

where $X' \in \{\rho', u'_x, u'_y, u'_z, P'\}$ and $C_j(z)$ are the Whittaker cardinal (sinc) functions. We can then construct an operator matrix $\underline{\underline{A}}$ consisting of block matrices for the linearized equations of gasdynamics. Finally, we solve the eigenvalue problem $\sigma \mathbf{U} = \underline{\underline{A}}\mathbf{U}$ for the eigenvalues σ and eigenvectors $\mathbf{U} = (\rho', \mathbf{u}'_x, \mathbf{u}'_y, \mathbf{u}'_z, P')^T$ numerically using a code we have written in PYTHON. (We refer to this code as our *spectral eigensolver*.)

3.2.2 WKBJ approach

We caution the reader that in this subsection only we employ the opposite sign convention to that employed in Section 3.1.1 (and in the remainder of this

thesis), a choice that was made in order to check our calculations more easily against those of Ruden et al. (1988) who carried out a similar calculation using a polytropic equilibrium. Thus in this subsection $N_B^2(z) < 0$ denotes convective *stability* whereas $N_B^2(z) > 0$ denotes convective *instability*.

Equations 3.12-3.16 can be combined into a second-order ODE of the form

$$\frac{d^2 U}{dz^2} + k_z^2(z)U = 0, \quad (3.29)$$

where

$$U = \left(\frac{\gamma P_0}{\sigma^2 + \kappa^2 + k_x^2 c_s^2} \right)^{1/2} u'_z, \quad (3.30)$$

in which $c_s(z)$ is the adiabatic sound speed, and the ‘vertical wavenumber’ of the perturbations is given by

$$k_z^2(z) \approx k_x^2 \frac{N_B^2(z) - \sigma^2}{\kappa^2 + \sigma^2}, \quad (3.31)$$

where $k_x = 2\pi/\lambda_x$ is the radial wavenumber of the perturbations (Ruden et al., 1988), and $\kappa^2 = 2\Omega(2\Omega - S)$ is the square of epicyclic frequency. In our calculations we assume a Keplerian disk, for which $S = (3/2)\Omega$ and thus $\kappa = \Omega$.

We obtain approximate solutions to Equation 3.29 analytically using WKBJ methods in the limit where $k_x H \gg 1$. Equation 3.31 has two turning points at $k_z^2(z) = 0$, which occur when $N_B^2(z) - \sigma^2 = 0$. Physically, one turning point occurs at the disk mid-plane ($z = 0$), while the other turning point occurs at the boundary of the convectively unstable region $z = L_c$. When $N_B^2(z) - \sigma^2 > 0$ then $k_z^2(z) > 0$ so the solutions of Equation 3.29 are spatially oscillatory, otherwise they are evanescent. The unstable modes, in which we are interested, are confined and oscillatory within the convectively unstable region, and spatially decay exponentially outside.

The standard WKBJ solution is given by

$$U(z) \sim c_1 \exp \left(+i \int^z k_z(z') dz' \right) + c_2 \exp \left(-i \int^z k_z(z') dz' \right),$$

where c_1 and c_2 are constants. By matching the interior (oscillatory) solution onto the exterior (exponentially decaying) solution, we obtain the eigenvalue

equation (dispersion relation):

$$\int_{z_0}^{z_1} k_z(z) dz = \left(n + \frac{1}{2}\right) \pi \quad n = 0, 1, 2, 3, \dots, \quad (3.32)$$

where lower and upper bounds of integration z_0 and z_1 , respectively, are related by the implicit equation $N_B^2(z_0) = N_B^2(z_1) = \sigma^2$, and n is the mode number. Substituting for $k_z^2(z)$ using Equation 3.31, we obtain an implicit equation relating the radial wavenumber k_x and the growth rate σ which we solve numerically via a root-finding algorithm.

Substituting Equation 3.31 into the left hand side of the dispersion relation Equation 3.32 obtains

$$\text{LHS} = \frac{k_x}{\sqrt{\kappa^2 + \sigma^2}} \int_{z_0}^{z_1} \sqrt{N_B^2(z) - \sigma^2} dz. \quad (3.33)$$

In order to place the integral in Equation 3.33 in dimensionless form, we make the substitution $z = yL_c$ to obtain

$$\text{LHS} = \frac{k_x L_c}{\sqrt{\kappa^2 + \sigma^2}} \int_{z_0/L_c}^{z_1/L_c} \sqrt{N_B^2(yL_c) - \sigma^2} dy. \quad (3.34)$$

The buoyancy frequency $N_B^2(z)$ for our Gaussian temperature vertical profile (see Section 3.1.1) is given by

$$\frac{N_B^2}{\Omega^2} = \frac{2}{\beta H^2} z^2 \left[1 - \left(1 - \frac{1}{\gamma} \right) \frac{\beta H^2 \Omega^2}{2T_0} e^{z^2/\beta H^2} \right], \quad (3.35)$$

Here T_0 , β and the ratio of specific heats γ are parameters with which we can control the width and depth of the buoyancy frequency profile.

Rescaling the vertical coordinate such that $z \rightarrow z/L_c$ we obtain

$$\frac{N_B^2}{\Omega^2} = \frac{2}{\beta H^2} \left(\frac{z}{L_c} \right)^2 \left[1 - \left(1 - \frac{1}{\gamma} \right) \frac{\beta H^2 \Omega^2}{2T_0} e^{z^2/\beta H^2 L_c^2} \right]. \quad (3.36)$$

Finally, substituting Equation 3.36 into Equation 3.34 we obtain

$$\text{LHS} = \frac{k_x L_c}{\sqrt{\kappa^2 + \sigma^2}} I(\tilde{\alpha}), \quad (3.37)$$

where the dimensionless integral $I(\tilde{\alpha})$ is given by

$$I(\tilde{\alpha}) \equiv \int_{\tilde{\alpha}}^{\tilde{\beta}(\tilde{\alpha})} \sqrt{N_B^2(yL_c) - \sigma^2} dy, \quad (3.38)$$




Figure 3.5: Numerical solution of the dimensionless integral given by Equation 3.38. Here $\tilde{\alpha} \equiv z_0/L_c$ is defined as the lower turning point z_0 rescaled by the width of the convective region L_c . The vertical profile parameters are given by $\{T_0 = 1.0, \rho_0 = 1.0, \beta = 3.0\}$. The adiabatic index is set to $\gamma = 5/3$.

and the (rescaled) lower and upper limits have been denoted as $\tilde{\alpha} \equiv z_0/L_c$ and $\tilde{\beta} \equiv z_1/L_c$, respectively. The limits of integration are related to each other, and to the growth rate, through the implicit equation $N_B^2(z_0) = N_B^2(z_1) = \sigma^2$. In order to calculate the dimensionless integral $I(\tilde{\alpha})$ we first solve the implicit equation $N_B^2(\tilde{\beta}) = N_B^2(\tilde{\alpha})$ to obtain the upper limit of integration $\tilde{\beta}$. The dimensionless integral can then be solved numerically, and the solution is shown in Figure 3.5.

Finally, we solve

$$\frac{k_x L_c}{\sqrt{\kappa^2 + \sigma^2}} I(\tilde{\alpha}) = \left(n + \frac{1}{2}\right) \pi, \quad (3.39)$$

in order to obtain, for a given mode number n , the growth rate σ as a function of the radial wavenumber k_x . The solutions are discussed in Section 3.2.3, and solutions for the first three modes ($n = 0, 1$, and 2) are indicated by the solid curves in Figure 3.6.

3.2.3 Eigenvalues

For a given radial wavenumber k_x , the eigenvalues calculated with the spectral eigensolver come in pairs that lie close together in the complex plane. These correspond physically to the *even* and *odd* modes of the convective instability. As described in Section 3.2.4, even modes are defined as those for which $u'_z(z)$ is symmetric about the mid-plane, whereas the odd modes are defined as those for which $u'_z(z)$ is antisymmetric about the mid-plane. In Fig. 3.6 we plot as black dashed curves the first few eigenvalue branches as functions of k_x , these corresponding to vertical quantum numbers $n = 0, 1$, and 2 .³ Note that in the figure k_x has been normalized by the width of the convectively unstable region L_c (given by Equation 3.7 in Section 3.1.1).

As the radial wavenumber k_x is increased, the difference between the growth rates of the even and odd modes vanishes, and the growth rates tend asymptotically to the maximum absolute value of the buoyancy frequency (indicated by a horizontal dot-dashed line). In other words, the maximum growth rate is limited by the depth of the buoyancy-frequency profile. The fastest growing modes are at arbitrarily large k_x (small radial wavenumber λ_x), and thus manifest as thin elongated structures. Thicker structures are not favored as they comprise greater radial displacements that are resisted by the radial angular momentum gradient.

Superimposed on the spectrally computed values are the WKBJ growth rates (solid lines). WKBJ methods cannot distinguish between even or odd modes, nonetheless we find very close agreement between the WKBJ and semi-analytical results for all radial wavenumbers, even at low k_x where the WKBJ approximation is, strictly speaking, invalid. As the radial wavenumber increases, the semi-analytical and WKBJ results converge. Note that associated with each mode branch is a minimum radial wavenumber $k_{x,\min}$ below which both the numerical and WKBJ growth rates are zero. This

³Note that the designations ‘ $n = 0$ (even)’ and ‘ $n = 0$ (odd)’ are due to how we decided to label the modes. In our calculation of the growth rates we observe pairs of modes with very similar growth rates. The first pair are labeled ‘ $n = 0$ (even)’ and ‘ $n = 0$ (odd)’ rather than $n = 0$ and $n = 1$, respectively, the next pair are labeled ‘ $n = 1$ (even)’ and ‘ $n = 1$ (odd)’ rather than $n = 2$ and $n = 3$, respectively, and so on.

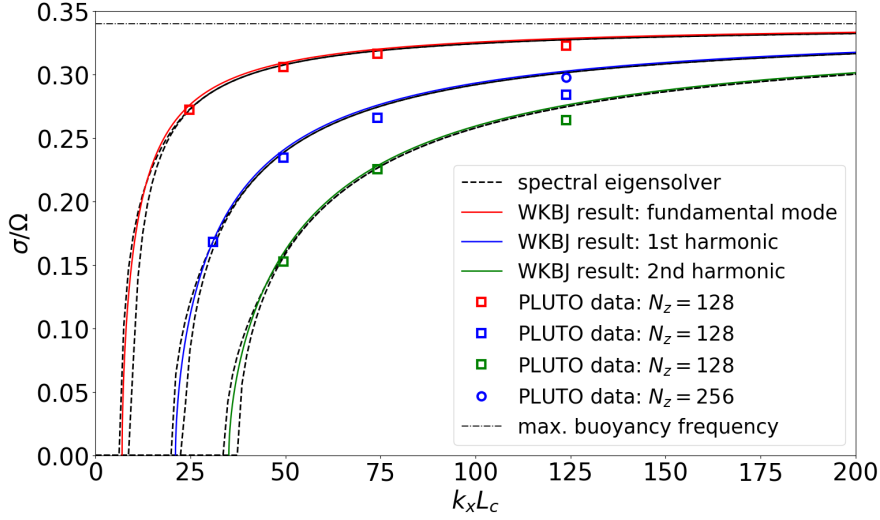


Figure 3.6: Growth rates as a function of radial wavenumber k_x (scaled by the width of the convectively unstable region L_c). The solid lines were calculated analytically using WKBJ methods, and the dashed lines were calculated semi-analytically using pseudo-spectral methods. The squares correspond to measurements taken from PLUTO simulations run at a vertical resolution of $N_z = 128$. The blue circle was taken from a PLUTO simulation run at $N_z = 256$.

feature is just visible in the bottom left-hand side of Figure 3.6.

We must stress that, because the fastest growing modes are on the shortest possible scales, inviscid simulations of convection are problematic, at least in the early (linear) phase of the evolution. It is here that the simulated behavior may depend on the varying ways that different numerical schemes deal with grid diffusion, something we explore in Section 3.3.3. Only with resolved physical viscosity can this problem be overcome. However, in the fully developed nonlinear phase of convection the short scale linear axisymmetric modes may be subordinate and this is less of an issue.




Figure 3.7: Vertical profiles of the perturbations $\text{Re}\{\rho'(z)\}$, $\text{Re}\{u'_z(z)\}$, $\text{Re}\{P'(z)\}$, and of $\text{Im}\{u'_x(z)\}$ and $\text{Im}\{u'_y(z)\}$ in the upper half plane $z \in [0, 2]H$, for the $n = 0$ odd mode at $k_x = 20.0$. The imaginary and real parts not shown are effectively zero. The black dot marks the upper extent of the convectively unstable region L_c , while the red dot marks the most convectively unstable points (see Equation 3.7). The eigenfunctions have been rescaled so that the maximum of each eigenfunction is unity.

3.2.4 Vertical structure of the eigenfunctions

The reader is reminded that here we have an unstably stratified fluid in a medium in which the gravitational field smoothly changes sign. This contrasts to the more commonly studied situation of convection at the surface of the Earth (or within the solar interior). It hence is of interest to study the structure of the convective cells in some detail. We consider first the fundamental ($n = 0$) odd mode with radial wavenumber $k_x = 20$.

In Figure 3.7 we plot the vertical profiles of the nonzero components of the eigenfunction. To provide greater visual clarity, we have rescaled the

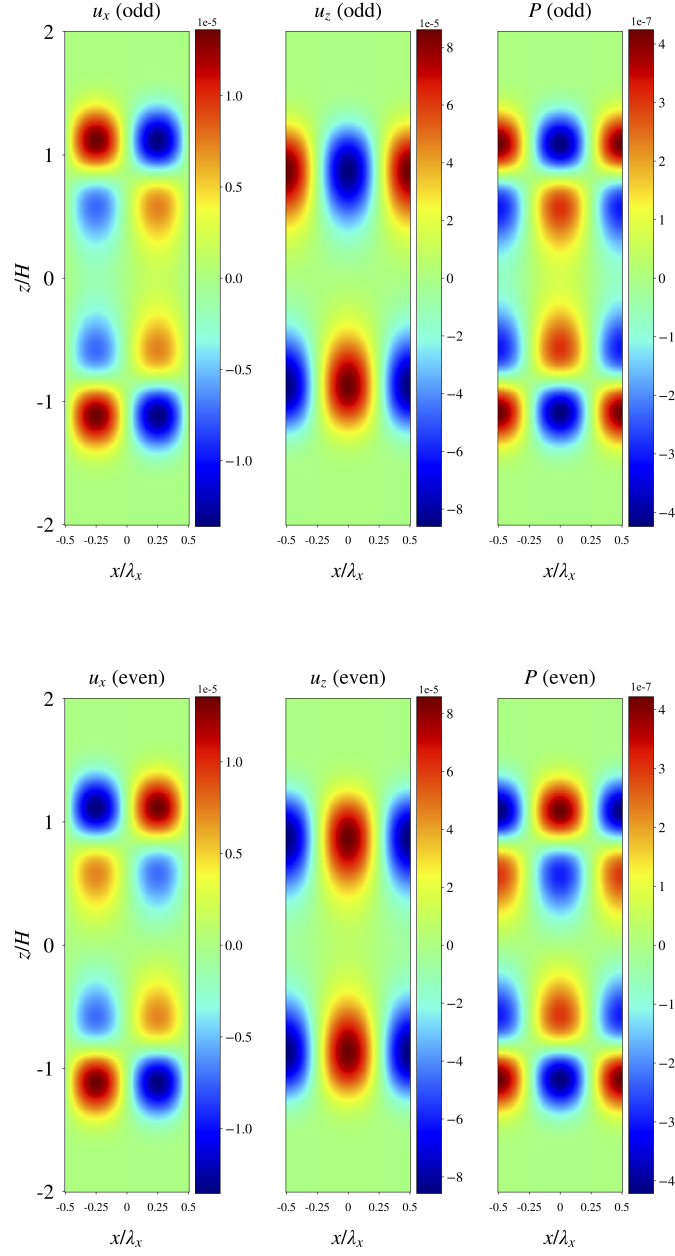


Figure 3.8: Two-dimensional profiles in the xz -plane of the real parts of the selected components of odd and even eigenfunctions for $n = 0$ and $k_x = 20.0$.

perturbations so that the maximum amplitude of each perturbation is unity: thus we can more easily observe the vertical structure of the perturbations, but not deduce their relative amplitudes. In Figure 3.8 we plot their real parts in the xz -plane with the correct amplitudes. We find that $|u'_z|/c_{s0} > |u'_y|/c_{s0} \sim |u'_x|/c_{s0} \gg |\rho'|/\rho_0 \gg |P'|/P_0$, where $c_{s0} = 1$, $\rho_0 = 1$, and $P_0 = 1$ are the background sound speed, density, and pressure at the mid-plane. Thus vertical velocity perturbations are greatest in magnitude, while pressure perturbations are smaller than vertical velocity perturbations by two-orders of magnitude, indicating that the convective cells are roughly in pressure balance with the surroundings.

In Figure 3.7, u'_z (solid black curve) is antisymmetric about the mid-plane, whereas the remaining perturbations are symmetric about the mid-plane. Figures 3.7 and 3.8 clearly show that both odd and even modes consist of a chain of convective cells above and below the mid-plane and localized near the most convectively unstable point (denoted by the large red dot in the former Figure). The peak in the amplitude of u'_z that occurs near this point tells us that the fluid elements reach their maximum acceleration where the buoyancy force is greatest. As the fluid perturbations move subsonically, they adjust their pressure to maintain a balance with the background pressure: thus where cool elements begin to rise (higher background pressure), the pressure perturbations (solid blue line) increase (i.e. $P' > 0$), and where hot elements begin to sink (lower background pressure), the pressure perturbations decrease (i.e. $P' < 0$). The vertical velocity perturbation (black solid line) is out of phase with the radial and azimuthal velocity perturbations (dashed lines), since vertical motion is converted into radial motion where the convective cells turn over (picture fluid motion at the top of a fountain).

The even modes possess a nonzero, but relatively small, u'_z at the mid-plane, and thus weakly couple the two sides of the disk. Such modes may permit some exchange of mass, momentum and thermal energy across the mid-plane when reaching large amplitudes.

As the radial wavenumber is increased from $k_x = 20.0$ to $k_x = 200.0$ (very small radial wavelengths), the perturbations become increasingly localized about the most convectively unstable point(s) (see Figure 3.9). The




Figure 3.9: Vertical profiles of the perturbations $\text{Re}\{\rho'(z)\}$, $\text{Re}\{u'_z(z)\}$, $\text{Re}\{P'(z)\}$, and of $\text{Im}\{u'_x(z)\}$ and $\text{Im}\{u'_y(z)\}$ in the full plane $z \in [-2, 2]H$, for the $n = 0$ odd mode at $k_x = 400.0$, showing localization of fast growing, short wavelength modes to the upper disk half-plane.

localization in z however is not as narrow as the radial wavelength. For larger k_x , in fact, each pair of even and odd modes change their character and become entirely localized to either the upper or lower disk. Activity in the two halves of the disk are hence completely decoupled for such small-scale (but fast growing) disturbances.

3.2.5 Comparison of theory with simulations

In this section we compare the growth rates previously calculated with those measured from PLUTO simulations initialized with the exact linear modes taken from the spectral eigensolver.⁴

Each PLUTO simulation is run in a box of size $\lambda_x \times 0.625H \times 4H$. We

⁴To estimate the numerical growth rates we first plotted (on a semi-log plot) the vertical kinetic energy density against time, and then used a least squares fit to measure the slope of the kinetic energy during the linear phase.

Table 3.1: Comparison of eigensolver growth rates with growth rates measured in PLUTO simulations. The simulations were initialized with the eigenfunctions corresponding to each mode n and radial wavenumber k_x . All simulations except the one marked with a footnote were run at a resolution of $64 \times 10 \times 128$. The standard deviation on the growth rates σ_{PLUTO} measured in the simulations is typically in the range ± 0.0002 to ± 0.0004 .

Mode	Parity	k_x	σ_{EIG}	σ_{PLUTO}	% error
$n = 0$	even	20.0	0.2727	0.2719	0.3
	odd	20.0	0.2724	0.2715	0.3
	even	40.0	0.3076	0.3059	0.6
	even	60.0	0.3186	0.3160	0.8
	even	100.0	0.3273	0.3223	1.6
$n = 1$	even	25.0	0.1694	0.1678	0.9
	odd	25.0	0.1643	0.1624	0.9
	even	40.0	0.2392	0.2343	1.1
	even	60.0	0.2747	0.2655	3.3
	even	100.0	0.3016	0.2840	5.8
$n = 2$	even	40.0	0.1593	0.1525	4.2
	odd	40.0	0.1580	0.1517	3.9
	even	60.0	0.2272	0.2251	0.9
	even	100.0	0.2749	0.2637	4

* Simulation run with a resolution of $N_z = 256$ grid points in the vertical direction.

maintain a fixed number of 32 grid cells per scale height in the z -direction, 16 grid cells per scale height in the y -direction and 64 grid cells per *radial wavelength* in the x -direction. Thus the radial wavelength of the modes is always well resolved in our simulations. No explicit diffusion coefficients (i.e. viscosity or thermal diffusivity) are used.

The results are plotted in Figure 3.6 as squares (and one circle). Numerical values appear in Table 3.1. We find excellent agreement between simulation and theory: the relative error is $< 1\%$ for the fundamental mode in the interval $k_x \in [20.0, 60.0]$ and for the first harmonic at $k_x = 25.0$. As the radial wavenumber is increased the simulation growth rates begin to deviate from the theoretical curves, although the percentage error remains less than 6% even at large wavenumbers. This behavior is expected due to

the effects of numerical diffusion which acts to *decrease* the growth rates as the radial wavelength approaches the size of the vertical grid. Although the radial wavenumber λ_x is always well resolved in the simulations, the vertical wavelength becomes increasingly less well resolved.

The numerical diffusion could be reduced by increasing the vertical resolution of the simulations. To check this we have rerun the PLUTO simulation initialized with eigenfunctions corresponding to the first harmonic ($n = 1$) even mode with $k_x = 100$ at *twice* the vertical resolution, i.e. $N_z = 256$ instead of $N_z = 128$. At a resolution of $N_z = 128$ we measured a growth rate of $\sigma_{\text{PLUTO}} = 0.2840 \Omega$, corresponding to a percentage error of 5.8% with the growth rate calculated semi-analytically ($\sigma_{\text{EIG}} = 0.3016 \Omega$). At a vertical resolution of $N_z = 256$, however, we measured a growth rate of $\sigma_{\text{PLUTO}} = 0.2974 \Omega$, corresponding to a percentage error of just 1.4%. Thus, the growth rates measured in the simulations appear to converge to those calculated analytically and semi-analytically as the resolution of the simulations is increased.

Finally, in all simulations of the axisymmetric modes we observed *inward* angular momentum transport. This is in agreement with the analytical argument presented in (Stone & Balbus, 1996) regarding axisymmetric flow.

3.3 Simulations of unforced compressible convection

In this section we describe fully compressible, three-dimensional simulations of hydrodynamic convection in the shearing box carried out with PLUTO. These simulations are ‘unforced’ in that they are initialized with a convectively unstable profile, but this unstable profile is not maintained by a separate process (e.g. turbulent heating via the magnetorotational instability). Thus convection rearranges the profile, shifting it to a marginally stable state, and thus ultimately hydrodynamical activity dies out. While the unforced convection studied in this section is essentially a transient phenomenon that depends on initial conditions, and is sensitive to the numerical scheme (as

we discuss in Sections 3.3.2 and 3.3.3), it serves as a good starting point for investigating the problem and illuminates a number of interesting features.

The key result of this section is that, even without the inclusion of explicit viscosity, we observe that hydrodynamic convection in the shearing box generally produces a *positive Reynolds stress*, and thus can drive outward transport of angular momentum. This is in direct contradiction to simulations carried out in ZEUS by SB96. In Section 3.3.1 we describe our fiducial simulations, and in Section 3.3.2 we compare ATHENA and PLUTO simulations. In Section 3.3.3 we describe the sensitivity of the sign of angular momentum transport to the numerical scheme. Finally, in Section 3.3.4 we investigate the effects of including explicit diffusion coefficients in our simulations, thus connecting to the Boussinesq simulations of Lesur & Ogilvie (2010).

3.3.1 Fiducial simulations

Set-up and initialization of fiducial simulations

The simulations described in this section were run at resolutions of 64^3 , 128^3 and 256^3 in boxes of size $4H \times 4H \times 4H$, where H is the scale height.⁵ All the simulations were initialized with the convectively unstable vertical profiles for density and pressure described in Section 3.1.1 with profile parameters $T_0 = 1.0$, $\rho_0 = 1.0$, $\beta = 3.0$ and an adiabatic index of $\gamma = 5/3$ (see Figure 3.2). The Stone and Balbus vertical profile SB96 was also trialed yielding very similar results, but we have omitted most of these in the interests of space. Vertical outflow conditions were employed for our fiducial simulations but periodic and free-slip boundary conditions in the vertical direction produced the same behaviour. Random perturbations to the velocity components were seeded at initialization, with a maximum amplitude $|\delta \mathbf{u}| \sim 10^{-5} c_{s0}$. Finally, we adopt a very small but finite thermal diffusivity of $\chi = 2 \times 10^{-6}$ to facilitate conduction of thermal energy through the vertical boundaries and to aid code stability. No explicit non-adiabatic heating, cooling, or thermal

⁵A simulation run at a resolution 512^3 is included in the list of fiducial simulations in Table B.1 (see Appendix B) but is not described in this section for brevity.




Figure 3.10: Left: semi-log plot of the time-evolution of the volume-averaged vertical kinetic energy density in PLUTO simulations at resolutions of 64^3 (blue line), 128^3 (red line) and 256^3 (black line). Right: time-evolution of volume-averaged xy -component of the Reynolds stress tensor.

relaxation are included in the simulations described in this section: therefore convection gradual dies away after non-linear saturation.

Time-evolution of averaged quantities

In the left panel of Figure 3.10 we track the time-evolution of the volume-averaged kinetic energy density associated with the vertical velocity for simulations run at resolutions of 64^3 , 128^3 and 256^3 , respectively.

Initially all simulations exhibit small-amplitude oscillations due to internal gravity waves excited in the convectively stable region at initialization. After some three orbits, the linear phase of the convective instability begins in earnest, characterised by exponential growth of the perturbation amplitudes. During this phase, internal energy is converted into the kinetic energy of the rising and sinking fluid motion that comprises the convective cells. As the resolution is increased, shorter scale modes are permitted to grow. Because they are the most vigorous the growth rate (proportional to the slope of the kinetic energy density) is slightly larger at better resolution.

The linear phase ends some 10 orbits into the simulation, and the flow becomes especially disordered. The peak in the vertical kinetic energy density occurs at this point, but this peak decreases with resolution, something we

discuss in Section 3.3.1. After non-linear saturation, the kinetic energy decreases gradually: the convective cells redistribute thermal energy and mass, thus shifting the thermal profile from a convectively unstable to a marginal state, and ultimately the convective motions die out. After about 38 orbits, the kinetic energy density has dropped to about one hundredth of its value at non-linear saturation in the lower resolution runs. The level of numerical diffusivity has an appreciable effect in damping activity after non-linear saturation, with the decrease in vertical kinetic energy successively smaller over the same period of time in the 64^3 , 128^3 and 256^3 simulations, respectively.

The behavior of the kinetic energy aligns relatively well with our expectations. The Reynolds stress, on the other hand, is more interesting. The evolution of the xy -component of the Reynolds stress is plotted in the right panel of Fig. 3.10. The stress is small, but perhaps its most striking feature is its positivity over the entire duration of all three simulations.

Perhaps most surprising is the positivity of the stress during the *linear* phase of the instability. We might expect that the axisymmetric modes (which send angular momentum inwards) dominate this period of the evolution, as non-axisymmetric disturbances only have a finite window of growth (about an orbit) before they are sheared out and dissipated by the grid. But the positivity of the stress suggests that instead it is the shearing waves that are the dominant players in the linear phase. Visual inspection of the velocity fields confirms that the flow is significantly non-axisymmetric, and we find several examples of strong shearing waves ‘shearing through’ $k_x = 0$ during this phase. It would appear these waves transport angular momentum primarily outward as they evolve from leading to trailing, behavior that is in fact consistent with Ryu & Goodman (1992), who find that inward transport only occurs at sufficiently long times when the waves are strongly trailing and hence effectively axisymmetric. (Even then the stress is extremely oscillatory.)

Why do the shearing waves dominate over the axisymmetric modes? One hypothesis is that 3D white noise seeds fast growing shearing waves that can outcompete the axisymmetric convective instability over their permitted window of growth. If this were to be true then the simulated nonlinear regime

is achieved before the dominant shearing waves become strongly trailing and begin to send angular momentum inward. Given that the typical timescale for shearing out is roughly a few orbits at best, this is marginal but not impossible. An alternative, more plausible, and more troubling hypothesis is that our inviscid numerical code misrepresents trailing shearing waves: more specifically, the code is artificially reseeding fresh leading waves from strongly trailing waves (‘aliasing’; Geoffroy Lesur, private communication). As a consequence, shearing waves outcompete the axisymmetric modes because they can shear through several times. If true, this is certainly concerning. But we hasten to add that this phenomenon should only be problematic in the low amplitude linear phase; once the perturbations achieve large amplitudes the aliasing will be subsumed under *physical* mode-mode interactions.

The linear phase ends in a spike in the stress. Time-averages of the volume-averaged value of R_{xy} from orbit 5 to orbit 15, a period spanning non-linear saturation, show that the Reynolds stress decreases as the resolution increases (see Table B.1 in Appendix B). The time- and volume-averaged value of R_{xy} over a period spanning non-linear saturation (orbits 5 to 15) is $+4.8 \times 10^{-6}$, $+2.7 \times 10^{-6}$ and $+1.7 \times 10^{-6}$ in the 64^3 , 128^3 and 256^3 simulations, respectively. In terms of the turbulent alpha-viscosity this corresponds (from Equation 2.67, but including the shear parameter q in the definition) to $\alpha \sim +2.3 \times 10^{-5}$, $+1.5 \times 10^{-5}$ and 9.6×10^{-6} , respectively. The dependence on the numerical viscosity can be explained by appealing to secondary shear instabilities (see next subsection).

The volume-averaged density remains roughly constant during the linear phase of the instability, followed by a drop after non-linear saturation as mass is lost through the lower and upper boundaries. Although the convectively unstable region at initialization is confined to the region $|z| < L_c \sim 1.11H$, and is therefore well within the box, convective overshoot deposits mass (and heat) outside the convectively unstable region. The overall decrease in density over the duration of the simulation is small, but appears to increase with greater resolution. The overall percentage change in the density is -1.4% , -3.4% and -8.8% in the 64^3 , 128^3 and 256^3 simulations, respectively.




Figure 3.11: From left to right: snapshots of u_z in the xz -plane taken at non-linear saturation (just after the linear phase) at resolutions of 64^3 , 128^3 and 256^3 , respectively. Blue denotes $u_z < 0$ (sinking fluid) and red denotes $u_z > 0$ (rising fluid).

Structure of the flow

The development of convective instability and the associated convective cells are best observed through snapshots of the vertical component of the velocity in the xz -plane, shown at resolutions of 64^3 , 128^3 and 256^3 just after non-linear saturation in Figure 3.11. A full set of convective cells is clearly visible at all resolutions. They are thin, filamentary structures several grid cells wide comprising alternating negative and positive velocities (updrafts and downdrafts).

The maximum vertical Mach number of the flow around non-linear saturation in the 256^3 run is about $M_z \sim 0.15$, with the largest vertical Mach numbers generally being measured near the vertical boundaries (where the temperature – and therefore the sound speed – is lowest).

The higher resolution simulations indicate that the plumes develop a wavy or buckling structure as they rise or sink, indicating the onset of a secondary shear instability. It is likely that the buckling of the convective plumes by these ‘parasitic modes’ limits the amplitudes of the linear modes, and ultimately leads to their breakdown. At lower resolution numerical diffusion smooths out the secondary shear modes, and so the convective plumes reach

large amplitudes before breaking down (blue curve in right panel of Fig. 3.10). At high resolution the shear modes are not so impeded and make short work of these structures (black curve in the same figure).

Vertical heat and mass flux

Figure 3.12 (a) shows the vertical profiles of horizontally-averaged heat and mass fluxes taken from snapshots just after non-linear saturation from the simulation with resolution 256^3 . For clarity, we have time-averaged the horizontally-averaged vertical mass and heat flux profiles between orbits 5 and 10, a period spanning non-linear saturation (see Figure 3.10).

Negative (positive) values for the fluxes for $z < 0$ and positive (negative) values for the fluxes for $z > 0$ correspond to transport of heat and mass away from (towards) the mid-plane. Overall, the heat flux is away from the mid-plane, peaking in the vicinity of the most convectively unstable points at initialization (indicated by the vertical dashed red lines in Figure 3.12). In addition, there is some mass flux towards the mid-plane within $-H < z < H$. Thus, convection is transporting mass and heat such as to erase the convectively unstable stratification, as expected. Note that the positive peaks in the heat flux near the most convectively unstable points are similar to those observed by Hirose et al. (2014) in the thermally dominated region of the disk (i.e. $P_{\text{thermal}} > P_{\text{magnetic}}$; cf. the middle-panel of Figure 5 in their paper).

Beyond $|z| > H$, the mass flux is directed away from the mid-plane, peaking just beyond the point which marks the boundary of the convectively unstable region. This outward mass flux might be due to convective overshoot, although we expect this effect to vanish if averaged over a suitable time interval. Alternatively, it is possible that heat transported towards the corona by the convective cells causes matter in the corona to heat up and become buoyant, or else that this heat generates a weak thermal wind.

Finally, in Figure 3.12 (b) we show the vertical profile of the horizontally-averaged R_{xy} taken from the 256^3 simulation and time-averaged between orbits 5 and 10. The Reynolds stress is clearly positive over all of the vertical




Figure 3.12: (a) Vertical profiles of the vertical heat flux (solid red line) and mass flux (solid black line). The vertical dashed blue line marks the boundary of the convectively unstable region at initialization, while the vertical dashed red line marks the most convectively unstable point at initialization. (b) Vertical profile of the xy -component of the Reynolds stress tensor R_{xy} . In all cases, the results have been time-averaged between orbits 5 and 10, spanning non-linear saturation.

domain, peaking just beyond the most convectively unstable points.⁶ Thus the bulk of outward angular momentum transport occurs where the convective cells begin to turn-over, resulting in radial mixing of the gas. Note, however, that we also observe rather large positive stresses near the vertical boundaries, with the stress at the vertical boundaries about 1.5 times the peak stress in the remainder of the domain.

3.3.2 Simulation of hydro convection in ATHENA

In the previous section we found that hydrodynamic convection in a vertically stratified shearing box in PLUTO (without explicit viscosity) can drive outward angular momentum transport. Here we verify this result using the finite-volume code ATHENA (Stone et al., 2008; Stone & Gardiner, 2010). To facilitate as close a comparison as possible between the two codes and

⁶Note however that at any given instant in time, we generally do observe regions where the Reynolds stress is negative.

also to the ZEUS runs of SB96, we initialize both codes with Stone and Balbus’s convectively unstable profile (see Section 3.3.1). Explicit diffusion coefficients were omitted and vertical periodic boundary conditions implemented. Both simulations were run in a shearing box of size $4H \times 4H \times 4H$ and at a resolution of $64 \times 64 \times 64$. PLUTO and ATHENA offer somewhat different suites of numerical schemes: here we settle on a combination that is slightly more diffusive than that employed in our fiducial simulations because this combination allows for as close a match as possible. Specifically, the numerical scheme employed in ATHENA is (second-order) piecewise linear interpolation on primitive variables, the HLLC Riemann solver and MUSCL-Hancock integration. In PLUTO we use (second-order) piecewise linear interpolation on primitive variables together with a Van Leer limiter function, the HLLC Riemann solver and MUSCL-Hancock integration. The angular frequency, and the sound speed at initialization, in both simulations were set to $\Omega = 10^{-3}$ and $c_s = 10^{-3}$, respectively.⁷

We find that angular momentum transport is directed *outwards* in both codes, demonstrating that the outward transport of angular momentum by hydrodynamic convection in the non-linear phase is robust to a change of code. Figure 3.13 compares the time-evolution of the volume-averaged Reynolds stress taken from the ATHENA simulation with that taken from the PLUTO simulation. Both simulations exhibit exponential growth followed by non-linear saturation, together with the development of convective cells (not shown). These results also contrast with the SB96 runs with ZEUS and thus demonstrate that the positive transport reported in previous sections is not special to the Gaussian temperature profile.

Although for this particular combination of schemes the overall result is the same, we noticed that different schemes resulted in qualitative differences between the two codes. For example, when we use the HLLC solver, piecewise parabolic reconstruction (PPM) and Corner-Transport-Upwind (CTU) inte-

⁷Note that the units differ from the values of $\Omega = 1$ and $c_s = 1$ employed elsewhere in this thesis. This is because we were unable to get ATHENA to work with units of $\Omega = 1$ and $c_s = 1$. Therefore in this section *only* we have used units of $\Omega = 10^{-3}$ and $c_s = 10^{-3}$ in both the PLUTO and ATHENA simulation.




Figure 3.13: Comparison of time-evolution of volume-averaged Reynolds stress taken from a simulation run in ATHENA (black) with the same quantity taken from a simulation run in PLUTO (red). Note that in these two simulations *only* we have employed units of $\Omega = 10^{-3}$ and $c_s = 10^{-3}$.

gration, ATHENA exhibits delayed onset of instability, a more gradual drop in kinetic energy density following non-linear saturation, and a slower drop in angular momentum transport compared to PLUTO. When we use the Roe solver, PPM reconstruction and CTU integration, the Reynolds stress is highly oscillatory in time in both codes, indicative, perhaps, of numerical instability. The different behavior (in both codes) based on which combination of numerical schemes is chosen is worrying, although we emphasize that the differences are probably accentuated by the transient nature of the evolution and its sensitivity to the initial conditions, and the fact that the linear phase is partially controlled by grid diffusion. The robustness of the results to changes in numerical scheme is investigated in more detail in the following section.




Figure 3.14: Left: time-evolution of volume-averaged Reynolds stress taken from a PLUTO simulation run with the less diffusive HLLC solver. Right: the same, but taken from a simulation run with the more diffusive HLL solver. The change in the sign of the stress tensor provides compelling evidence that whether convection can transport angular momentum inwards or outwards depends on the diffusivity of the underlying numerical scheme.

3.3.3 Sensitivity of sign of angular momentum transport to numerical scheme

An important result of Sections 3.3.1-3.3.2 is that purely hydrodynamic convection in PLUTO and in ATHENA resulted in $R_{xy} > 0$, i.e. outward angular momentum transport. This is in disagreement with the ZEUS results of SB96, who reported a Reynolds stress of $R_{xy} < 0$. Given that ZEUS is a non-conservative, finite-difference code, our hypothesis for explaining the discrepancy is that ZEUS run at comparatively low resolution (as in SB96) is sufficiently diffusive that it imposes an artificial near-axisymmetry on the flow (which will send angular momentum inward). In this section we report our attempts to assess the numerical diffusivity of various algorithms and their impact on convection.

First we reran the fiducial simulations described in Section 3.3.1 but with different combinations of numerical schemes. We found that the sign of angular momentum transport is indeed sensitive to our choice of scheme. This is best illustrated in Figure 3.14: the left-hand panel shows the time-evolution of the volume-averaged R_{xy} taken from a simulation run with the HLLC Rie-

mann solver (simulation NSTR22e11a in Table B.2), while the right-hand panel shows the same quantity but taken from a simulation run with the (more diffusive) HLL solver (simulation NSTR22e12a in Table B.2). Both simulations exhibit similar exponential growth in the vertical kinetic energy during the linear phase of the instability, and the velocity field shows the development of convective cells in both cases, but the sign of the Reynolds stress is radically different. The HLL run is also far more laminar and axisymmetric. We next repeated the HLL runs with higher resolutions, up to 256^3 , but with no change in the sign of R_{xy} during the linear phase. It must be stressed that going to higher resolutions does not necessarily help in the problem of convection; this is because the fastest growing linear modes are always near the grid scale and hence it is impossible (in the linear phase at least) to escape grid effects. Following non-linear saturation, however, we do briefly observe outward transport ($R_{xy} > 0$) in our 256^3 HLL run, though we observe this only for a few orbits before the entropy gradient is erased and convection ceases (after which R_{xy} drops and fluctuates around zero).

The results of different combinations of schemes is summarized in Table B.2 in Appendix B. They indicate that the sign of R_{xy} appears to be robust to changes in the interpolation and time-stepping schemes, but sensitive to the Riemann solver. In particular, the less diffusive Riemann solvers (Roe and HLLC) gave $R_{xy} > 0$, while the more diffusive Riemann solvers (HLL and a simple Lax-Friedrichs solver) gave $R_{xy} < 0$. Altogether we explored twelve different configurations of Riemann solver, interpolation scheme, and time-stepping method. These ranged from the most accurate (least diffusive) set-up which was identical to that employed in the simulations described in the previous section (a Roe Riemann solver, third-order-in-space WENO interpolation and third-order-in-time Runge-Kutta time-stepping), to the least accurate (most diffusive) set-up (a simple Lax-Friedrichs Riemann solver, second-order-in-space linear interpolation and second-order-in-time Runge-Kutta time-stepping).

Thus we have compelling evidence that, as suspected, angular momentum transport due to convection can be sensitive to the diffusivity of the underlying numerical scheme. It appears that over-smoothing of the flow

by diffusive Riemann solvers, such as HLL, or by codes such as ZEUS that employ artificial viscosity and finite-differencing of the pressure terms, impose a spurious axisymmetry on the flow. This axisymmetry picks out the axisymmetric convective modes, which in turn transport angular momentum inwards.

3.3.4 Viscous simulations

Given the concerns raised in the last section regarding numerical schemes, as well as the fact that the fastest growing inviscid modes are on the shortest scales, we expand our study to include explicit viscosity (and thermal diffusivity). A properly resolved viscous simulation should exhibit none of the numerical problems encountered above, and the fastest growing mode occurs on a well defined scale above the (resolved) viscous scale. Our main aim in this section is to test whether the results of our fiducial simulations are solid: mainly, if angular momentum transport can be positive in the presence of viscosity. Additionally, the inclusion of explicit diffusion coefficients enables us to investigate the Rayleigh number dependence of fully compressible hydrodynamic convection in the shearing box, and thus connect to previous work by Lesur & Ogilvie (2010).

We carry out a suite of simulations at a resolution mainly of 256^3 investigating the effects on hydrodynamic convection when the Rayleigh number Ra is increased, but the Prandtl number Pr is fixed at 2.5.⁸ Thus we decrease both the viscosity and thermal diffusivity in order to increase the Rayleigh number, but keep their ratio fixed at 2.5. Note however, that within any individual simulation we keep the viscosity and thermal diffusivity constant in both space and time. The Rayleigh numbers of the simulations are $Ra = 10^5, 10^6, 10^7, 10^8, 10^9$, and 10^{10} . (Note that the simulations undertaken at the two highest Ra are probably underresolved, as explored later.) The Richardson number at initialization is fixed by the initial vertical profile, which is described in detail in Section 3.1.1 (and shown in Figure 3.2).

⁸We have also repeated some of the simulations at a Prandtl number of unity. The differences with the $Pr = 2.5$ simulations are nominal.

For the profile parameters chosen, the Richardson number at initialization is $Ri \sim 0.05$. Further details of the simulations are given in Table B.3 in Appendix B.

Rayleigh number dependence

As the Rayleigh number is increased from low to high values the system proceeds through the same sequence of states found by Lesur & Ogilvie (2010). We observe no instability for $Ra = 10^5$. At $Ra = 10^6$ instability occurs but the flow appears relatively laminar and axisymmetric; in particular the Reynolds stress is *negative* throughout the linear and nonlinear phases. We conclude that the critical Rayleigh number for the onset of convection lies in the range $10^5 < Ra_c < 10^6$. At $Ra = 10^7$ the instability is more vigorous and the flow field significantly more chaotic and non-axisymmetric in the nonlinear phase, at which point the Reynolds stress has become positive. We conclude that the critical Ra at which the sign of R_{xy} flips lies between 10^6 and 10^7 . At higher Ra the flow appears even more turbulent and non-axisymmetric. It is a relief that in the nonlinear phase of the instability we find agreement with the inviscid simulations of previous sections at sufficiently high Ra .

Our critical Rayleigh number for the onset of convection is considerably higher than in Lesur & Ogilvie (2010), who report a value of $Ra_c = 6900$, but this is perhaps explained by their much larger Richardson number = 0.2 which remains constant throughout their box. On the other hand, the critical Ra at which R_{xy} changes sign is much closer to ours, which suggests that the breakdown into non-axisymmetry may be controlled by secondary shear instabilities that are more sensitive to the viscosity than the convective driving.

While the transport of angular momentum is unambiguously outward after nonlinear saturation when $Ra \geq 10^7$, it is almost always inwards during the linear phase, as is illustrated by the red curve in Fig. 3.15. Moreover, during this early stage, the simulation is dominated by the axisymmetric modes of Section 3.2. This disagrees with our inviscid fiducial simulations




Figure 3.15: Time-evolution of volume-averaged xy -component of Reynolds stress tensor from two simulations run at resolutions of 256^3 (black curve) and 512^3 (red) curve at a Rayleigh number of $Ra = 10^9$. The vertical dashed lines mark the end of the linear phase in the 256^3 (black) and 512^3 (red) simulations, respectively.

(see discussion in Section 3.3.1), which are non-axisymmetric from the outset. We confess that our viscous simulations are more in tune with our physical intuition: (a) in the linear phase growing axisymmetric modes out-compete shearing waves, (b) at sufficiently large amplitudes the modes are subject to secondary shear instabilities in the xz -plane that buckle the rising and falling plumes, (c) perhaps concurrently or on a short time later (and viscosity permitting), secondary *non-axisymmetric* instabilities also attack the plumes because they exhibit significant shear in the xy -plane as well, (d) at this point, the flow degenerates into something more disordered, and importantly, non-axisymmetric and the Reynolds stress flips sign. We conclude that viscosity preferentially damps shearing waves vis-a-vis axisymmetric modes, or effectively kills off the artificial aliasing of shearing waves in inviscid simulations (if this is present).

Convergence with resolution

A final issue is whether our viscous simulations are adequately resolved. More specifically: above what critical Ra is a grid of 256^3 points inadequate? We conducted simulations at $Ra = 10^8$ with 256^3 and 512^3 grid points and found generally good agreement between the two. The linear growth rates were almost identical and the ultimate nonlinear state statistically similar. The only noticeable difference was in the peak Reynolds stress, which was somewhat larger in the higher resolution run. Overall, we conclude that 256^3 grid points are adequate to resolve a simulation with $Ra = 10^8$.

Things start to deteriorate at a Rayleigh number of $Ra = 10^9$. In Figure 3.15 we plot the time-evolution of the xy -component of the volume-averaged Reynolds stress for two simulations run at 256^3 (black) and 512^3 (red). The lower resolution run possesses no extended period of negative R_{xy} , in contrast to the runs at $Ra = 10^8$. We speculate that at this resolution physical viscosity is subdominant to the grid and non-axisymmetric disturbances are artificially enhanced, probably via aliasing. At 512^3 , the stress is definitely negative in the linear phase and there is a strong negative peak shortly afterward. The physical viscosity is now permitted to work properly and appears to prohibit artificial non-axisymmetric disturbances. As a consequence, the axisymmetric modes preserve their control of the simulation for significantly longer and the simulation is in accord with those at lower Ra . More reassuring is the nonlinear phase a little later in which the two flows closely resemble each other. We conclude that in the linear phase 256^3 is insufficient to describe a simulation of $Ra = 10^9$, but in the later nonlinear phase it is probably adequate.

3.4 Structures in forced compressible convection

In Section 3.3 we initialized our simulations with a convectively unstable vertical profile, but otherwise did not include any source of heating or cooling.

Convection (and to a much lesser extent, conduction) transferred heat and mass vertically so as to zero the buoyancy frequency and send the box into a convectively stable equilibrium. As a result, we were only able to probe non-linear convection for a short period of time. Now we aim to continually sustain convection in order to explore this phase in greater depth. In the absence of self-sustaining convection, this means we have to maintain the convectively unstable profile artificially.

3.4.1 Set-up

SB96 perpetuated convection by forcing the temperature at the mid-plane to adjust to its value at initialization. This strategy, however, raises problems in conservative codes, such as PLUTO and ATHENA, because the energy injection at the mid-plane has no way to leave the box, except through the distant vertical boundaries (in ZEUS numerical losses on the grid supplied a turbulence-dependent ‘cooling’). In practice, we found that the disk heats up to an inordinate level and, more importantly, it settles into a marginally unstable state, rather than a driven convective state.

Rather than forcing the code in this way, we mimic the effects of *both* heating and cooling through thermal relaxation. The idea is to add a source term to the energy equation such that the vertical internal energy profile relaxes to its value at equilibrium on a timescale τ_{relax} . Although this technique is artificial, it serves as a very basic tool for approximating the effects of realistic heating and cooling, as might be supplied by MRI turbulence and radiative transfer. The timescale τ_{relax} then would be the characteristic time that the radiative MRI system achieves thermodynamical quasi-equilibrium. In our simulations, however, we take the relaxation timescale to be equal to the linear growth time of convection.

In addition, to mitigate the effects of mass outflows, we incorporate a mass source term to the simulations. At the end of the n th time-step we calculated the total vertical mass loss through the upper and lower vertical boundaries. We then added this mass back into the box in the $(n+1)$ th time-step with the same vertical profile used to initialize the density (cf. Equation

3.5).

We implement the thermal relaxation term by making slight modifications to PLUTO’s built-in optically thin cooling module. The thermal energy equation is updated during a substep to take into account user-defined sources of heating and cooling. The resolution employed in the simulation described in this section is $256 \times 256 \times 256$ in a box of size $6H \times 6H \times 4H$, corresponding to about 43 grid cells per scale height in the x - and y -directions and 64 grid cells per scale height in the vertical direction. The numerical set-up, boundary conditions and initial conditions are the same as those describe in Section 3.1.1. The thermal relaxation time is taken to be equal to the inverse of the growth rate of the convective instability which we measured to be $\sigma = 0.2404 \Omega$. Thus the thermal relaxation time is $\tau_{\text{relax}} = 4.2 \Omega^{-1}$, i.e. the internal energy is relaxed back to its equilibrium profile on about 0.7 orbits. We have also included an explicit viscosity and thermal diffusivity of $\nu = \chi = 1.075 \times 10^{-5}$ corresponding to a Rayleigh number of $\text{Ra} = 10^9$ and a Prandtl number of $\text{Pr} = 1$.

Finally, we have repeated the simulations in a cube of resolution 64^3 and size $4H \times 4H \times 4H$, as well as at a resolution of $128 \times 128 \times 64$ in a box of size $6H \times 6H \times 4H$, and also with periodic boundary conditions in the vertical direction, and observed similar results. We have also run a simulation with the HLL solver to partially explore any code dependence.

In the left-hand panel of Figure 3.16 we show the time-evolution of the volume-averaged kinetic energy density associated with the vertical velocity. As in Section 3.3, exponential growth in the linear phase is followed by non-linear saturation. The forced simulations, however, do not subsequently decay. Instead the vertical kinetic energy increases at a slower rate until about orbit 20, at which point oscillations in the kinetic energy density begin to develop. The cycles increase in frequency and amplitude until about orbit 41 at which point the system settles into a quasi-equilibrium, with the volume-averaged vertical kinetic energy density oscillating about $\langle E_{\text{kin},z} \rangle \sim 2 \times 10^{-3}$ and the period remaining steady at $\Delta t = 0.88$ orbits.




Figure 3.16: Left: semi-log plot of time-evolution of volume-averaged vertical kinetic energy density in a simulation with thermal relaxation. Right: time-evolution of volume-averaged xy -component of the Reynolds stress tensor. The thermal relaxation time was set to $\tau_{\text{relax}} = 4.2 \Omega^{-1}$.

3.4.2 Large-scale oscillatory cells

Associated with the oscillations are large fluctuations in the xy -component of the Reynolds stress tensor (shown in the right-hand panel of Figure 3.16). Instantaneous fluctuations in $\langle R_{xy} \rangle$ and in $\langle \alpha \rangle$ may be either positive or negative, but the time-averaged values over this cyclical phase (orbit 20 to the end of the simulation) are $R_{xy} \sim +9.9 \times 10^{-6}$ and $\alpha \sim +3.9 \times 10^{-5}$, respectively. Furthermore, comparing the oscillations in the vertical kinetic energy density to the fluctuations in $\langle R_{xy} \rangle$ and also to the volume-averaged gas pressure $\langle P \rangle$ (not shown here) it is apparent that *peaks* in the kinetic energy density are correlated with both $\langle R_{xy} \rangle < 0$ and troughs in $\langle P \rangle$, while *troughs* in the kinetic energy density are correlated with $\langle R_{xy} \rangle > 0$ and peaks in $\langle P \rangle$.

A clearer picture of the behaviour of the system emerges when we study the structure of the flow during the cyclical phase. Figure 3.17 shows snapshots of the z -component of the velocity in the xz -plane. The first panel is taken just after the end of the linear phase (orbit 7.4); the subsequent panels are taken at five successive peaks and troughs in the kinetic energy. As the kinetic energy rises in the non-linear phase, the thin convective cells




Figure 3.17: Snapshots at different times of the z -component of the velocity in the xz -plane taken from a simulation in which convection was sustained using thermal relaxation. Narrow convective cells shown in (a) just after non-linear saturation merge to form large scale structures shown in (b), which are destroyed (c) and recreated (d) in a cyclical manner with the opposite rotation.

with radial wavelengths $\lambda_x \sim 5 \Delta x \sim 0.177H$, where Δx is the size of grid a cell in the x -direction, slowly begin to merge into larger coherent structures which couple the two halves of the disk together. By orbit 20 (start of the cyclical phase), the radial wavelength of the convective cells has increased to $\lambda_x \sim H$, and by the time the quasi-steady equilibrium state sets in (at around orbit 40) the radial wavelength of the convective cells is $\lambda_x \sim 3.4H$.

Comparing the snapshots of u_z to the peaks and troughs in the kinetic energy, it becomes apparent that peaks in the convective energy are associated with the large-scale axisymmetric convective cells, hence $\alpha < 0$. Troughs in the kinetic energy are associated with destruction of those convective cells

and with positive stress (outward angular momentum transport). Thus we are observing large-scale convective eddies that appear to be created and destroyed cyclically. Furthermore, it is evident from Figure 3.17 that after the eddies are destroyed, they are re-formed with the opposite rotation.

The reader may be alarmed that the large eddies extend all the way to the vertical boundaries of the domain, and indeed there is an uncomfortable level of mass loss during this phase. To check that the flows are not an artefact of our box size, we ran additional simulations with a vertical domain of $\pm 3H$. A density floor had to be activated in such runs, which unfortunately compromised our thermal equilibrium and complicated the onset of convection early in the simulation. However, they ultimately settled onto the cyclical state described above, but now the large-scale eddies peter out before reaching the vertical boundaries. As a result, the mass loss drops to negligible amounts. This confirms that these flows are physical and robust, though only marginally contained within our smaller boxes.

Next, to explore any code dependence of this final outcome, we have rerun the simulation with the more diffusive HLL solver. We observe a negative Reynolds stress during the linear phase and well into the non-linear phase. Associated with this is a remarkably axisymmetric flow field, as confirmed by viewing slices of the pressure perturbation δP in the xy -plane at different times. The simulation, nonetheless, enters the cyclical phase during which this axisymmetry is broken. As expected there is a flip in the sign of the Reynolds stress from negative to positive. The behavior thereafter mirrors that of the simulation of forced compressible convection run with the Roe solver: the cyclical creation and destruction of large scale convective cells and an oscillatory Reynolds stress. We conclude that for forced runs, the ultimate quasi-steady state depends negligibly on the algorithm.

3.4.3 Zonal Flows

In Figure 3.18 we plot space-time diagrams in the xt -plane of the yz -averaged pressure perturbation $\delta P(x, t) \equiv (\langle P \rangle_{yz} - \langle P \rangle) / \langle P \rangle$ and of the yz -averaged perturbation to the y -component of the velocity $\langle \delta u_y \rangle_{yz}(x, t) \equiv \langle u_y + (3/2) \Omega x \rangle_{yz}$.




Figure 3.18: Space-time diagrams of yz -averaged y -component of the perturbed velocity δu_y (left) and of the pressure perturbation δP (right). (The perturbed pressure is defined in the first paragraph of Section 3.4.3).

It is immediately evident from Figure 3.18 that around the onset of the cyclical phase (orbit 20) alternating streaks in δP and δu_y begin to develop. Respectively, these mark alternating bands (in x) of high ($\delta P > 0$) and low ($\delta P < 0$) pressure, and of super-Keplerian ($\delta u_y > 0$) and sub-Keplerian ($\delta u_y < 0$) flow.

From Figure 3.18 it is clear that the pressure and velocity perturbations are 90 degrees out of phase, which is characteristic of *zonal flows*. We hesitate, however, to claim that these structures are in geostrophic balance (i.e. that $\partial_x P' \sim \rho_0 \Omega \delta u_y$) because, as is evident from the space-time diagrams,

the flows are not stationary in time but appear to fluctuate over about one orbit – the same timescale over which the large convective cells are created and destroyed. Indeed the creation and destruction of the zonal flows tracks precisely that of the large-scale convective cells, showing clearly that the two phenomena are connected. The axisymmetry observed during the formation of large-scale convective cells and zonal flows is consistent with the inward transport of angular momentum observed while these structures remain coherent: axisymmetric convective modes dominate during the lifetime of these structures and transport angular momentum inwards.

In Figure 3.19 (right-hand panels) we plot snapshots in the xy -plane taken during the cyclical phase of the perturbation of the y -component of the velocity δu_y and of the perturbed pressure δP . The axisymmetric structure of the zonal flows is clearly visible in panels (b) and (d), as is the $\pi/2$ phase difference between the pressure perturbation δP and the perturbed y -component of the velocity δu_y .

3.4.4 Vortices

Because the zonal flows consist of alternating bands of sub- and super-Keplerian motion – with strong shear between the bands – we expect that they could give rise to vortices via the Kelvin-Helmholtz instability (modified by rotation and stratification). In Figure 3.20 we plot a snapshot in the xy -plane taken during the cyclical phase of the z -component of the vertical component of vorticity ω_z . Small but coherent anti-cyclonic blobs of vorticity are observed during the cyclical phase (shown in blue in Figure 3.20), and these occur precisely where the flow transitions from sub-Keplerian to super-Keplerian (i.e. at the edges of the zonal flows). We observe both anti-cyclonic and cyclonic vortices, though the anti-cyclonic vortices greatly outnumber the cyclonic ones, which is consistent with the fact that cyclonic vortices tend to be more unstable in Keplerian shear flows.

The vortices tend to be elongated at the start of the cyclical phase (around orbit 20), with an aspect ratio of about 0.4, and grow increasingly circular with time; in fact, once quasi-steady equilibrium has been reached at around




Figure 3.19: Two-dimensional slices in the xy -plane of the perturbation to the y -component of the velocity δu_y , and of the fractional pressure perturbation δP from a simulation of forced compressible convection at $z \approx 0.5H$. Left column: slices taken from a snapshot just after non-linear saturation (orbit 7.4). Right column: slices taken from snapshot generated during the cyclical phase (orbit 41.6).




Figure 3.20: Two-dimensional slice in the xy -plane of the z -component of the vorticity ω_z taken from snapshot generated at the start of the cyclical phase (orbit 20) and at $z \approx 0.5H$. For clarity we have zoomed in on the upper-right quadrant of the xy -plane.

orbit 41, many of the vortices have aspect ratios approaching unity. In principle various instabilities could influence the evolution of the vortices (Lesur & Papaloizou, 2009), but sufficiently high grid (or explicit) diffusion will inhibit these instabilities. Thus it is likely that the evolution of the vortices in our simulation is governed by other factors, such as the surrounding turbulent flow field or possibly viscosity. Finally, the vortices appear to have limited three-dimensional extent, and we did not observe any vortices that remained coherent for depths exceeding half a scale height.

3.4.5 Discussion

It is tempting to link our results to the large-scale emergent structures in recent simulations of rotating hydrodynamic convection with uniform (rather than Keplerian) rotation and in the Boussinesq (rather than compressible) regime (Julien et al., 2012; Rubio et al., 2014; Favier et al., 2014; Guervilly

et al., 2014). Both Favier et al. (2014) and Guervilly et al. (2014) observe growth in the vertical component of the kinetic energy at high Rayleigh numbers ($Ra \sim 10^7 - 10^9$), as we do. In addition, Favier et al. (2014) and Guervilly et al. (2014) notice the formation of depth-invariant large-scale vortices in the xy -plane of their simulations (i.e. in the plane perpendicular to the rotation axis).

There are both qualitative and quantitative differences between our results and those of Favier et al. (2014) and Guervilly et al. (2014). Although the Rayleigh number of our simulation with thermal relaxation was $Ra = 10^9$, which is consistent with the Rayleigh numbers at which large-scale vortices were observed in the simulations of Favier et al. (2014) and Guervilly et al. (2014), we observe large-scale *convective cells* (in the xz -plane) rather than large-scale vortices (in the xy -plane). We, too, observe anti-cyclonic vortices in the xy -plane, but these are small and do not appear to have any three-dimensional extent compared to the depth invariant vortices observed in the Boussinesq simulations with uniform rotation. It is possible that they are prevented from merging and thus growing in size because of the turbulence associated with repeated destruction of the large-scale convective cells. The strong shear might also inhibit their growth.

Due to the cyclical nature of the large-scale convective cells, it is tempting to link our results to the intermittent convection reported in the MRI shearing box simulations of Hirose et al. (2014) and Coleman et al. (2018)). In our runs the forcing is due to explicit thermal relaxation, whereas in the runs of Hirose et al. (2014) the forcing is self-consistently provided by MRI heating and radiative cooling. Thus the forcing, and in particular the time-scales associated with the forcing are rather different. However, both thermal relaxation and the MRI limit cycles are similar in that they lead to a cyclical build up of heat and its subsequent purge through vertical convective transport, with the role of opacity in the simulations of Hirose et al. (2014) being simply to modulate this cycle. Thus our results demonstrate that strongly non-linear hydrodynamic turbulent convection has a cyclical nature that might be generic, and that might therefore be robust to the inclusion of more realistic thermodynamics.

3.5 Conclusions

Motivated by recent radiation magnetohydrodynamic shearing box simulations that indicate that an interaction between convection and the magnetorotational instability in dwarf novae can enhance angular momentum transport, we have studied the simpler case of purely hydrodynamic convection, both analytically and through three-dimensional, fully-compressible simulations in PLUTO.

For the linear phase of the instability, we find agreement between the growth rates of axisymmetric modes calculated theoretically and those measured in the simulations to within a percentage error of $< 1\%$, thus providing a useful check on our PLUTO code. The linear eigenmodes are worth examining not only to help understand the physical nature of convection in disks, but also because they may appear in some form during the nonlinear phase of the evolution, especially on large-scales, and during intermittent or cyclical convection.

We then explored the nonlinear saturation of the instability, both when convection is continually forced and when it is allowed to reshape the background gradients so that it ultimately dies out. We focussed especially on the old problem of whether hydrodynamic convection in a disk leads to inward ($\alpha < 0$) or outward ($\alpha > 0$) angular momentum transport. In both forced and unforced convection we found $\alpha > 0$ in general in the nonlinear phase. These results were confirmed by a separate run using the code ATHENA, but contradict the classical simulations of SB96 who reported inward transport in both cases using the code ZEUS.

This discrepancy reveals a set of unfortunate numerical difficulties that complicate the simulation of convection in disks. These, in large part, issue from the fact that the inviscid linear modes of convection grow fastest on the shortest possible scales. Thus, no matter the resolution, the nature of the code's grid dissipation will always impact on the system's evolution, certainly in the linear phase and possibly afterwards. We argue that a more diffusive numerical set-up, such as supplied by ZEUS at low resolution or Riemann solvers such as HLL, imposes an axisymmetry on the flow which

leads to generally inward transport of angular momentum. But, on the other hand, we suspect that less diffusive solvers such as Roe and HLLC artificially alias shearing waves in the linear phase of the evolution, leading to spurious non-axisymmetric flow early in a simulation. Though concerning, we believe this is only a problem in the low amplitude linear phase because physical mode-mode interactions will dominate once the perturbations achieve larger amplitudes. Nonetheless, shearing wave aliasing certainly deserves a separate study.

To properly dispense with these numerical issues one must add explicit viscosity (and thermal diffusivity), as this regularizes the linear problem. We find that at a Richardson number of $Ri \sim 0.05$, onset of convection is observed for a critical Rayleigh number $10^5 < Ra_c \leq 10^6$. Just above this value convection is largely axisymmetric and $\alpha < 0$. At a larger second critical Ra between 10^6 and 10^7 , the sign of α switches and the flow becomes more turbulent and nonaxisymmetric. This sequence of states mirrors that simulated by Lesur & Ogilvie (2010). At large (resolved) Rayleigh numbers, viscous simulations are initially controlled by the axisymmetric modes; these are then attacked by secondary shear instabilities in both the xz and xy -planes, which break the axisymmetry and order of these structures, leading to a more chaotic state. At lower Ra , viscosity suppresses the non-axisymmetric shear instabilities and axisymmetry is never broken. (At even lower Ra , convection never begins, of course.)

In forced convective runs, rather than maintaining the convection by a fixed heating source at the mid-plane, we instead allowed the vertical equilibrium to relax to its initial, convectively unstable, state. Our thermal relaxation is artificially imposed, but its overall effect is to mimic the heating of the mid-plane and cooling of the corona due to physical mechanisms that maintain the convectively unstable entropy profile, such as the MRI and radiative losses present in the simulations of Hirose et al. (2014). We observed in the non-linear stage now the formation of large-scale convective cells (similar in some respects to elevator flow) that emerge and break down cyclically, in addition to zonal flows and vortices.

Despite our demonstration that hydrodynamic convection can lead to

positive stress and outward transport of angular momentum, the fact remains that the time-averaged stresses are small (typically we measured $\alpha \sim 10^{-6} - 10^{-5}$). Having said that, the magnitude of α is sensitive to the depth of the buoyancy frequency profile, and a deeper profile could increase α by an order of magnitude or more.

Finally we have not observed self-sustaining hydrodynamic convection in any of our unforced simulations. By self-sustaining convection we mean that (when $\alpha > 0$) energy extracted from the shear by convection might itself cause convective motions, which in turn extract energy from the shear, closing the loop. It is more likely that if convection is to occur in disks it will be as a byproduct of other processes, such as heating by density waves, emitted in the presence of a planet, or by dissipation of magnetorotational turbulence. We will investigate the latter mechanism and its instigation of convection in Chapter 4, and we intend to investigate the former mechanism in future work.

Chapter 4

MHD convection in disks

In this chapter we explore the interaction between convection and the MRI in accretion disks. In particular our aim is to isolate the generic features of the interplay between the two instabilities in a controlled manner. We do so through numerical simulations, working in the fully compressible, vertically stratified shearing box approximation. We include both ideal and non-ideal (via an explicit resistivity) magnetohydrodynamics, and employ a perfect gas equation of state and optically thin cooling. We omit much of the complicated physics that has been included in recent work on the subject such as radiative transfer, radiation pressure, and an equation of state that takes into consideration the ionization state of the gas, thus ensuring that our results are as general as possible.

The structure of the chapter is as follows: first, in Section 4.1 we provide a brief overview of the numerical set-up, including initial conditions, boundary conditions, and mass source term (the diagnostics we use throughout this chapter have already been discussed in Section 2.5). We also discuss in more detail the criteria for the onset of convection in a turbulent fluid, and how we expect convection and the MRI to interact. In order to investigate the effect of the MRI alone on the vertical structure of the disk, we first explore the non-linear regime of the MRI (without convection) in vertically stratified simulations without explicit cooling in Section 4.2. Then, in Section 4.3, we investigate the interplay between convection and the MRI in simulations with

a simple cooling prescription. Finally, we discuss our results and conclusions in Sections 4.4 and 4.5, respectively.

4.1 Parameters and set-up

4.1.1 Numerical set-up

Initial conditions and units

All our simulations are initialized from an equilibrium exhibiting a Gaussian density profile

$$\rho = \rho_0 \exp \left[-\frac{z^2}{H_0^2} \right], \quad (4.1)$$

where ρ_0 is mid-plane density, and H_0 is the scale height at the mid-plane at initialization (formally defined below). The equilibrium is isothermal and convectively stable. Note that although the initialization equilibrium is isothermal, the simulations themselves are run with a *perfect gas* equation of state.

The background velocity is given by $\mathbf{u} = -(3/2)\Omega_0 x \mathbf{e}_y$. At initialization we usually perturb all the velocity components with random noise exhibiting a flat power spectrum. The perturbations $\delta\mathbf{u}$ have maximum relative amplitude of about $5 \times 10^{-3} c_{s0}$ and can be either positive or negative.

All simulations are initialized with a *zero-net-flux* (ZNF) magnetic field configuration (more specifically, at initialization we employ a sinusoid in the vertical component of the magnetic field B_z with the other two components being zero).¹ The field strength at initialization is controlled through the ratio of gas pressure to magnetic pressure at the mid-plane $\beta_0 \equiv P/(B^2/2)$.

¹The primary motivation for choosing a ZNF configuration as opposed to a net vertical flux (NVF) configuration was to compare our results more easily with the ZNF simulations of Bodo et al. (2012) and Hirose et al. (2014). In actual dwarf novae it is possible that the magnetic field of the primary threads the disk (though the actual field configuration is unknown), in which case a NVF configuration is more appropriate. Such NVF simulations have been carried out by Scepi et al. (2018a).

(Note that here and in the remainder of the chapter we have chosen units such that the factor μ_0 disappears.) In our simulations we set $\beta_0 \equiv 1000$. Note however, that the onset of MRI turbulence should erase the initial conditions after some period of time. Thus we could also have employed a lower or higher plasma beta at initialization and obtained the same result in the non-linear phase.

As in our hydrodynamic simulations, units are selected so that $\Omega = 1$. The length unit is chosen so that the initial mid-plane isothermal sound speed $c_{s0} = 1$, which in turn defines a reference scale height $H_0 \equiv c_{s0}/\Omega = 1$. Note, however, that the sound speed (and the scale height) is generally a function of both space and time.

Box size and resolution

We measure box size in units of initial mid-plane scale height H_0 , defined above. For our fiducial simulations we employ a resolution of $128 \times 128 \times 196$ in boxes of size $4H_0 \times 4H_0 \times 6H_0$, which corresponds to a resolution of 32 grid cells per H_0 in all directions.

In order to test the convergence of some of our results with resolution we employ resolutions of $64 \times 64 \times 96$ and $256 \times 256 \times 392$, corresponding to 16 and 64 grid-cells per scale height in each direction, respectively. To investigate the effect on our results of narrower boxes, we employ a radial extent $L_x = H_0$ and $L_x = 2H_0$, respectively (keeping the number of grid-cells per scale height fixed at $32/H_0$).

Boundary conditions

We use shear-periodic boundary conditions in the x -direction (see Hawley et al. (1995a)) and periodic boundary conditions in the y -direction. In the vertical direction, we keep the ghost zones associated with the thermal variables in isothermal hydrostatic equilibrium, in the manner described in Zingale et al. (2002) (in addition, the temperature of the vertical ghost zones is updated at each time-step to match the temperature in the active cells bordering the ghost cells). For the velocity components we use mostly standard

outflow boundary conditions in the vertical direction, whereby the vertical gradients of all velocity components are zero (i.e. $\partial u_x/\partial z = 0, \partial u_y/\partial z = 0, \partial u_z/\partial z = 0$) and variables in the ghost zones are set equal to those in the active cells bordering the ghost zones. For the magnetic field we employ ‘vertical field’ boundary conditions, i.e. $B_x = 0, B_y = 0, \partial B_z/\partial z = 0$ (see, for example, Riols & Latter (2018)).

Mass source term

As our simulations employ open boundary conditions in the vertical direction (see Section 4.1.1), mass can escape from the domain leading to a gradual depletion of mass in the box. To mitigate this we incorporate a simple mass source term. At the end of the n th step, we subtract the total mass in the box at the end of that step M_n from the total mass in the box at initialization M_0 . This mass difference $\Delta M_n \equiv M_0 - M_n$ is added back into the box with the same profile used to initialize the density (cf. Equation 4.1).

For our fiducial MRI only simulation (see Section 4.2.2, we compared ΔM_n to the actual mass flux through the vertical boundaries at the end of each step. The results generally agree to within about 1%. The cumulative mass lost (and added back into box) over 200 orbits is about $6M_0$.

4.1.2 Criteria for convective instability in a viscid fluid

As we have discussed at the end of Section 1.3.3, in a viscid fluid convection is mitigated by the effects of viscosity ν and thermal diffusivity χ , with the ratio of destabilizing and stabilizing processes quantified by the Rayleigh number. Thus in a viscid fluid, the criterion that the square of the buoyancy frequency is negative, i.e. $N_B^2 < 0$, is a necessary condition but not a sufficient one. When explicit viscosity and thermal diffusivity are included, convective instability requires both that $N_B^2 < 0$, *and* that the Rayleigh number exceed some critical value Ra_c .

There remains a subtlety in the interpretation of the Rayleigh number, however, in particular as to the origin of the diffusion coefficients for viscosity ν and thermal diffusivity χ . The onset of convection on top of a

laminar background is determined by the *microscopic* Rayleigh number (i.e. the viscosity ν and thermal diffusivity χ are microscopic in nature). The existence of a *turbulent* background, however (such as might be supplied by MRI turbulence) may result in an *effective* (or turbulent) viscosity and thermal diffusivity due to turbulent transport of momentum and of heat, respectively. The idea that disk turbulence, and the MRI in particular, can transport momentum has been long established, but the ability of turbulence to transfer heat has only been discussed sporadically (e.g. Rudiger (1987); Heinrich (1994); D'Alessio et al. (1998); Gu et al. (2000); Blaes et al. (2011b); the last suggesting flows associated with rising flux tubes can transport heat). Thus for a given entropy gradient, or, equivalently, a fixed N_B^2 , the enhanced effective viscosity and thermal diffusivity due to MRI turbulence act to lower the effective Rayleigh number, which might then impede the onset of convection. Thus in our investigations of the interplay between convection and the magnetorotational instability the fluid should be characterized by a *turbulent* Rayleigh number Ra_{eff} , which must be sufficiently large so that convection resists the disordered background flow. In any case, the sign of N_B^2 (as defined in Equation 1.21) is certainly insufficient to assign convection to MRI-turbulent flows, as is often done in recent work.

We caution, however, that effective diffusion coefficients are just a model for the actual turbulent transport, and not necessarily always a good one: in particular turbulent fluxes may not always behave in a diffusive manner, especially on short length or timescales. A complementary way of thinking about how the MRI might inhibit convection, then, is to consider the timescales involved: the turn-over time of the most vigorous MRI eddies will be of the order $\tau_{\text{MRI}} \sim 1/\Omega$, i.e. on a dynamical timescale, whereas the fastest growing convective modes grow on timescales of the order $\tau_{\text{conv}} \sim 1/|N_B|$ (see Figure 3.6). Unless $\tau_{\text{conv}} \sim \tau_{\text{MRI}}$, the convective modes will get disrupted by the MRI turbulence before they can grow and form coherent plumes. In addition to sharing similar timescales, the MRI and convective eddies also share similar length scales. Thus there is no separation of scales between the two phenomena, suggesting that they must interact. Note that some of these ideas regarding turbulent heat transport and turbulent damping of convec-

tion were foreshadowed in Gu et al. (2000), a somewhat neglected work.

It must also be said that the existence of magnetic fields will also impact on the onset and development of convection (e.g. Weiss & Proctor (2014) and references therein). The magnetic tension associated with a large-scale magnetic field will impede nascent convective motions in a relatively straightforward way, but the small-scale magnetic fluctuations associated with the zero-net flux MRI may have an effect that can be packaged away in the effective viscosity discussed above.

Finally we wish to discuss a very important point about how the two instabilities interact. The reader might be tempted at this stage to think of convection occurring ‘on top of’, or simultaneously with, MRI turbulence if the effective Rayleigh number is sufficiently large. Indeed several authors who have investigated the interplay between the MRI and convection in disks imply that the two instabilities simply interact in an ‘additive’ manner, in other words that convection and the MRI are operating together at the same time (Bodo et al., 2012; Hirose et al., 2014). As we find in controlled numerical experiments set up to sustain both convection and the MRI, however, the two instabilities *do not* appear to operate simultaneously, though they can interact in non-trivial ways (see Section 4.3.2). Thus it is not the case that once the effective Rayleigh number exceeds some critical value convection will start on top of MRI turbulence: rather the magnitude of the effective Rayleigh number should be interpreted as a condition on whether the flow will be dominated by convection *or* by the MRI.

In addition to the physical issues discussed above, there is, of course, the practical problem of identifying the two instabilities diagnostically in the first place, which is no easy task. As we have already mentioned $N_B^2 > 0$ can rule out convection, but $N_B^2 < 0$ is only a necessary, not a sufficient, criterion for establishing whether there is convection. Another common diagnostic for identifying convection that has been used in the literature is the vertical heat flux, or some measure thereof. However, this too is problematic, because (as we find in Section 4.2), the MRI itself is quite efficient at transporting heat vertically. Physically this can be understood by thinking in terms of an effective thermal diffusivity due to MRI, though as we pointed out earlier,

it is not exactly clear whether heat transport by MRI turbulence acts in a diffusive manner.

4.2 Stratified MRI simulations without explicit cooling

4.2.1 Motivation

In this section we describe the results of our fiducial vertically stratified ZNF MRI simulation. We employ open vertical boundaries (see Section 4.1.1) but do not include explicit cooling: thus cooling is facilitated by advection of thermal energy across the vertical boundaries only. We refer to simulations like this as *box-cooled* simulations. In order to isolate the behavior of the MRI in vertically stratified boxes with a perfect gas equation of state, we omit explicit cooling and diffusion coefficients. Key questions that we wish to address in this section include: ‘What is magnitude and behavior of α and how does it compare to previous work?’; ‘What vertical equilibrium structure does the disk settle into in the MRI turbulent state?’, and ‘Is this vertical structure convectively unstable?’. The aim is for our fiducial simulation to serve as a benchmark against which we can compare both simulations that include spatially uniform explicit cooling, as well as simulations with piecewise (i.e. height-dependent) cooling (see Section 4.3). For the sake of brevity, we restrict our discussion in this section to a single simulation. A more comprehensive investigation of the effects of resolution and radial box size on the ZNF vertically stratified MRI is deferred to Appendix A.2.

In addition to acting as a bench-mark against which to compare later results, the simulations presented in this section (and in the appendices) are intended to fill a gap in the literature. Most 3D MHD vertically stratified shearing box simulations of the MRI have included thermodynamics in either a very simple or in a very complicated manner. On the simple end of the spectrum, there is a large body of work that is restricted to an *isothermal* equation of state (e.g. Davis et al. (2010); Bai & Stone (2013); Bodo

et al. (2014); Salvesen et al. (2016); Ryan et al. (2017)), thus precluding the existence of convection. On the more complicated end of the spectrum, several authors employed a perfect gas equation of state together with *radiative transfer* and thus radiation pressure (through flux-limited diffusion, e.g. Turner et al. (2002)), while others have employed both radiative transfer and an equation of state that takes into consideration the ionization state of the gas (e.g. Hirose et al. (2006); Blaes et al. (2011a); Hirose et al. (2014)). Flaig et al. (2010) investigated protoplanetary disks using flux limited diffusion radiative transfer and opacities suitable for temperatures between 1000 K and 2000 K. In a follow-up paper they added explicit resistivity as well (Flaig et al., 2012). Perhaps somewhat surprisingly, only a handful of vertically stratified ZNF simulations have been carried out with a *perfect gas equation of state* (Stone et al., 1996; Bodo et al., 2012; Gressel, 2013; Bodo et al., 2013a, 2015; Riols & Latter, 2018). Of these Stone et al. (1996) and Bodo et al. (2012, 2013a, 2015) employed closed boundary conditions in the vertical direction, Gressel (2013) used very narrow boxes (with $L_x = H_0$), and Riols & Latter (2018) included *self-gravity* in most of their simulations.

In Section 4.2.2, we describe our fiducial box-cooled simulation. In Appendix A.1 we discuss some robustness issues related to the non-linear MRI in stratified shearing boxes. In particular we find that a diffusive implementation of the $\nabla \cdot \mathbf{B} = 0$ condition using Constrained Transport can weaken or even switch off the non-linear MRI. Finally, brief convergence studies (both in resolution, and in radial box size) are presented in Appendix A.2.

4.2.2 Fiducial simulation

Set-up and initialization

The simulation described in this section was run at a resolution of $128 \times 128 \times 196$ in a box of size $4H_0 \times 4H_0 \times 6H_0$, where H_0 is the scale height at the mid-plane at initialization. Thus the resolution is about $32/H_0$ in all directions. The simulation was initialized from the isothermal, vertically stratified equilibrium described in Section 4.1.1, with zero-net magnetic flux, and an initial mid-plane plasma beta of $\beta_0 = 1000$. The simulation was run

for 200 orbits ($1257 \Omega^{-1}$).

A crucial element of this simulation concerns the vertical boundaries. Bodo et al. (2012) employed closed vertical boundaries (and no explicit cooling) in their vertically stratified MRI simulations, which resulted in the box heating up, until the disk eventually filled the entire box. This resulted in a distinct tent-like vertical profile for temperature and a flat vertical profile for the density throughout their box (see Figures 4 and 5, respectively, in their paper). More seriously, Gressel (2013) demonstrated that the enhancement in α that Bodo et al. (2012) observed was greatly reduced when open boundary conditions were employed. Thus to mitigate uncontrolled heating of our box, we employ open boundary conditions (as described in Section 4.1.1), enabling advection of mass and energy (thermal, kinetic, and magnetic) across the vertical boundaries. One disadvantage of this approach is that the thermal equilibrium will be influenced by the size of the box and the vertical boundary conditions. To check the extent to which this affects the results we have rerun this simulation with uniform (i.e. height-independent) cooling (see Section 4.3.1), and found that the results (i.e. the time-evolution of volume-averaged quantities, the flow field, and convective stability) are the same.

Time evolution of averaged quantities

In Figure 4.1 we show the time-evolution of the z -component of the volume-averaged vertical kinetic energy density (top) and the total magnetic energy density (bottom). We look at the vertical component of the kinetic energy rather than the total kinetic energy because we are primarily interested in how the MRI moves fluid vertically and rearranges the vertical structure of the disk (which we can later compare to how convection affects the vertical disk structure). After initialization, there is exponential growth (corresponding to the linear phase of the MRI) over the first few orbits, followed by non-linear saturation when $\langle E_{\text{kin},z} \rangle \sim 10^{-2}$. The system then settles into a quasi-steady state with fluctuations around an average vertical kinetic energy of about 10^{-3} . The time-evolution of the magnetic energy density $\langle E_{\text{mag}} \rangle$ is




Figure 4.1: Semi-log plot of the time-evolution of volume-averaged vertical kinetic energy density (red curve) and volume-averaged total magnetic energy (black curve) from the fiducial inviscid, box-cooled, vertically stratified MRI simulation without explicit cooling.

very closely correlated with the vertical kinetic energy. It peaks at around $\langle E_{\text{mag}} \rangle \sim 5 \times 10^{-2}$ at non-linear saturation, and subsequently fluctuates around an average value of around 1.5×10^{-2} .

The thermal energy density exhibits similar behavior to the kinetic and magnetic energies. After initialization $\langle E_{\text{th}} \rangle$ increases rapidly as the MRI dissipates heat, eventually settling into a quasi-steady state in which it fluctuates around a mean value of $\langle E_{\text{th}} \rangle \sim 0.82$. The thermal energy balance is determined by turbulent dissipation of heat due to the MRI and cooling due to advection of thermal energy across the vertical boundaries. The horizontally-averaged mid-plane scale height $\langle H \rangle_{xy}(z = 0, t) \equiv \sqrt{\gamma \langle T(z = 0, t) \rangle_{xy}} / \Omega$ is a good preliminary diagnostic of the variation in ‘disk thickness’ over the course of the simulation. After initialization $\langle H \rangle_{xy}$ increases rapidly from




Figure 4.2: Plot of time-evolution of volume-averaged Reynolds stress (blue), magnetic stress (red), and α (black) from the fiducial inviscid box-cooled vertically stratified MRI simulation without explicit cooling. The volume-averaged thermal pressure (divided by 30) is also plotted (green) in order to track the phase relationship between stress and pressure.

$\langle H \rangle_{xy} = H_0$ reaching $\langle H \rangle_{xy} \sim 1.5H_0$ at non-linear saturation, before leveling out and fluctuating just under $1.5H_0$ for the remainder of the simulation, thus reaching half the vertical box semi-size of $L_z/2 = 3H_0$. We interpret this behavior as evidence that the disk rapidly heats up and expands due to the dissipation of heat by MRI turbulence, with the expansion being halted once turbulent heating is balanced by cooling due to advection of disk fluid across the vertical boundaries. For a more quantitative estimate of the effect of box cooling we also calculated the wind cooling time (cf. Equation 2.73) and find this to be $\tau_w \sim 112 \Omega^{-1}$ compared to the value of $\tau_w \sim 150 \Omega^{-1}$ measured in the box-cooled simulation MRI-S1 of Riols & Latter (2018) (though their simulation was run in a box half the radial size of ours).

In Figure 4.2 we plot the time-evolution of the volume-averaged Reynolds stress $\langle R_{xy} \rangle$, magnetic stress $\langle M_{xy} \rangle$, and of $\langle \alpha \rangle$. The time-averaged values

(from orbit 40 to orbit 200) are: $\langle\langle\alpha\rangle\rangle_t \sim 0.0141$, $\langle\langle M_{xy}\rangle\rangle_t \sim 6.2 \times 10^{-3}$, and $\langle\langle R_{xy}\rangle\rangle_t \sim 1.4 \times 10^{-3}$. The ratio of magnetic stress to Reynolds stress is ~ 4.5 , demonstrating that most of the stress in the non-linear phase is due to the magnetic field. The stresses measured in our simulations are somewhat larger than those measured in the ‘small box’ simulation (MRI-S1) of Riols & Latter (2018) (RL2018): they reported $\langle\langle\alpha\rangle\rangle_t \sim 0.0076$, $\langle\langle M_{xy}\rangle\rangle_t \sim 2.5 \times 10^{-3}$, and $\langle\langle R_{xy}\rangle\rangle_t \sim 7.5 \times 10^{-4}$, with the time-average in their simulation being taken over 238 orbits. We are not certain as to the exact nature of the discrepancy between the stresses measured in our simulation and those of RL2018. The simulation of RL2018 employed a box of half the radial size compared to ours, but as we discuss in more detail in Appendix A.2 we expect the stresses to be higher in that case, which we’ve confirmed through a $L_x = 2H_0$ simulation of our own. Because RL2018 used the same code as we did, it is possible that the lower stresses they report are due to numerical issues with the implementation of $\nabla \cdot \mathbf{B} = 0$, which we discuss in greater detail in Appendix A.1.

Vertical structure of the disk

In Figure 4.3 we plot the horizontal- and time-averaged vertical profiles of various quantities. The top panel shows the vertical heat flux (in red) and mass flux (in black). Both the mass and heat flux are directed away from the mid-plane. The heat flux peaks at a value of around $F_{\text{heat}} \sim \pm 5 \times 10^{-3}$ at around $\pm 2H_0$ on either side of the mid-plane. The vertical temperature profile (not shown) is fairly constant (i.e. isothermal) within around $1.5H_0$ of the mid-plane, but drops rapidly outside of this region. The buoyancy frequency $\langle\langle N_B^2 \rangle\rangle_{xy,t}$ (bottom panel of Figure 4.3) is positive everywhere, except exactly at the mid-plane where the disk is marginally stable. The positivity of N_B^2 precludes convection, and indeed the flow field (i.e. u_z in the xz -plane) appears to be dominated by MRI turbulence: we do not observe any noticeable vertical structure or plumes characteristic of convection. Thus the combination of MRI heating (which occurs predominantly at the mid-plane) and the weak implicit cooling facilitated through advection of fluid





Figure 4.3: Horizontal-and-time-averaged vertical disk profiles from the fiducial inviscid box-cooled vertically stratified MRI simulation without explicit cooling. Top: vertical mass (in black) and heat (in red) fluxes. Bottom: (square) of the buoyancy frequency. The time-averages are taken from orbit 50 to orbit 196.

across the boundaries is insufficient to set up a convectively unstable entropy gradient.

The vertical stress profiles for the time- and horizontally-averaged Reynolds stress $\langle\langle R_{xy}\rangle_{xy}\rangle_t$ and magnetic stress $\langle\langle M_{xy}\rangle_{xy}\rangle_t$ (not shown) are relatively constant within $\pm 1.5H_0$ of the mid-plane. Both profiles drop off rapidly beyond around $1.5H_0$ and are only about one tenth of their peak values at the boundaries. Note that in our preliminary analysis, in which we time-averaged our results over just 60 orbits, we observed that the stresses displayed a slightly double-peaked profile across the mid-plane, with the magnetic stress reaching peak values of just under 0.01 at around $\pm 1.2H_0$, compared to values just under 0.008 around the mid-plane, and the Reynolds stress peaking at 0.02 around $\pm 1.2H_0$ compared to around 0.0016 near $z = 0$. This double peaked nature vanished, however, when we took the time-averages over longer intervals (orbit 50 to orbit 200).

Given that the critical wavelength of the linear MRI is a function of the Alfvén speed $v_A = B/\sqrt{\rho}$ (which depends on density and therefore on height in vertically stratified simulations), we have also checked whether we are resolving the MRI at the mid-plane. The critical, or fastest growing, linear MRI modes are related to the Alfvén speed by $\lambda_c = 2\pi\sqrt{16/15}v_A/\Omega$, which can be more conveniently written in terms of the plasma beta as $\lambda_c \sim 9.18/\sqrt{\beta}$ in units of the initial mid-plane scale height (Hawley et al., 1995a). In the literature a diagnostic that is commonly used to test whether this wavelength is resolved is to calculate the Q-parameter, or ratio of critical MRI wavelength to the grid size, i.e. $Q \equiv \lambda_c/\Delta z$ in the vertical direction (Riols & Latter, 2018). The plasma beta decreases with height and thus the critical MRI wavelength is generally smallest at the mid-plane in vertically stratified simulations. In our run, however, we find that the horizontally-averaged mid-plane plasma beta during the non-linear phase (i.e. averaged from orbit 50 to orbit 196) is $\langle\langle\beta(z=0)\rangle_{xy}\rangle_t \sim 58$, dropping monotonically to a value of around $\langle\langle\beta(|z|=L_z/2)\rangle_{xy}\rangle_t \sim 10$ near the vertical boundaries. This corresponds to a Q-parameter of around $Q \sim 39$, and so the MRI should be well resolved at the mid-plane in our simulations run at a resolution of $32/H_0$. Thus the double peaked nature of the profiles is more likely to be a

physical effect rather than a numerical one, and, as we found when we took longer time-averages, a transient one at that.

Finally, the vertical profile of α shows $\langle\langle\alpha\rangle_{xy}\rangle_t$ peaking strongly near the vertical *boundaries* with $\langle\langle\alpha\rangle_{xy}\rangle_t \sim 0.06$ at the boundaries compared to $\langle\langle\alpha\rangle_{xy}\rangle_t \sim 0.01$ at the mid-plane. This is, however, entirely an artifact of the vertical stratification and not an indication of enhanced MRI turbulence near the vertical boundaries. The pressure in our vertically stratified simulation decreases for $z > 0$ and, provided that the pressure decreases faster than the stresses do, this causes $\langle\langle\alpha\rangle_{xy}\rangle_t$ to appear to peak at the boundaries. Thus α makes more sense as a *volume*-averaged quantity rather than a *horizontally*-averaged quantity.

4.3 Stratified MRI simulations with cooling

4.3.1 Inviscid simulation with uniform cooling

The inviscid simulations discussed in Section 4.2 (and in Appendix A.2) are *box-cooled*, in the sense that cooling is facilitated by advection of thermal energy across the vertical boundaries only. Thus the balance of heating and cooling is set by the vertical extent of the box. It is not clear a priori whether these simulations can be compared in any meaningful way to the simulations in this section which employ an explicit cooling prescription. To test whether the box cooling significantly alters the results compared to explicit cooling, we have run a simulation (NSTRMC47) with identical resolution ($128 \times 128 \times 196$ or $32/H_0$), box size size ($4H_0 \times 4H_0 \times 6H_0$), and initial conditions to our fiducial box-cooled simulation, but *including* explicit cooling via a linear cooling prescription (see Section 2.1). The cooling timescale is $\tau_c = 10$ orbits. In contrast to the simulations discussed in Section 4.3.2, cooling is employed *uniformly* everywhere in the box in this simulation, and no explicit diffusion coefficients are used. The simulation was run for 100 orbits ($629 \Omega^{-1}$).

In our uniformly cooled simulation, it is not clear from the outset whether the explicit cooling term or advection of energy across the vertical boundaries




Figure 4.4: Top: semi-log plot of the time-evolution of volume-averaged vertical kinetic energy density for inviscid simulations with (a) an explicit cooling prescription implemented uniformly in the z -direction with a cooling timescale of $\tau_c = 10$ orbits (blue curve), and (b) no explicit cooling (red curve). Bottom: time-evolution of the volume-averaged total stress $T_{xy} \equiv R_{xy} + M_{xy}$ normalized by initial mid-plane pressure $P_0 = 1$ in the two simulations. The blue vertical dashed line denotes the time at which cooling was switched on in the simulation with explicit cooling.

dominates the cooling. The primary diagnostic we use to estimate the effects of cooling due to advection of energy across the vertical boundaries on the energy budget is the wind cooling timescale τ_w (see Equation 2.73 in Chapter 1). In our box-cooled simulation we measured this to be $\tau_w = 112.4 \Omega^{-1}$, compared to the value of $\tau_w = 156.3 \Omega^{-1}$ that we measure in our uniformly cooled simulation. Thus the wind cooling timescale is significantly longer than the explicit cooling timescale of $\tau_c = 62.8 \Omega^{-1}$, which reassures us that the explicit cooling rather than the box-cooling dominates the cooling budget in our uniformly cooled simulation. However, a further difficulty is that for thermal equilibrium to be maintained in the non-linear phase, the explicit cooling needs to be fine-tuned so that the total cooling balances the viscous heating, while simultaneously ensuring that the explicit cooling exceeds the wind cooling. As we have already discussed, if the heating exceeds the total cooling the disk will heat up and expand until box cooling establishes thermal equilibrium.² But the opposite can also occur: if the cooling is too strong the disk will cool and collapse, thus extinguishing the MRI (because the critical wavelength of the MRI will exceed the disk scale height). Because we wish to compare this simulation to those discussed in the next section we have opted to use the same cooling timescale employed in those simulations ($\tau_c = 10$ orbits), rather than adjust the cooling to maintain thermal equilibrium. This results in a slight imbalance between cooling and heating in this simulation, however, and we find that the disk is contracting to the mid-plane as it cools, though the contraction is slow enough that the MRI does not appear to be affected.

In Figure 4.4 we compare the time-evolution of the volume-averaged vertical kinetic energy density (upper panel) and total stress (lower panel) between the uniformly cooled simulation (blue curve), and the box-cooled simulation (red curve). The thermal energy density (not shown) in the explicitly cooled

²The relation between α and τ_c that should hold in thermal equilibrium is explicitly given by $\alpha(\gamma - 1)S = 1/\tau_c + 1/\tau_w$, where the right-hand-side corresponds to contributions due to viscous heating and the left-hand-side corresponds to contributions due to explicit cooling and cooling due to advection of thermal energy across the vertical boundaries. See (Riols & Latter, 2018).

simulation is, as expected, lower than in the box-cooled simulation, with $\langle E_{\text{th}} \rangle \sim 0.5$ in the explicitly cooled simulation compared to $\langle E_{\text{th}} \rangle \sim 0.8$ in the box-cooled run. Otherwise the results of the box-cooled and explicitly cooled simulations look very similar. In particular we find that $\langle \langle N_B^2 \rangle_{xy} \rangle_t > 0$ everywhere in our explicitly cooled run (the time-average was taken between orbit 40 and orbit 100), and thus the disk is convectively stable everywhere in the explicitly cooled simulation (as it is in the box-cooled simulation). Thus the box-cooled simulations, despite their obvious drawback that the thermal equilibrium is controlled by the size of the box, nevertheless appear to exhibit behavior that is qualitatively and even quantitatively similar to a run in which the equilibrium is controlled by the cooling timescale.

Finally, for completeness we have also carried out an inviscid simulation with *height*-dependent cooling (i.e. the cooling is activated only above $z = \pm 0.75H_0$) and a cooling timescale of $\tau_c = 10$ orbits. The buoyancy frequency is negative within around $\pm H_0$ in this simulation, but we cannot detect any other signs of convection. Otherwise the results are very similar to those just discussed for the uniformly cooled simulation.

4.3.2 Resistive simulations with cooling above $|z| > 0.75H_0$

Motivation

In this section we explore the interplay between the MRI and *convection*. In our fiducial simulation (see Section 4.2.2) we investigated the MRI in stratified shearing boxes without any explicit cooling, and found that the disk was convectively stable. Since neither MRI heating together with box cooling, nor MRI heating together with explicit cooling that is uniform space, can set up an unstable entropy gradient, in this section we include height-dependent explicit cooling using a linear cooling prescription (see Section 2.1). Thus we aim to imitate the vertical structure of dwarf novae disks in the high-state, in which an optically thick disk (in which accretion is driven by MRI turbulence, and, possibly, enhanced by convection) is surrounded by

an optically thin corona.

The MRI is a local instability. As discussed in Section 4.1.2, the onset of convection on top of the background MRI turbulence likely depends not just on the sign of the entropy gradient (or, equivalently, on the sign of N_B^2), but also on the effective Rayleigh number $\text{Ra}_{\text{eff}} = |N_B|^2 H^4 / (\nu_{\text{eff}} \chi_{\text{eff}})$. This provides us with an avenue for controlling whether we can get convection or not in our simulations: at a sufficiently large Ra_{eff} we expect the onset of convection to occur. Increasing Ra_{eff} beyond that point should result in an increased vertical heat flux relative to the MRI, and thus provides a way of controlling the ‘strength’ or intensity of convection relative to the MRI. One approach is to increase Ra_{eff} by *decreasing* the effective viscosity ν_{eff} . In order to reduce the extent to which the MRI impedes the onset of convection, we mitigate the MRI by employing a small but finite *explicit resistivity* η . This decreases the effective viscosity thus increasing the effective Rayleigh number.

A second approach is to increase the effective Rayleigh number by increasing the magnitude of N_B^2 . We achieve this by decreasing the *cooling timescale* τ_c . This increases the temperature difference between the corona and the mid-plane which in turn increases the unstable entropy gradient in that region. Thus by specifying both the cooling timescale τ_c and an explicit resistivity η in our simulations, we can effectively control the strength of convection and of the MRI, respectively, allowing us to explore their interaction in different regimes of parameter space.

A key result is that we cannot find evidence that convection and the MRI occur additively and continuously in time, as has been claimed in recent works (Bodo et al., 2012; Hirose et al., 2014). At best the two instabilities seem to feed off of one another, existing as alternating and relatively well-separated MRI-dominated and convective-dominated cycles (see Section 4.3.3) in which convection appears to seed the MRI while dissipation of MRI turbulence provides residual heating that subsequently triggers convection (see Section 4.3.3 for a more detailed discussion). Otherwise we observe straight MRI (see Section 4.3.4), exhibiting only the occasional convective burst, though these outbursts do occur more frequently as the minimum value of the buoyancy

frequency is decreased (i.e. as the cooling timescale is lowered). We discuss two runs that clearly exhibit one or the other behavior in Sections 4.3.3 and 4.3.4, respectively. More marginal cases are discussed in Section 4.3.5 where we conduct a parameter survey (in the space of (τ_c, η)) to determine the separation between the different regimes.

Set-up

In the simulations described in Sections 4.3.3-4.3.5, we employ an explicit resistivity η (which is kept constant in space and time in any individual simulation). In our inviscid simulation in which cooling was implemented *uniformly* in the domain (see Section 4.3.1) we found that the disk was convectively stable everywhere. Here we implement piecewise cooling in space, i.e. cooling is turned on only within a region $|z| > 0.75H_0$ above the mid-plane. This allows for an unstable entropy gradient to develop in the vicinity of the mid-plane.

We initialize all simulations from around orbit 5 of the non-linear MRI turbulent state of our fiducial simulation (see Section 4.2). This is sufficiently close to non-linear saturation of the linear MRI that the disk has not had time to heat up sufficiently to fill the box. At initialization the mid-plane scale height is $H \sim 1.55H_0$. Each simulation was run for 100-200 orbits ($628 - 1257 \Omega^{-1}$).

4.3.3 MRI/convective cycles

Time-evolution of volume-averaged quantities

For a cooling timescale of $\tau_c = 10$ orbits and resistivities $\eta \geq 5 \times 10^{-4}$ we observe a state characterized by cyclical outbursts in which the flow appears to switch between convection and the MRI. As our fiducial example of a run exhibiting this behavior, we have chosen simulation NSTRMC44e1b. This simulation was run with a cooling timescale of $\tau_c = 10$ orbits and an explicit resistivity of $\eta = 5 \times 10^{-3}$. Although the resistivity is relatively large, this simulation exhibits the clearest behavior of MRI/convective cycles. In




Figure 4.5: Top: semi-log plot of the time-evolution of volume-averaged vertical kinetic energy density (in black) and ratio of vertical to radial kinetic energy (in red) from a simulation exhibiting MRI/convective cycles (NSTRMC44e1). Bottom: time-evolution of volume-averaged magnetic energy density (black). The thermal pressure (divided by 10) is superimposed in green. The simulation employed a constant cooling timescale of $\tau_{\text{relax}} = 10$ orbits above $|z| = 0.75H_0$, and a uniform explicit resistivity of $\eta = 5 \times 10^{-3}$.




Figure 4.6: Time-evolution of volume-averaged α parameter (black), magnetic stress (red), and Reynolds stress (blue) from a simulation exhibiting convective/MRI cycles (NSTRMC44e1). The simulation employed a constant cooling timescale of $\tau_{\text{relax}} = 10$ orbits above $|z| = 0.75H_0$, and a uniform explicit resistivity of $\eta = 5 \times 10^{-3}$.

Section 4.3.5 we carry out a parameter survey in which we find MRI/cycles for smaller resistivities and at lower cooling timescales, though these cycles can have different properties to the ones discussed here.

In Figure 4.5 we plot the time evolution of the vertical kinetic energy density (top) and magnetic energy density (bottom). Compared to the behavior of the kinetic energy in our box-cooled simulation (see Figure 4.1) in which $\langle E_{\text{kin},z} \rangle$ fluctuated around 10^{-3} in the non-linear phase, in this simulation we instead observe bursts in the kinetic energy density during which $\langle E_{\text{kin},z} \rangle$ can reach as high as 10^{-2} . Between the outbursts, the kinetic energy density is found to be highly oscillatory, reminiscent of the behavior of the kinetic energy observed in simulations of forced compressible (hydro) convection in Section 3.4. To more clearly track the changes in the vertical kinetic energy, we have also plotted the time evolution of the ratio

of vertical to radial kinetic energy. This is anti-correlated with magnetic and vertical kinetic energies. The vertical kinetic exceeds the radial kinetic energy during convection dominated periods (see below), while during MRI-dominated $\langle E_{\text{kin},z} \rangle \sim 0.7 \langle E_{\text{kin},x} \rangle$. Note that this is large than the ratio of $\langle E_{\text{kin},z} \rangle \sim 0.3 \langle E_{\text{kin},x} \rangle$ reported in the vertically stratified ZNF shearing box simulation IZ1 of Stone et al. (1996) (see Table 2 of their paper), although their simulation was isothermal (and thus stably stratified), and therefore likely impeded vertical motions. An important point here, however, is that the MRI itself naturally gives rise to stronger vertical flows than convection (as evidenced by the peaks in the vertical kinetic energy during the MRI-dominated bursts in α).

In Figure 4.6 we plot the time evolution of the stresses and of α . The initial spike in α around orbit 2 is due to the linear MRI, followed by non-linear saturation. The vertical dashed line indicates the time at which resistivity was turned on. The stress abruptly drops as resistivity appears to quench the non-linear MRI. However this quenching is followed by alternating periods of low stress (‘quiescence’) and high stress (‘outbursts’). During outburst α can reach as high as 0.08 (we have observed values up to 0.1 in some of our preliminary simulations), though typical values of α in outburst are around $\langle \alpha \rangle \sim 0.04 - 0.06$.

Structure of the flow

The periods of quiescence and outburst observed in the time-evolution of volume-averaged quantities (see Figure 4.6) correspond to *convection*-dominated and *MRI*-dominated phases, respectively. This can be seen most clearly in Figure 4.7 where we plot the z -component of the velocity in the xz -plane. In the left-most panel, taken from a snapshot during the first *quiescent* phase (around orbit 15), the onset of convective instability is clearly visible within about $\pm H_0$ of the mid-plane through thin convective cells consisting of hot rising fluid (in red) and cool sinking fluid (in blue). These thin convective cells are reminiscent of linear convective instability observed in our hydro simulations (see Figure 3.11). Although at this point in the simulation the




Figure 4.7: Snapshots of the vertical component of the velocity in the xz -plane taken at different times from a simulation exhibiting MRI/convective cycles (NSTRMC44e1). (a) convection-dominated phase, (b) MRI-dominated phase, (c) convection-dominated phase with large-scale convective cells. The simulation employed a constant cooling timescale of $\tau_c = 10$ orbits above $|z| = 0.75H_0$, and a uniform explicit resistivity of $\eta = 5 \times 10^{-3}$.

MRI has been quenched by the explicit resistivity, the combination of cooling from above and residual heating from before the MRI was quenched has triggered what appears to be the linear phase of the convective instability. Conversely, during *outbursts* (middle-panel of Figure 4.7 taken at orbit 35) the flow is dominated by MRI turbulence. No convective cells are visible in the flow field during outbursts, though the buoyancy frequency is still negative during these phases. Finally, during subsequent quiescent (i.e. convection-dominated) phases we see the emergence of large scale convective cells (right-most panel of Figure 4.7, taken from a snapshot at orbit 49). These are reminiscent of the large-scale cyclical convective cells reported in hydrodynamic simulations of non-linear forced compressible convection (see Figure 3.17). The cells are cyclically created and reformed with the opposite orientation, and their manifestation in the flow field coincides with short term oscillations (with periods of order 1 orbit) in the vertical kinetic energy and




Figure 4.8: 2D y -averaged power spectra of specific vertical kinetic energy from a simulation exhibiting MRI/convective cycles (NSTRMC44e1). The colorbars are logarithmic. Left: spectrum taken from snapshots time-averaged over a *convection*-dominated phase. Right: same but time-averaged over an MRI-dominated phase. The simulation employed a constant cooling timescale of $\tau_{\text{relax}} = 10$ orbits above $|z| = 0.75H_0$, and a uniform explicit resistivity of $\eta = 5 \times 10^{-3}$. For each spectrum the data has been taken within $\pm H_0$ of the mid-plane and averaged in time over several tens of orbits.

in the stresses. However here the heating is supplied self-consistently through the dissipation of MRI turbulence from the preceding MRI outburst.

The difference in flow structure between the MRI and convection-dominated phases is also seen in a more quantitative manner in 2D power spectra of the specific vertical kinetic energy (see Figure 4.8). During convection-dominated phases (a time-average over one of which is shown in the left-hand panel), the kinetic energy is clearly contained in a narrow band along the k_x axis, representative of the updrafts and downdrafts that dominate the flow structure within $\sim 2H_0$ during the convection dominated phase. During MRI-dominated phases on the other hand, the energy is distributed more evenly in the (k_x, k_z) plane.

Vertical structure of the disk

In Figure 4.9 we show the vertical profiles (averaged over the x - and y -directions and over time) of the vertical mass and heat fluxes (top row) and buoyancy frequency squared (bottom row). The left-hand column consists of profiles time-averaged over snapshots from an MRI-dominated phase

(orbits 30 to 37), whereas the right-hand column consists of profiles time-averaged over snapshots from a convection-dominated phase (orbits 144 to 153). There are both qualitative and quantitative differences in the vertical profiles compared to those measured in our box-cooled MRI simulation (see Figure 4.3), indicating that the MRI/convective cycles can rearrange the background structure of the disk, though as we explain below, the results should be interpreted with caution.

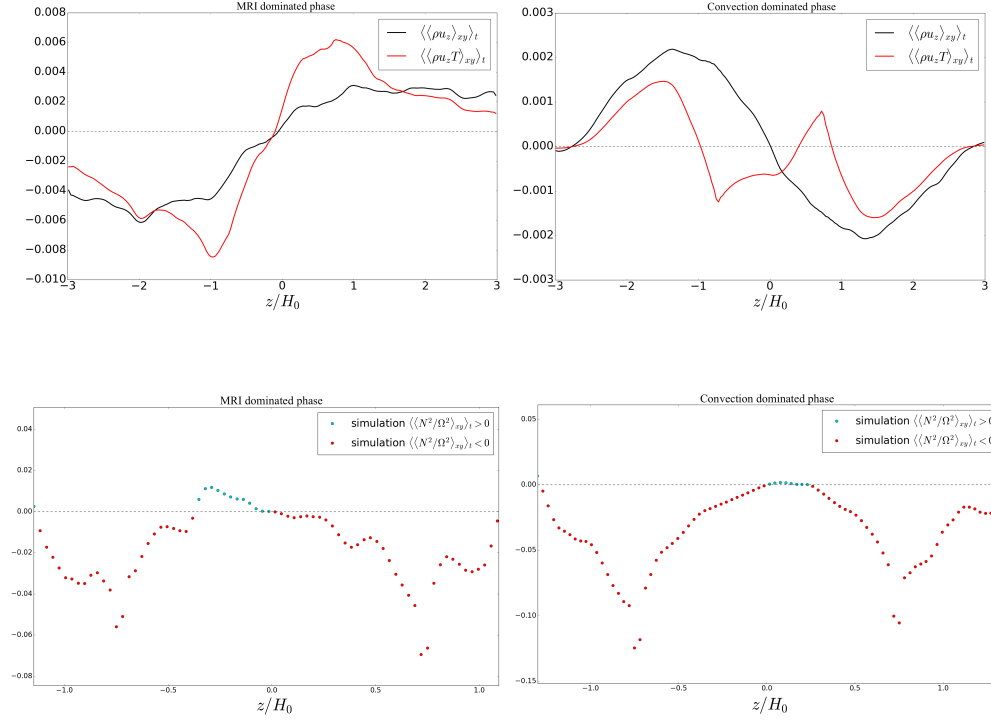


Figure 4.9: Vertical disk structure from a simulation exhibiting MRI/convective cycles (NSTRMC44e1). Left column: horizontal- and time-averaged vertical mass and heat fluxes (top panel) and buoyancy frequency squared (bottom panel) taken from an MRI-dominated phase (orbits 30 to 37). Right column: the same diagnostics but taken from a *convection*-dominated phase (orbits 144 to 153). The simulation was run with a constant cooling timescale of $\tau_c = 10$ orbits above $|z| = 0.75H_0$, and a uniform explicit resistivity of $\eta = 5 \times 10^{-3}$.

During the MRI-dominated phase (left-hand column of Figure 4.9) the heat and mass flux are both directed outwards, consistent with the transport properties of MRI turbulence (see the top panel of Figure 4.3 taken from our box-cooled simulation). Unlike in our box-cooled simulation, however, in which MRI heating and box-cooling alone could not set-up an unstable entropy gradient, the buoyancy frequency during the MRI-dominated phase in the simulation with explicit cooling is negative within $\sim 1.2H_0$ of the mid-plane, reaching a minimum value of around $\min(\langle\langle N_B^2 \rangle_{xy} \rangle_t) \sim -0.07$ around $|z| \sim 0.75H_0$. As we have already emphasized, however, $N_B^2 < 0$ does not automatically imply convection, and thus the negativity of the buoyancy frequency during the MRI-dominated phase does not mean that there is convection occurring in tandem with the MRI. In the *convection*-dominated phase (right-hand column) on the other hand, the vertical heat flux is also directed away from the mid-plane but the vertical mass flux is directed *towards* the mid-plane, which we have often found to be a reliable, if transient, sign of the onset of convection. The buoyancy frequency is negative within $\sim 1.2H_0$, but its minimum value is twice what it is in the MRI-dominated phase with $\min(\langle\langle N_B^2 \rangle_{xy} \rangle_t) \sim -0.12$ around $|z| \sim 0.75H_0$. Finally, the temperature profile time-averaged over the duration of the simulation (orbit 50 to orbit 200; not shown), drops monotonically from the mid-plane compared to the flat topped, isothermal profile observed within around $2H_0$ of the mid-plane in our box-cooled simulation.

An interesting feature is that the vertical heat flux during the convective-dominated phase actually appears to be less than that during the MRI-dominated phase, suggesting that turbulent convection is not so efficient at transporting heat as to completely remove the residual heat left over from the previous MRI-dominated phase. For example during the third MRI burst (orbits 89 to 95) the maximum vertical heat flux is $\max(\langle\langle F_z \rangle_{xy} \rangle_t) \sim 0.0033$, whereas during the following convection-dominated phase the maximum heat flux is just $\max(\langle\langle F_z \rangle_{xy} \rangle_t) \sim 0.0009$ (orbits 103 to 115). Conversely, the buoyancy frequency is consistently higher when averaged over the convection-dominated phases than when averaged over the MRI-dominated phases: $\min(\langle\langle N_B^2 \rangle_{xy} \rangle_t) \sim -0.13$ during the convection-dominated phase (or-

bits 103 to 115), compared to $\min(\langle\langle N_B^2 \rangle_{xy} \rangle_t) \sim -0.04$ during the preceding MRI dominated phase (orbits 89 to 95). This pattern, of large (minimum, i.e. most negative) N_B^2 but small heat flux when convection is dominant and small (minimum) N_B^2 but large heat flux when the MRI is dominant, is one that we find consistently in our simulations, including those run at shorter cooling timescales (see Section 4.3.5). We caution, however, that taking meaningful time-averages is difficult here, first due to the relatively short-intervals over which the time-averages have been taken (because the phases themselves only last several tens of orbits at best), and second because the disk is not in thermal equilibrium in either phase. During the MRI bursts heat is being built up due to the dissipation of MRI turbulence, whereas during the convective-dominated phase heat is no longer being built up but redistributed. Thus the buoyancy frequency is changing during both phases, something which might explain the asymmetry in N_B^2 in Figure 4.3.

4.3.4 MRI-dominated

For lower ($\eta < 5 \times 10^{-4}$) resistivities (but at the same cooling time of $\tau_c = 10$ orbits), we do not observe MRI/convective cycles reported in the previous section. More importantly we also fail to find evidence that the MRI and convection ‘coexist’, in the sense that they are continuously present at the same time. Rather we simply find MRI turbulence, though we do occasionally observe what appear to be convective bursts in the flow field, and these become more frequent in our more marginal simulations (e.g. at $\eta = 2.5 \times 10^{-4}$).

An example of an MRI-dominated run is given by simulation NSTRMC44f1 (with parameters $\tau_c = 10$ orbits and $\eta = 10^{-5}$).³ The time-evolution of the kinetic and magnetic energy densities, and of the stresses, in the non-linear phase are similar to those measured in our box-cooled and uniformly cooled simulations with $E_{\text{kin},z} \sim 10^{-3}$ and $E_{\text{mag}} \sim 10^{-2}$, respectively. While these quantities fluctuate with time in the non-linear phase, we observe no dis-

³We have rerun this simulation with double the resolution (i.e. 64 cells per H_0) to ensure that the resistive scales were resolved; see NSTRMC44f1HR in Table C.3.

cernible bursts in the stresses as we do in simulations with larger explicit resistivity (see Section 4.3.3). The time- and volume-averaged value of alpha is $\langle\langle\alpha\rangle\rangle \sim 0.015$ (compared to $\langle\langle\alpha\rangle\rangle \sim 0.014$ in our box-cooled simulation).

NSTRMC44f1 is an important example of the limitations of relying solely on the sign of the buoyancy frequency as a diagnostic for convection. The horizontal- and time-averaged vertical buoyancy frequency profile is negative within around H_0 of the mid-plane, reaching a minimum value of $\min(\langle\langle N_B^2 \rangle_{xy} \rangle_t) \sim -0.14$. As we have already alluded to, however, this does not imply convection. The motion of any one fluid element is governed by the net force on the fluid element. This will be due to a combination of turbulent stresses due to the MRI and the buoyancy force due to the unstable entropy gradient. In the limit of small resistivity (i.e. large turbulent stresses due to the MRI) and a long cooling timescale (i.e. a small entropy gradient), the forces on the fluid element due to the MRI will greatly exceed whatever buoyancy force there is on the element due to the unstable entropy gradient. Thus it does not make much sense on physical grounds to argue that the MRI and convection are present at the same time. Indeed despite the negativity of the buoyancy frequency we cannot find any evidence for convection in the flow field.

To investigate the effects of cooling prescription on vertical structure more quantitatively, in Figure 4.10 we plot the 2D power spectrum of the specific kinetic energy for our box-cooled simulation (top row), uniformly cooled simulation (second row), and height-dependent cooling simulation with explicit resistivity $\eta = 10^{-5}$ (third row). We have also included the spectrum from the MRI-dominated phase of the height-dependent cooling simulations exhibiting MRI/convective cycles which was run with a resistivity of $\eta = 5 \times 10^{-3}$ (bottom row). The colorbars are logarithmic, and each row corresponds to a separate simulation. The left-hand column shows the spectra calculated by taking a y -average over the domain, whereas the right-hand column shows the spectra from a single y -slice (we will discuss the latter below). For each spectrum the data has been taken within $\pm H_0$ of the mid-plane and averaged in time over several tens of orbits.




Figure 4.10: Comparison of 2D specific kinetic energy spectra from simulations with different cooling prescriptions. See text for further details.

The y -averaged spectra for the box-cooled and uniformly cooled simulations are practically indistinguishable, and power is distributed more or less evenly along both the k_x and k_z directions. The y -averaged spectrum in the height-dependent cooling simulation (third row) is slightly, but visibly, flatter than the other two simulations (first and second rows), however, the flatness is far less pronounced than in the spectrum captured during the convection-dominated phase of the MRI/convective cycles simulation we discussed in the last section (see the left-hand panel of Figure 4.8). We have also plotted the spectrum from the MRI-dominated phase of the simulation NSTRMC44e1 exhibiting MRI/convective cycles (fourth row). This too has power slightly more distributed along the k_x direction compared to the box-cooled and uniformly cooled simulations, but is less elongated than the spectrum from the $\eta = 10^{-5}$ simulation. The slightly elongated spectrum in $\eta = 10^{-5}$ simulation (third row) is somewhat surprising result given that we cannot find evidence of coherent vertical structure (e.g. convective plumes) in the flow field in this simulation. It is possible that because the height-dependent simulation is unstably stratified around the mid-plane, vertical motions are less suppressed than they are in the stably stratified box-cooled and uniformly-cooled simulations, hence accounting for the more elongated profile. In any case, the key point is that the difference between the spectra is very marginal in all four cases, suggesting that the vertical structure in the simulations with height-dependent cooling is very similar to that in the convectively stable box-cooled and uniformly cooled runs.

Next we compare our height-dependent cooling simulations. We find similar results for simulations with resistivities $10^{-5} \leq \eta \lesssim 2.5 \times 10^{-4}$. The left-hand column of Figure 4.11 compares the y -averaged 2D spectra from three simulations all run with a cooling timescale of $\tau_c = 10$ orbits, but with resistivities of $\eta = 10^{-5}$, 10^{-4} , and 2.5×10^{-4} , respectively (thus as we move down a column in Figure 4.11 we are moving along a horizontal skewer in the (τ_c, η) parameter space; see Figure 4.13). There is very little difference between the spectra for the $\eta = 10^{-5}$ and $\eta = 10^{-4}$ runs, and both appear significantly more isotropic compared to the spectrum that we observed during the convective-phase of the simulation exhibiting MRI/convective cycles

(see the left-hand panel of Figure 4.8), though still slightly more elongated in the k_x direction than the box-cooled and uniformly cooled runs. To detect vertical structure associated with convection we have also tried looking at the 2D spectra of different quantities (vertical mass flux, vertical heat flux, density, and pressure) but we cannot detect any appreciable difference between the $\eta = 10^{-5}$ and $\eta = 10^{-4}$ simulations, nor enhanced structure in the vertical direction, for any of these quantities either.




Figure 4.11: 2D power spectra of specific vertical kinetic energy from a simulations with cooling implemented above $|z| = 0.75H_0$. The colorbars are logarithmic. Each row corresponds to a separate simulation with height-dependent cooling, with the resistivity increasing from top to bottom, and all simulations employed a cooling timescale of $\tau_c = 10$ orbits. The left-hand column shows y -averaged spectra, while the right-hand column shows the spectra calculated from an individual y -slice. For each spectrum the data has been taken within $\pm H_0$ of the mid-plane and averaged in time over several tens of orbits.

One drawback of using y -averaged spectra, is that these only pick up axisymmetric structures. Thus in the right-hand column of Figure 4.11 we plot data taken from a fixed y -slice in each simulation (though we still average the same y -slice over many successive snapshots to form a time-average). These spectra appear to have power noticeably more distributed in the k_y direction compared to the y -averaged spectra, revealing that non-axisymmetric structures have more pronounced vertical structure than axisymmetric ones, and that the former have a different k_x/k_z profile. However, as was the case for the y -averaged spectra, there is little difference between the $\eta = 10^{-5}$ and $\eta = 10^{-4}$ simulations when comparing the spectra taken from individual y -slices. Furthermore the extra structure in k_z observed in the y -slices (compared to y -averaged data) is also apparent in the box-cooled and uniformly cooled simulations which are convectively stable (see Figure 4.10). Thus it is unlikely that the extra vertical structure we observe in the y -sliced data compared to the y -averaged data is due to the presence of convective cells.

For the most marginal case (the $\eta = 2.5 \times 10^{-4}$ run), we do observe what looks like the occasional convective burst in the flow field, and the spectrum appears flatter than in the previous two examples, but only very marginally so. It is somewhat difficult to classify this run, but from the time-evolution of α we do not observe bursts as distinct as we do in the $\eta = 5 \times 10^{-4}$ simulation (which we classified as exhibiting MRI/convective cycles). In particular the reader should compare the blue curve (showing the time-evolution of the volume-average of α) in Figure 4.12, corresponding to the $\eta = 2.5 \times 10^{-4}$ run, to the blue curve in Figure 4.16, which is for a run at $\eta = 5 \times 10^{-4}$. There is not much resemblance between the two. On the other hand, we find significant overlap between time-evolution of α in the runs with $\eta = 10^{-5}, 10^{-4}$, and 2.5×10^{-4} , respectively (see the black, red, and blue curves, respectively, in Figure 4.12). This suggests that in the marginal case, it is still MRI turbulence that dominates the flow rather than than convection.




Figure 4.12: Time-evolution of volume-averaged alpha viscosity from simulations run with height-dependent cooling and explicit resistivity. The cooling timescale is the same at $\tau_c = 10$ orbits in all three simulations, but the resistivity is different. There is significant overlap in the time variable behavior of $\langle \alpha \rangle$ between all three simulations.

4.3.5 Parameter survey

In Sections 4.3.3-4.3.4 we found that for a fixed cooling timescale of $\tau_c = 10$ orbits we could achieve two distinct outcomes for the interaction between the MRI and convection by varying the resistivity (interpreted as ‘MRI strength’). At high resistivities ($\eta \geq 5 \times 10^{-4}$) we found MRI/convective cycles. At lower resistivities ($\eta \leq 2.5 \times 10^{-4}$) we found the MRI, essentially recovering the results of our inviscid simulations, though we emphasize that marginal cases ($\eta \sim 2.5 \times 10^{-4}$ in particular) are difficult to classify, because on the one hand they do not exhibit bursts in the volume-average of α characteristic of MRI/convective cycles, on the other hand we do sporadically see some vertical structure in the flow and also in the 2D spectra. According to the argument presented in Section 4.3.2, each simulation should be charac-




Figure 4.13: Parameter space of simulations run with beta cooling (plotted on the ordinate as the inverse of the cooling timescale τ_c) and explicit resistivity η . Blue dots correspond to simulations that exhibited MRI only (see Section 4.3.4), and red triangles correspond to simulations that exhibited MRI/convective cycles (see Section 4.3.3). The black diamond indicates a simulation in which the disk was disrupted (see text for further details).

terized by an effective Rayleigh number, whose value can be controlled by changing either the explicit resistivity as in Section 4.3.2 (thus controlling the ‘MRI strength’), *or* by changing the cooling timescale (thus controlling the ‘convective strength’, at least once the Rayleigh number is large enough for convection to begin in the first place). Thus in the regime of low resistivity but short cooling timescales we expect to recover the MRI/convective cycles observed at high resistivity and long ($\tau_c \sim 10$ orbits) cooling timescales. To investigate this, and to determine the parameter-space boundary between the different outcomes of MRI and convection, we present in this section a parameter survey in the space of cooling timescale τ_c against explicit resistivity η .

The results of our parameter survey are plotted in Figure 4.13 which

shows simulations in the space of inverse cooling timescale $1/\tau_c$ against explicit resistivity η . For a fixed resistivity, we run a series of simulations at successively lower cooling timescales, with each simulation being restarted from the end state of the previous simulation (thus for each resistivity we move upwards along a vertical skewer in Figure 4.13). For example, the simulation with $(\eta = 10^{-4}, \tau_c = 5 \text{ orbits})$ was restarted from the end of the simulation with $(\eta = 10^{-4}, \tau_c = 10 \text{ orbits})$.

An example of the typical behavior we observe as the cooling timescale is lowered (keeping the resistivity fixed) is given by the vertical skewer at $\eta = 10^{-4}$ (simulations NSTRMC44c1 through NSTRMC44c7 in Table C.3). As expected, as we decrease the cooling timescale, we move from a regime in which the simulations exhibit only MRI at $\tau_c = 10$ orbits (e.g. simulation NSTRMC44c1) to a regime characterized by MRI/convective cycles at $\tau_c = 1$ orbit (simulation NSTRMC44c6). In the latter regime, we observe outbursts in the vertical kinetic energy density and in α (see Figure 4.14) as well as large-scale convective cells in the flow field (see Figure 4.15), both of which are characteristic of the MRI/convective cycles observed in our most resistive run (see Figure 4.5). The pressure (green curve) increases during outbursts as the disk heats up due to the dissipation of MRI turbulence, and drops during the convection-dominated phases. We also plot (in red) the ratio of vertical to radial kinetic energy (for clarity we have plotted a rolling average over 2.5×10^4 time-steps or about 8 orbits). This is generally anti-correlated with the vertical kinetic energy and α (notice, in particular, the peaks in the ratio of vertical to radial KE around orbit 695, 725, 810, and 875, which coincide with troughs in the vertical KE). This anti-correlation mirrors that observed in our most extreme MRI/convective cycles simulations (see Figure 4.5).⁴ The behavior is consistent with what we expect from having a convection-dominated phase (during which the ratio of vertical KE

⁴During troughs in the ratio of vertical to radial KE (MRI-dominated phases) $E_{\text{kin},z}/E_{\text{kin},x} \sim 0.7$, which is consistent with the value of $E_{\text{kin},z}/E_{\text{kin},x} \sim 0.7$ measured during the troughs in simulation NSTRMC44e1 (see Figure 4.5). During the peaks (convection-dominated phases) $E_{\text{kin},z}/E_{\text{kin},x} \sim 0.8\text{-}0.9$, smaller than the value of $E_{\text{kin},z}/E_{\text{kin},x} \sim 1.2$ measured in NSTRMC44e1.




Figure 4.14: Top: time-evolution of volume-averaged α (in black) and pressure (in green) from simulation NSTRMC44c6 ($\tau = 1$ orbit, $\eta = 10^{-4}$), exhibiting the outbursts characteristic of MRI/convective cycles. Bottom: time-evolution of vertical kinetic energy density (black) and a rolling average of the ratio of vertical to radial kinetic energy (red); the red curve is plotted on a semi-log plot. The red dashed line indicates where the cooling timescale was lowered to $\tau_c = 1$ orbit from $\tau_c = 2$ orbits. The blue dashed lines indicate restart positions.




Figure 4.15: Colorplots of the vertical component of the velocity in the xz -plane taken from snapshots at different times from a simulation with a very low cooling timescale, which exhibited MRI/convective cycles (NSTRMC44c6). (a) convection-dominated phase, (b) MRI-dominated phase. The simulation employed a constant cooling timescale of $\tau_c = 1$ orbit above $|z| = 0.75H_0$, and a uniform explicit resistivity of $\eta = 10^{-4}$.

to radial KE should be enhanced), and an MRI-dominated phase (where we expect radial KE to dominate). The flow in NSTRMC44c6 is quite turbulent, however, and the phase shift between the two curves is not easy to see for the first 100 orbits after the cooling timescale was lowered (i.e. up to 700). Therefore, rather than having two distinct MRI-dominated and convective dominated phases (as in our most resistive simulation), we think it is more likely that the two processes (MRI and convection) are competing with one another in this case, with convection periodically overtaking the MRI.

Although we have classified both NSTRMC44c6 (Figure 4.14) and NSTRMC44e1 (Figure 4.6) as simulations exhibiting MRI/convective cycles, it is clear just from the time-evolution of α that the cycles are of different nature

in each simulation. In NSTRMC44e1 the resistivity is large enough to kill the MRI entirely during the ‘quiescent’ periods. Residual heat left over from the MRI-dominated phase then leads to the onset of convection, which can exist because there is no longer any MRI to disrupt the convective updrafts. Convection itself then appears to reseed the MRI, and so on (we discuss this in more detail in Section 4.4.1). In simulations with $\eta < 10^{-3}$, on the other hand, the MRI is never quite killed off for lengthy periods of time by the resistivity alone. This means that in this case it is harder to get convection compared to the highly resistive run, because the resistivity doesn’t inhibit the MRI to same degree, and the latter continues to redistribute heat vertically (via an effective thermal diffusivity), thus impeding the onset of convection. At small resistivities, but low cooling timescales, the unstable entropy gradient is sufficiently large to enable the onset of convection for short periods of time, resulting in intermittent bursts of convection. These might then potentially reseed the MRI again, in a manner similar to what is occurring in the NSTRMC44e1, which leads to the bursts in α . During the burst the MRI becomes dominant again and thus convection ceases, and so on.

For the simulation with the lowest cooling timescale (NSTRMC44c7 with $\tau_c = 0.5$ orbits) the disk is disrupted: specifically, the disk appears to collapse towards the mid-plane near the radial center of the box at $x = 0$, leaving two large clumps in density either side of $x = 0$ with very little material in between. We observe no outbursts for the duration of the simulation, although there remain occasional signs of large-scale convective cells in the flow field. Finally, we observe a similar transition, i.e. from a run exhibiting behavior similar to that of the MRI-only runs at long cooling timescales ($\tau_c = 10$ orbits) but transitioning to MRI/convective cycles at short ($\tau_c = 1$ -2 orbits), along the vertical skewer at $\eta = 2.5 \times 10^{-4}$. In this case the transition occurs at a slightly longer cooling timescale ($\tau_c = 2$ orbits) than it did along the $\eta = 10^{-4}$ skewer, which is not surprising given that the ($\eta = 2.5 \times 10^{-4}$, $\tau_c = 10$ orbits) run was already quite marginal.

4.4 Discussion

We found that vertically stratified simulations of the MRI with a perfect gas equation of state, a height-dependent explicit cooling prescription, and uniform explicit resistivity can lead to convective/MRI cycles. Several questions naturally arise. To what extent are the MRI bursts (i.e. enhancements in α) actually related to convection? Furthermore, if the outbursts are related to convection, to what extent can we find support for the hypothesis presented in Hirose et al. (2014) that convection can ‘seed’ net-vertical flux MRI? Can narrow boxes (in radius) mitigate the large-scale convective cells (which typically have a width $\sim H_0$) observed in the convective/MRI cycles simulations? Finally, what are the applications to dwarf novae and to previous work on MHD convection in disks? These questions are discussed in the following subsections.

4.4.1 Effect of explicit resistivity

Simon et al. (2011) (SHB2011) carried out vertically stratified *isothermal* MRI simulations with explicit resistivity and viscosity, and found that at certain magnetic Prandtl numbers Pm (where $Pm \equiv \nu/\eta$) accumulation of toroidal field could switch the MRI back on again after it was initially killed by the resistivity. To what extent are the results of our MRI/convective cycles simulation a manifestation of this phenomenon? To investigate this, we reran the simulation described in Section 4.3.3 *without* any explicit cooling, but otherwise with exactly the same set-up, including an explicit uniform resistivity of $\eta = 5 \times 10^{-3}$ (see simulation NSTRMC44e1NoCool in Table C.5). Note that since there is no explicit cooling in this simulation, any heating present (e.g. viscous heating due to MRI turbulence) will cause the disk to expand until thermal equilibrium is established between viscous heating and cooling due to advection across the vertical boundaries (i.e. ‘box cooling’).

We find that vertical kinetic energy density drops by an order of magnitude immediately after initialization (from $\sim 5 \times 10^{-3}$ to $\sim 10^{-4}$), as does

the magnetic energy density (from $\sim 3 \times 10^{-2}$ to $\sim 4 \times 10^{-3}$). Thereafter the kinetic energy continues to decrease steadily (though fluctuating) over the duration of the simulation (100 orbits). There is a gradual increase in the magnetic energy, which saturates at around 10^{-2} by the end of the simulation. There is a steep drop in both the Reynolds and Maxwell stresses after initialization and both remain of order 10^{-5} for the duration of the simulation. In addition the flow field appears laminar and there are neither visible signs of convection, nor of MRI turbulence. We conclude that in the absence of convection the explicit resistivity kills the MRI which is not re-seeded by some other process for the duration of the simulation. Thus we can confirm that, at least for the largest resistivity we investigate, the outbursts and MRI/convective cycles are not being driven by resistivity.

At lower resistivities the picture is somewhat more complicated. In Figure 4.16 we plot the time-evolution of α (in blue) from a simulation (NSTRMC44b1) with both height-dependent cooling (with $\tau_c = 10$ orbits) an explicit resistivity of $\eta = 5 \times 10^{-4}$, which is an order of magnitude smaller than that used in NSTRMC44e1. This simulation also exhibits clear bursts in α , though the peaks are smaller ($\alpha \sim 0.015$ - 0.025 , about half the size of the peaks in α observed in the $\eta = 5 \times 10^{-3}$ run), and spaced closer together. Other than that the results are broadly similar to those in NSTRMC44e1, and we observe both large scale convective cells between the bursts in α and MRI turbulence during the bursts. One major difference however, is the behavior of this simulation when the cooling is turned off. We have rerun NSTRMC44b1 with exactly the same set-up, but without cooling (NSTRMC44b1NoCool in Table C.5), and plotted the time-evolution of α from this run as a black curve in Figure 4.16. The buoyancy frequency is positive everywhere in this run, and thus there is no convection. Unlike at a resistivity of 5×10^{-3} , where turning the cooling off killed the MRI for good, in this simulation we find bursts in α , though these are smaller in magnitude than those in the simulation with cooling.

It is tempting to draw parallels between this (resistive only) simulation and some of the vertically stratified isothermal resistive simulations of SHB2011 (in particular see the red curve on the right-hand panel of their




Figure 4.16: Time-evolution of volume-averaged alpha viscosity from a simulation run with height-dependent cooling and explicit resistivity with ($\tau_c = 10$ orbits, $\eta = 5 \times 10^{-4}$) (blue curve), and from a simulation run with the same explicit resistivity but no cooling (black curve).

Figure 10). Thus NSTRMC44b1 without cooling might be the perfect gas analogue of what SHB2011 found in their isothermal simulations, where accumulation of toroidal magnetic field, rather than convection, was reseeding the MRI. This comparison should be made with caution, however, as the magnetic Reynolds number is smaller in our simulation ($R_m = 2000$) than in the aforementioned simulation of SHB2011 (where $R_m = 3200$). In addition SHB2011 varied the explicit viscosity (and thus the magnetic Prandtl number) in their simulations rather than the resistivity. Finally the time scales of the bursts are very different between the two simulations. In the simulations of SHB2011 the stress builds up and then decreases again over several hundred orbits, whereas the bursts observed in our resistive-only simulation recur much more rapidly, on timescales of just 10-15 orbits. In any case, while it appears that resistivity is contributing to the bursty behavior in this simulation, the fact that the bursts are consistently larger when height-dependent

cooling is added seems to suggest that convection can enhance this behavior.

4.4.2 Is convection seeding net-vertical-flux MRI?

This brings us to our second question: is convection potentially reseeding net-vertical-flux MRI leading to an enhancement in α as claimed by Hirose et al. (2014)? Hirose et al. (2014) hypothesize that vertical convective motions drag magnetic field lines upwards, thus creating net vertical field that seeds the (powerful) net-vertical field MRI (see (Hawley et al., 1995b)). Note that the addition of an explicit resistivity in our simulation complicates this scenario because it breaks the approximation of ideal MHD, and thus mitigates fluid-field line coupling. The net-vertical flux MRI is characterized by radial *channel flows* during its linear phase. It is unlikely, however, that channel flows are being seeded in this simulation given the short radial wavelength of the B_z fluctuations, and indeed visual inspection of the flow field (i.e. B_x and B_y in the xz -plane) reveals no streaky motions (in x) characteristic of these channel flows. In addition we find no noticeable difference in B_x and B_y at the onset of an outburst compared to the middle of an outburst. More revealing however, is the time-evolution of the root-mean-square vertical magnetic field component $\langle \sqrt{B_z^2} \rangle$ (see Figure 4.17). This exhibits a clear phase shift compared to the time-evolution of $\langle \alpha \rangle$: the vertical magnetic field appears to increase just before $\langle \alpha \rangle$ does, which might be indicative that convection is building up vertical field just before an outburst. The lag between the rms vertical magnetic field and the stress is approximately a few orbits, which is comparable to the MRI growth time. Note that having stronger vertical magnetic field, even if it varies with x , still aids in seeding the MRI, even if doesn't quite seed the channel modes associated with net vertical flux MRI.

It should also be mentioned that, quite aside from the generation of a mean B_z , it has long been established that convection can itself work as a turbulent dynamo, producing disordered magnetic field up to equipartition strengths (see review by Rincon (2019) and references therein). The zero net flux MRI might be able to use this dynamo field as a seed, though we don't




Figure 4.17: Time-evolution of volume-averaged alpha viscosity (black), root-mean-square vertical component of the magnetic field (red), and pressure divided by 10 (green) from a simulation exhibiting MRI/convective cycles (NSTRMC44e1) with ($\tau_c = 10$ orbits, $\eta = 5 \times 10^{-3}$).

pursue this angle explicitly, concentrating on the more straightforward and probably more dominant effect of the large-scale B_z generated.

Another difference between our simulation and that of Hirose et al. (2014) is the relationship between the time-evolution of $\langle \alpha \rangle$ and the time-evolution of the pressure $\langle P \rangle$. Hirose et al. (2014) find that in one of their upper branch simulations (where convective/radiative outbursts were reported), $\langle \alpha \rangle$ is *anti*-correlated with thermal pressure. When heat builds up at the mid-plane in their simulations the pressure increases, which, in the absence of an increase in stress, leads to a drop in alpha. Conversely, when heat is removed from the mid-plane the pressure drops, and in the absence of a decrease in stress, this contributes to the burst in $\langle \alpha \rangle$ (see Figure 7 of their paper). In our simulation the opposite is true: $\langle \alpha \rangle$ is well correlated with the stresses (see Figure 4.6). Thus it is the rise in *stress* that is enhancing $\langle \alpha \rangle$ in our simulation rather than a drop in pressure.

4.4.3 Dependence on box size

All the flux-limited-diffusion simulations exhibiting convective/radiative cycles in Hirose et al. (2014) were carried out in very narrow boxes (H_0 in the radial direction). The disadvantage of narrower boxes is that there exists the possibility that the radial box size can preclude the turn-over of convective cells. This might particularly be true in the case of the large-scale convective cells observed in NSTRMC44e1: the typical size of the large-scale convective eddies in this run was found to be $\sim 1.3H_0$. Thus we briefly investigate the dependence of the results of this simulation on the radial box size by comparing the results of NSTRMC44e1 ($L_x = 4H_0$) with NSTRMC46 ($L_x = 2H_0$) and NSTRMC45a3 ($L_x = H_0$), keeping all other parameters the same.

As the box size is reduced, we find that the number of outbursts is reduced. We observed 3 outbursts during the first 100 orbits in the $L_x = 4H_0$ simulation, 2 outbursts in the $L_x = 2H_0$ simulation, and just one outburst in the $L_x = H_0$ simulation. The onset of the first outburst is also delayed as the radial box size is reduced: from around orbit 25 (at $L_x = 4H_0$) to around orbit 40 (at $L_x = 2H_0$) up to orbit 60 (at $L_x = H_0$). The magnitude of the outbursts (measured in terms of the maximum value of $\langle\alpha\rangle$ measured during the outbursts) is about the same at all three radial box sizes. Crucially large-scale convective cells are observed in the flow field only in the $4H_0$ simulation: the $2H_0$ and H_0 simulations exhibit only narrow convective cells during their quiescent phases. Thus it appears that radial box sizes of $2H_0$ and H_0 are insufficient to observe the large-scale convective cells.

4.4.4 Applications to dwarf novae

The convective/radiative cycles observed in the radiative transfer simulations of Hirose et al. (2014) (HBK2014) have outbursts that recur on timescales similar to those in our MRI/convective cycles runs. However, besides the more complicated radiative transfer physics, which we mock up using a simple height-dependent cooling prescription, there are appreciable differences between our results and those of HBK2014. One major difference is that the cycles in HBK2014 are driven by opacity changes (which we do not include

in our models), specifically the dependence of opacity on temperature and pressure. In brief, first a drop in pressure and temperature leads to an increase in opacity. Heat is trapped at the mid-plane as a result, setting up an unstable entropy gradient. Convection ensues (at least according to the diagnostics HBK2014 use to detect convection, mainly the sign of the buoyancy frequency and the ratio of advective to total [i.e. radiative + advective] heat flux), and the convective plumes seed vertical magnetic fields. This leads to enhanced MRI turbulence and an enhancement in stress. The dissipative heating increases as the stress does, causing the temperature and pressure to increase. This now results in a *decrease* in opacity which eventually suppresses convection. The dissipation rate then drops (because convection is no longer seeding net vertical flux MRI), the temperature and pressure drop accordingly, and the cycle is repeated. In fact, stress and pressure are nearly out of phase in their simulations, such that the pressure decreases during an outburst, whereas in our simulations the thermal pressure increases during outbursts (see Figure 4.17). Furthermore, convection appears to be present *during* the outbursts in their simulations, whereas we find the opposite in our runs, i.e. that convection is suppressed during outbursts (when MRI turbulence is dominant), but dominant during quiescence.

Another difference between the cycles observed in HBK2014 and those in our simulations, is that the heat transport in their simulations is alternately dominated by advection (or convection, according to their criteria) and radiation, whereas in our simulations we have not included radiative transport so the heat flux is necessarily always advective (which can be due to convection, or due to the MRI, which we have shown is quite efficient at transporting heat vertically on its own). Coleman et al. (2018) carried out local, radiative MHD simulations very similar to those of HBK2014, but employed opacity and equation of state tables relevant to the helium-rich AM CVn disks. A key difference in most of their simulations compared to those of HBK2014 is that the heat transport is persistently dominated by advection (or convection, according to their criteria), which they attribute to the system lying between the two opacity maxima resulting from the two ionizations of helium at $\sim 3 \times 10^4$ K and $\sim 7 \times 10^4$ K, respectively. Thus these results are inde-

pendent of changes in opacity, and might therefore provide a better point of comparison to ours.

Three different outcomes (each corresponding to a different initial surface density) are shown in Figure 7 of Coleman et al. (2018), which we have reproduced in Figure 4.18 for convenience. The top panel shows ‘persistent convection’, the middle panel shows ‘mostly persistent convection’, and the bottom panel shows ‘intermittent convection’ (i.e. convective/radiative cycles as seen in HBK2014). Both the top and the middle cases have a ratio of advective to total heat flux close to unity for the duration of the simulations, and are thus similar in that respect to our simulations. In particular note that in the middle panel (the ‘mostly persistent convection’ case), the advective Mach number is mostly anti-correlated with α , and that α is quite bursty. This behavior is reminiscent of the simulation NSTRMC44c6 discussed in Section 4.3.5 in which α was very bursty, the ratio of vertical to radial KE was largely out of phase with α , and in which we observed frequent bursts of convection. Indeed the opacity in their run is mostly clustered around the second Helium ionization maximum (see Figure 6 of their paper). The large opacity likely leads to a steeper entropy gradient which is more favorable for convection to dominate over the MRI. The top panel of their Figure 7, on the other hand, exhibits less bursty behavior and the advective Mach number is largely in phase with α . This is similar to the behavior of our resistive simulations with long ($\tau_c \sim 10$ orbits) cooling timescales, in which we found the flow was dominated by MRI turbulence rather than convection. Indeed in their run the opacity is largely concentrated in the trough between the two opacity maxima, and the lower opacity likely leads to a shallower entropy gradient, which is less favorable to convection. Thus what Coleman et al. (2018) refer to as ‘convection’ in this run might really just be MRI turbulence.

Apart from the time variable behavior just described, note that Coleman et al. (2018) find that $\alpha \gtrsim 0.1$ for the duration of their simulations (in all three cases). As they point out, this calls into question the role of the cycles observed in HBK2014 in enhancing α , but it is also significantly higher than the values we find in our simulations (we have seen $\alpha \sim 0.08 - 0.1$ only during

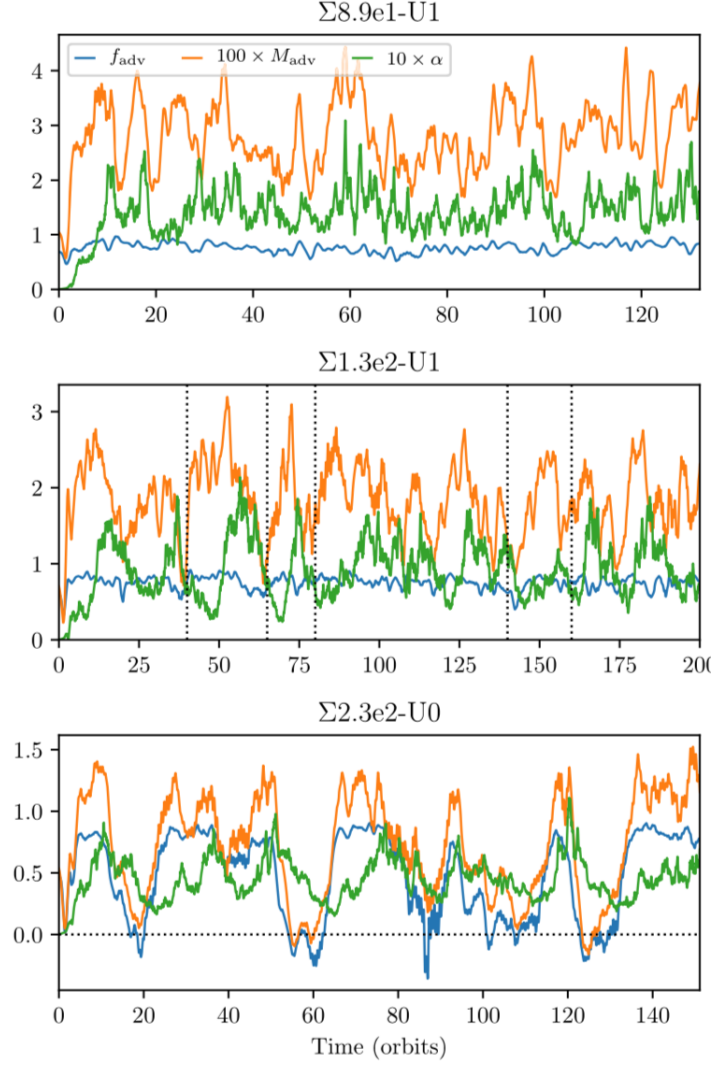


Figure 4.18: Time-evolution of $\alpha \times 10$ (green curve), ratio of vertical advective to radiative flux (blue curve), and vertical advective Mach number (orange curve) from local radiative MHD simulations modelling AM CVn disks. Taken from Coleman et al. (2018).

the very peaks of the MRI bursts in our most resistive simulations.

A further difference between our simulations and those of HBK2014 (and of Coleman et al. (2018)) is that their simulations are ideal, whereas we employ an explicit resistivity. Scepi et al. (2018a), however, did carry out radiative MHD simulations with explicit resistivity. In a run in which the magnetic Reynolds number within the mid-plane was around $Rm \sim 2 \times 10^4$, they did not observe radiative/convective cycles. This is consistent with our run NSTRMC44c1 (with $Rm \sim 10^4$) in which we also do not find any cycles. At lower magnetic Reynolds numbers ($Rm < 10^4$), however, Scepi et al. (2018a) found that the MRI died, whereas we observe convection which reseeds the MRI, leading to the aforementioned MRI/convective cycles.

Finally, we wish to comment on the relevance of our results to dwarf novae. With respect to the resistivity, Gammie & Menou (1998) have estimated that the magnetic Reynolds number in the U Gem type dwarf nova SS Cyg in quiescence is around $Rm \sim 10^3$. This falls squarely within the regime in which we observe MRI/convective cycles and powerful MRI outbursts: for example, our run NSTRMC44b1 had $Rm = 2 \times 10^3$ (see Figure 4.16). An interesting corollary of our results is that we have found that convection appears to prolong MRI activity to lower Rm than the critical value of $Rm \sim 10^4$ quoted by Gammie & Menou (1998) and found in previous ZNF simulations (Fleming et al., 2000; Scepi et al., 2018a).⁵ Resistivity is not the entire picture, however, and during the quiescent phase of dwarf novae other non-ideal MHD effects, such the Hall effect, might also play a role (Coleman et al., 2016), though they are likely subdominant compared to Ohmic resistivity (Scepi et al. (2018a); see also Coleman et al. (2018) for a brief discussion). Finally, in our MRI/convective cycle runs we have referred to the convection-dominated phase as the ‘quiescent’ phase, and the MRI-dominated phase as the ‘outburst’ phase, because the stresses are low during the convection-dominated phase and high during the outburst phase. The shortest recurrence time between dwarf novae outbursts is around 7 days (see Section 1.2.3), whereas the interval between successive MRI-dominated out-

⁵Note that the simulations of Fleming et al. (2000) were unstratified.

bursts in our most resistive simulations is around 50 dynamical timescales. For a disk of size 5×10^{-3} AU around a white dwarf of mass $1 M_{\odot}$ the outermost annulus will execute 50 orbits in about 6.4 days. Interestingly, this is comparable to the recurrence time. While it is tempting to link these cycles with the outburst cycles observed in dwarf novae, however, the two phenomena generally share very different timescales, at least in the parameter regimes we explored.

4.5 Conclusions

Motivated by our findings in Chapter 3 that hydrodynamic convection can transport angular momentum outward in disks, that the MRI can act as a heat source for convection, and by claims from radiation MHD simulations that angular momentum transport can be enhanced when the two instabilities interact in dwarf novae (Bodo et al., 2012; Hirose et al., 2014), we have undertaken a study of the interplay between convection and the MRI through controlled, three-dimensional, vertically stratified, zero-net-flux (ZNF), fully-compressible MHD simulations in PLUTO.

Previous work on the topic has used the negativity of the buoyancy frequency to determine whether convection is present or not (Hirose et al., 2014). We have discussed how in a turbulent fluid this is only a necessary condition, not a sufficient one, because the turbulent transfer of heat and momentum by the MRI leads to an effective thermal diffusivity and viscosity, and therefore an effective Rayleigh number. MRI turbulence might limit the onset of convection by lowering the effective Rayleigh number below some critical value at which the unstable entropy gradient can overcome the diffusive effects of the turbulent diffusivities (though this does assume that turbulent transport coefficients behave in a diffusive manner). Furthermore we have argued on physical grounds that the two instabilities are unlikely to be additive, or to coexist simultaneously, and that we expect either one or the other to dominate the flow at any given time. The two can certainly interact in non-trivial ways, however, since the MRI provides residual heating

that can, given sufficient cooling, set up the unstable entropy gradient that is necessary for convection. Conversely we find in some of our simulations that convection itself can reseed the MRI.

To isolate the effects of the MRI alone on vertical disk structure and heat transport we carried out inviscid numerical simulations without any explicit cooling, though we still allowed for advection of mass and energy across the vertical boundaries in order to prevent the disk from heating up and expanding until it filled the box, as was the case in the simulations of Bodo et al. (2012). We find that MRI heating and box cooling alone are unable to set up an unstable entropy gradient, and the buoyancy frequency is positive everywhere in the disk, thus ruling out the possibility of convection under these conditions. Surprisingly, however, we find that the MRI itself is rather efficient at transporting heat in the vertical direction, which calls into question the use of the vertical heat flux as a diagnostic for convection, as was done in Bodo et al. (2012) and Hirose et al. (2014). Our measurements of the time-averaged values of α (and of the stresses) are broadly consistent with previous vertically stratified ZNF MRI simulations with $\langle\langle\alpha\rangle\rangle \sim 0.015$. In narrower boxes of size $L_x = 2H_0$ and $L_x = H_0$, respectively, we find that the time-averaged stresses increase as the radial box size is decreased. Furthermore the stresses becomes significantly more bursty, which might have important implications for the narrow box radiative transfer MHD simulations of Hirose et al. (2014) and Scepi et al. (2018a), and is something that should be investigated in future work.

To obtain a negative entropy gradient in MRI simulations, and therefore assure that at least a necessary condition for convection is satisfied, necessitates the use of an explicit cooling prescription. We carried out simulations with a height-dependent cooling prescription and explicit resistivity, which enabled us to control both the unstable entropy gradient (by controlling the cooling timescale) and the strength of the MRI (by controlling the resistivity). As anticipated by our physical arguments, we find that convection and the MRI tend not to coexist, or at least do not interact in an additive manner. At low resistivities ($\eta < 2.5 \times 10^{-4}$) the MRI is too strong and disrupts any convective modes before they can become coherent plumes. Despite the

negativity of the buoyancy frequency, the results look similar to those of our convectively stable box-cooled and uniformly cooled simulations. In marginal cases (e.g. $\eta \sim 2.5 \times 10^{-4}$), however, we do see intermittent convective bursts in the flow field, though these are not sustained for any appreciable length of time, nor do they lead to any enhancements in the stress.

At higher resistivities ($\eta \geq 5 \times 10^{-4}$), however, the picture is more complicated. We have observed convective/MRI cycles in which there are alternating convection-dominated and MRI-dominated phases. At the highest resistivities, the resistivity kills the MRI but the combination of residual MRI heating and height-dependent cooling leads to convection. Furthermore, we observed the same cyclical large-scale convective cells that we found in our hydrodynamic simulations (see Section 3.4), demonstrating that these structures are robust when the heat source is implemented in a self-consistent manner by residual heating from the MRI, rather than through a heat source that was put in by hand, as in our hydro simulations. Furthermore convection appears to reseed the MRI, leading to very strong MRI bursts in which α can reach as high as $\alpha \sim 0.08$. While we cannot find any evidence of the strong radial ‘channel modes’ associated with the net vertical flux MRI prior to these outbursts, we do see a clear increase in the rms vertical magnetic field prior to each outburst, suggesting perhaps that the large scale convective cells can create vertical field lines as suggested by (Hirose et al., 2014).

Chapter 5

The stress-pressure relationship

In this chapter we investigate the stress-pressure relationship. We present some of our key preliminary results here, but plan on expanding on these in the near future, and thus we regard this chapter as a work in progress. The stresses accompanying MRI turbulence are related to the pressure in the disk, and have been shown to increase and decrease with pressure. We examine the time-lag associated with this dependence and discuss its implications for thermal instability. In order to motivate our investigations we briefly review various applications of the stress-pressure lag in Section 5.1, focusing in particular on thermal stability. In Section 5.2 we discuss our set-up, initial conditions, and time-dependent cooling prescription. Finally in Section 5.3 we present our results, first in simulations without any explicit cooling (Section 5.3.1), and then in simulations where we use our cooling prescription to vary the mean pressure and study resultant stress response (Section 5.3.2). We discuss our results and planned future work in Section 5.4.

5.1 Applications

One of the main assumptions of the alpha disk model of accretion disks is a closure scheme (commonly referred to as the ‘alpha disk prescription’) for the effective viscosity which posits that the total stress is linearly proportional to the pressure, i.e. $\Pi_{xy} \sim P^q$ with $q = 1$ (Shakura & Sunyaev, 1973;

Pringle, 1981). Various local unstratified MHD numerical simulations have been carried out to test this assumption, with many finding only a weak scaling between stress and pressure (Hawley et al., 1995a; Sano et al., 2004; Simon et al., 2008; Minoshima et al., 2015). These results were called into question by Ross et al. (2016), who showed that the weak scaling reported in the aforementioned simulations was in fact a numerical artifact due to insufficiently large box size. As the box heated up (due to thermalization of MRI turbulence) the pressure scale height $H = c_s/\Omega$ would increase until it exceeded the vertical box size L_z . Thus the largest eddies (which contribute most to the stress) were confined to the box, so the stress would plateau while the pressure continued to increase. Ross et al. (2016) employed larger boxes (such that $L_z > H$ for long enough for a clear scaling relationship to be established) and explicit diffusion coefficients, and found a near linear relationship between stress and pressure with $q \sim 0.9$. Inviscid simulations and strong net vertical or toroidal magnetic fields were found to weaken this relationship, however.

Embedded in the alpha disk prescription is not only an algebraic assumption about the relationship between the stress and the pressure, but also the causal assumption that the stress responds *instantly* to the pressure. Mathematically this means that there is no separate evolution equation for the stress. The stress-pressure relationship, and the existence of a lag between the stress and the pressure in particular, has important consequences for thermal stability. Shakura & Sunyaev (1976) and Piran (1978) showed that the alpha disk model is thermally unstable in radiation pressure dominated, geometrically thin, optically thick disks (such as the inner radii of LMXBs in certain phases of their outburst) if stress scales linearly with the *total* pressure (gas + radiation) and the cooling is due to electron scattering. Physically this occurs when the heating and cooling rates have different dependencies on the mid-plane temperature. If the heating rate is sufficiently sensitive to the temperature that a small increase in the temperature leads to a large increase in the heating rate, the resultant viscous heating will overwhelm the cooling leading to thermal runaway. The heating rate is linked to the stress, and so according to the classical theory of radiation dominated disks,

a change in pressure should lead to a change in stress.

However early local MHD simulations of radiation-dominated disks failed to observe thermal runaway (Hirose et al., 2009). Furthermore they found that the stress was leading the pressure rather than vice versa, which Hirose et al. (2009) (incorrectly) ascribed as the reason for the lack of thermal instability in their simulations. Jiang et al. (2013), however, did observe thermal instability in their vertically stratified MHD simulations, and determined the real reason for the suppression of thermal instability in the simulations of Hirose et al. (2009) was the use of small box sizes. Note that the radial wavenumber of the fastest growing mode of the thermal instability is $k_x = 0$, and thus the instability has no radial length scale associated with it. Rather small boxes prevent the stress from responding to the pressure, which in turn inhibits thermal instability, as discussed in more detail below. Setting these numerical issues aside, however, the linear theory of the alpha disk model with a time-delayed response of stress to pressure does indicate that a lag of 1-10 orbits reduces the growth rate of the instability (Lin et al., 2011; Ciesielski et al., 2012).

In addition, there is a second, physically motivated, subtlety to the response between the stress and pressure. Ross et al. (2017) also carried out unstratified local simulations (in boxes that were sufficiently large to mitigate the aforementioned numerical effects) with an optically thin, power law cooling prescription (but without radiation pressure). By tuning the cooling they could set up states that were thermally stable or unstable (at least according to linear theory). In their stable equilibrium runs, Ross et al. (2017) found that when there was very little variation in the mean pressure, fluctuations in the stress could react back on the pressure through thermalization of turbulent energy (cf. Figure 4 of their paper), thus causing the pressure to act like a function of the stress, as observed by Hirose et al. (2009). Conversely thermally unstable runs underwent runaway heating or cooling, indicating that on longer timescales the stress was responding to the pressure, as required for thermal instability. A key point here is that there are two separate regimes of stress-pressure behavior: on shorter timescales fluctuations in the stress thermalize orbital energy and result in fluctuations in the pressure,

while on longer timescales it is the stress that tracks variations in the mean pressure.

Though in this chapter we restrict our attention to the response of the stress to variations in pressure, a lag between stress and pressure might be indicative that the stress lags behind other quantities, such as strain and density, too. Thus a second application of our investigations is to the *viscous overstability*. The instability arises when a component of the viscous stress tensor is in phase with the longitudinal oscillations of a density wave (or f-mode), enabling the wave to tap into the free energy of the orbital shear (Kato, 1978). A dynamical interpretation is that the viscous stress acts as an overcompensating restoring force on a fluid element that has been displaced from its orbit, causing the fluid element to overshoot its equilibrium position (Latter & Ogilvie, 2008). In this case wavecrests correspond to regions of enhanced viscous stress, which extracts orbital energy from the background shear and transfers this to the wave, enabling the wave to grow in amplitude. The net result is that density waves grow rather than being dampened. This can affect density waves in eccentric accretion disks (Ogilvie, 2001), in planetary rings (Borderies et al., 1985; Papaloizou & Lin, 1988; Schmit & Tscharnuter, 1995) and in the ‘decretion’ disks around rapidly rotating Be stars (Okazaki, 2000). The instability can be quenched, however, when the stress behaves in a more complicated manner than that prescribed by an alpha viscosity, or, more generally, when the stress response is more complicated than that prescribed by a Navier-Stokes model, in which the stress is linearly proportional to (and responds instantly to) the strain. Indeed Ogilvie (2001) has shown semi-analytically that overstability of non-axisymmetric waves is suppressed when the stress responds to the strain over a (prescribed) relaxation timescale rather than instantaneously. Thus if there is a lag between the stress and strain, the synchronization between the viscous stress and the density oscillations is lost and the mechanism enabling viscous overstability fails.

5.2 Set-up

All the results described in this chapter were obtained using 3D MHD unstratified PLUTO simulations (see Sections 2.4 and 2.5 for a more general discussion of the code and diagnostics). Most of our simulations are run without explicit diffusion coefficients (which we refer to as our *ideal* runs) and employ a resolution of 128^3 in a box of size $4H_0 \times 4H_0 \times 4H_0$ (corresponding to 32 cells per scale height). To check that our results are converged we also run simulations with explicit viscosity and resistivity; for these we employ a higher resolution of 256^3 in a box of size $4H_0 \times 4H_0 \times 4H_0$ (i.e. 64 cells per scale height).

In our 256^3 simulations we employ both an explicit resistivity of $\eta = 2 \times 10^{-4}$ and an explicit viscosity of $\nu = 8 \times 10^{-4}$, which we refer to as our *non-ideal* runs. For an initial sound speed of $c_{s0} = 1$ and scale height of $H_0 = 1$, this corresponds to a Reynolds number of $\text{Re} = c_{s0}H_0/\nu = 1250$, a magnetic Reynolds number of $\text{Rm} = c_{s0}H_0/\eta = 5000$, and a magnetic Prandtl number of $\text{Pm} = \nu/\eta = 4$. These values are the same as those employed in the 256^3 simulations with explicit diffusion coefficients of Ross et al. (2016).

Finally we use periodic boundary conditions in the vertical direction (in addition to the usual shear-periodic and periodic boundary conditions in the x - and y -directions, respectively). Thus mass should not escape the box, and we expect the total mass to be conserved in these runs. However we have found that periodic boundary conditions in PLUTO paired together with the FARGO algorithm for orbital advection (see Section 2.4) does lead to some small mass loss. Although the reasons for this are unclear,¹ the mass loss is small over the runtimes that we use. For example, we measured a percentage change in mass of around 0.3% in a 128^3 simulation run for 500 orbits, too small to warrant the implementation of a mass source term.

¹We have confirmed that the total mass *is* conserved when FARGO is turned off, so the bug is probably due to the implementation of the FARGO algorithm.

5.2.1 Cooling Prescription

In order to investigate the lag between MRI turbulent stresses and pressure, we employ an explicit cooling prescription using the built-in optically thin cooling module in PLUTO. Our aim is to control the mean pressure so that it executes quasi-periodic variations over long timescales. We expect the stress to be very bursty on short timescales, but over many cycles of the mean pressure it should be possible to extract a meaningful lag in the stress. Note that this approach differs both from that employed in Ross et al. (2016) (in which the box heated up uncontrollably), and from Ross et al. (2017) who used explicit cooling in their runs, but who did not vary the mean pressure over long timescales.

To vary the volume-averaged pressure we use a piecewise time-dependent linear cooling function of the form

$$\Lambda = \begin{cases} -P/\tau_H, & \text{('long' timescale)} \\ -P/\tau_C, & \text{('short' timescale)} \end{cases} \quad (5.1)$$

where τ_H and τ_C are two different cooling time-scales (with $\tau_H > \tau_C$). When the ‘long’ cooling timescale τ_H is activated, heating by viscous dissipation (due to the MRI turbulence) overwhelms the cooling and the mean pressure tends to rise (the ‘heating phase’). Once the volume-averaged pressure exceeds some maximum critical value $\langle P \rangle = \langle P \rangle_+$ the cooling is switched to the ‘short’ cooling timescale τ_C . Now the cooling rate exceeds the heating rate due to viscous dissipation and the mean pressure drops (the ‘cooling phase’). Once the mean pressure falls below some minimum critical value $\langle P \rangle = \langle P \rangle_-$ the cooling is switched back to the ‘long’ cooling timescale and the cycle is repeated. Using this simple time-dependent cooling module we can control the volume-averaged pressure on timescales comparable to τ_H and τ_C , and force it to oscillate between a minimum and maximum value of $\langle P \rangle_+$ and $\langle P \rangle_-$, respectively, and then investigate the subsequent stress response to the mean pressure variation.

5.3 Results

5.3.1 Heating runs

Before we implement our explicit cooling prescription to investigate the time-lag between stress and pressure, we carry out several runs with heating only. In these runs we expect the pressure to increase monotonically (due to the dissipation of MRI turbulence), and our aim is to measure the power law dependence of stress on pressure, i.e. $\Pi_{xy} \propto P^q$.

Our fiducial simulation (NSTRMC10h12) was run with explicit resistivity and viscosity (such that $\text{Re} = 1250$, $\text{Rm} = 5000$, and $\text{Pm} = 4$), and was initialized with a radial wavenumber for the vertical magnetic field of $k_x = 4(2\pi/L_x) = 2\pi$ and velocity perturbations of size $\delta u < 0.1c_{s0}$. Note that this simulation has the same dimensionless parameters (i.e. Re , Rm , and Pr), box size, resolution, and initial conditions as Ross et al. (2016). We found that choosing the right initial conditions was very important. For example with a radial wavenumber of $k_x = 2\pi/L_x = 0.5\pi$ our box would heat up too quickly after initialization, reaching a volume-averaged pressure of around $\langle P \rangle \sim 3$ after just 10 orbits, compared to $\langle P \rangle \sim 1.4$ after the same amount of time when initialized with a vertical magnetic field of radial wavenumber $k_x = 8\pi/L_x = 2\pi$. Similarly if the initial velocity perturbations δu were too small (we tried $\delta u < 0.04c_{s0}$) non-linear saturation would set in later and at a higher peak stress, resulting in the box heating up significantly before the end of the linear phase.

We have also investigated the effects of different initial conditions for k_x and δu . Several runs were carried out without explicit diffusion coefficients to investigate the effects of grid diffusion. All simulations described in this section were run at a resolution of 256^3 (64 grid cells per scale height).

In the left- and right-hand panels of Figure 5.1 we show the time-evolution of the volume-averaged pressure $\langle P \rangle$ and of the total stress $\langle \Pi_{xy} \rangle$, respectively, from our fiducial simulations. The pressure increases monotonically by a factor of around 7 over 40 orbits. The stress exhibits more complicated behavior: it fluctuates over short timescales (< 0.1 orbit), but over the du-




Figure 5.1: Left: time-evolution of pressure from our 256^3 simulation without cooling and with explicit diffusion coefficients. Right: time-evolution of the total stress (Reynolds + magnetic) from the same simulation, normalized by the pressure at initialization $P_0 = 1$.

ration of the simulation the stress exhibits three distinct kinds of behavior, similar to what was reported in Ross et al. (2016). During the first phase, corresponding to the linear MRI, the stress increases exponentially, reaching a peak value at non-linear saturation of around 0.015. The stress drops briefly after non-linear saturation, but begins to increase fairly rapidly from about orbit 6 onwards. Between orbits 6 and 33 we record a five-fold increase in the stress, which is consistent with the increase in stress over the 60 orbits following non-linear saturation reported by Ross et al. (2016). After orbit 33 we record a large jump in the stress (from about 0.05 to 0.065), after which the stress remains quasi-steady in time, though undergoing fairly large fluctuations of around 0.01 in magnitude. We interpret this as the simulation having reached the *box-dominated* regime. The pressure just before the simulation entered the box-dominated regime was $\langle P \rangle \sim 4.8$.

In Figure 5.2 we plot the volume-averaged total stress as a function of the pressure. This plot should be compared to the red curve in Figure 6 of Ross et al. (2016) which shows the equivalent plot from their 256^3 with explicit diffusion coefficients. Between the initial spike in stress and the box-dominated regime where the stress plateaus as the pressure continues




Figure 5.2: Log-log plot of total stress against pressure from our 256^3 simulation without cooling and with explicit diffusion coefficients. From left to right, the dashed lines have slopes of $q \approx 2.4$ and $q \approx 1.45$, respectively.

to increase, we find two phases where the pressure has a distinct power law dependence on the pressure. The stress rises more steeply with pressure during the first phase (which lasts from around orbit 5 to orbit 15) and we measure a power law index of $q \approx 2.4$. During the second phase (lasting from around orbit 15 to orbit 33) the stress rises less steeply with pressure compared to the first phase and we measure a power law index of $q \approx 1.45$. Ross et al. (2016) also observe two different scaling laws before the simulation enters the box-dominated regime, though the first is very steep and short and was not fitted. For the second, gentler, rise in stress they measure a slope of $q \approx 0.9$ in their simulation with explicit diffusivities. In any case we caution that not too much weight should be conferred on the exact value of the slope measured either by us or by Ross et al. (2016). In both cases the slope was measured from a single run in which strong random fluctuations are bound to obscure clean power law behavior.




Figure 5.3: Time-evolution of total stress multiplied by 100 (in black) and total pressure (in blue) from a 128^3 ideal simulation with ‘intermediate’ cooling timescales for the heating and cooling phases of $\tau_H = 30$ orbits and $\tau_C = 7.5$ orbits, respectively. A rolling average of the total stress is shown in red. The horizontal dashed black lines indicate the maximum and minimum values that volume-averaged pressure is allowed to reach before the cooling prescription switches to a shorter (or longer) timescale. The vertical dashed red line indicates the restart position. Note the clear lag in stress compared to the pressure starting from around orbit 75.

5.3.2 Runs with cooling

In the previous section we investigated the relationship between stress and pressure by allowing the box to heat-up without any cooling. Thus in those runs, thermalization of MRI turbulence resulted in a monotonic pressure increase which the stress then tracked, until the simulation reached the box-dominated regime. In this section we wish to investigate how the stress responds to variations in the mean pressure in greater detail. To do so we run our simulations with the explicit cooling prescription described in Section 5.2.1. By adjusting the two cooling timescales τ_H and τ_C we can control




Figure 5.4: Same as Figure 5.3 but for a 256^3 simulation with explicit diffusion coefficients. The heating phase and cooling phase cooling timescales were set to $\tau_H = 30$ orbits and $\tau_C = 4.5$ orbits, respectively.

the behavior of the mean pressure, allowing it to increase when the cooling timescale τ_H is sufficiently long so that heating is greater than cooling (the ‘heating phase’), and forcing the pressure to decrease when τ_C is sufficiently short so that cooling overwhelms heating (the ‘cooling phase’). We discuss three different cases: a simulation in which the mean pressure variation is ‘fast’ such that the stress cannot keep up with the changes in the pressure, a simulation in which the mean pressure variation is ‘slow’ and in which the stress very closely tracks the pressure, and an ‘intermediate’ case in which the stress lags behind the pressure.

In Figure 5.3 we show in black the time-evolution of the stress normalized by the mid-plane pressure at initialization (where $P_0 = 1$) and pressure (in blue) from our intermediate case simulation. The total stress has been multiplied by 100 to compare more easily to the pressure. The long and short cooling timescales were set to $\tau_H = 30$ orbits and $\tau_C = 7.5$ orbits, respectively, and the pressure was permitted to vary between $\langle P \rangle_- = 1$ and

$\langle P \rangle_+ = 2$, thus ensuring that the box never heated up to the box-dominated regime. To better track the dips and increases in the stress we also carry out a rolling-average of the total stress using a width of 10^4 steps, or about 7 orbits. This simulation was restarted from an MRI turbulent state shortly after non-linear saturation from a simulation that had previously been run without any cooling. At the start of the simulation (orbit 4.7) the pressure continues to increase, though at a slower rate than before due to the explicit cooling. At orbit 19.3 the pressure reaches its maximum permitted value at which point the cooling timescale is switched to the shorter timescale τ_C and the pressure begins to drop. The decrease is not monotonic however, and shortly before orbit 50 the pressure plateaus before reaching its minimum value, and even begins to increase again due to the increasing stress. In fact until orbit 75 the stress does not appear to be responding very strongly to the pressure at all. After the first pressure cycle is completed, however, the stress begins to respond much more sensitively to the pressure, and for the remainder of the simulation the stress can be seen to track the mean pressure, albeit with a lag of several orbits.

In Figure 5.4 we show the time-evolution of stress and pressure from a 256^3 simulation with explicit diffusion coefficients. This was restarted with cooling turned on from shortly after non-linear saturation of the 256^3 run (which did not employ cooling) discussed in Section 5.3.1. The long cooling time is the same as in the 128^3 ‘intermediate’ run discussed in the last paragraph. When we attempted to run this simulation using the same short cooling timescale as was employed in the 128^3 simulation ($\tau_C = 7.5$ orbits), however, we found that the cooling was frequently overwhelmed by the turbulent heating, so that the mean pressure would never quite drop to its minimum value when the short cooling timescale was used, but would instead flatten out and begin to respond to bursts in the stress rather than vice versa. Thus we were forced to lower the short cooling timescale to $\tau_C = 4.5$ orbits in our 256^3 run in order to get the mean pressure to decrease more or less monotonically during the cooling phases. As in the 128^3 case over the first 60 orbits or so there is very little correlation between the stress and the pressure, and again after orbit 190, with the volume-averaged stress plateauing over the next cycle-




Figure 5.5: Cross-correlation coefficient between stress and pressure from a 128^3 inviscid simulation with ‘intermediate’ cooling timescales of $\tau_H = 30$ orbits and $\tau_C = 7.5$ orbits. To the right of the black vertical dashed line (at 0 orbits on the x -axis) is the correlation for a lag (in orbits) in which the stress follows the pressure, while to the left of the line the correlation is for the stress leading the pressure. The peak correlation (red vertical dashed line) occurs when the stress lags behind the pressure by about 3.34 orbits.

and-a-half of the mean pressure variation. Between orbit 60 and orbit 190, however, the stress does begin to respond to the pressure and we observe a clear correlation between the two with the stress lagging behind the pressure.

In our heating-only run we found that the simulation entered the box-dominated regime once the stress reached around 0.05-0.065 (cf. the right-hand panel of Figure 5.1) after which the stress plateaued. Note that in both our 128^3 and 256^3 cooling runs the total stress is generally below this range, reaching values above $\Pi_{xy} \sim 0.05$ only at the very tips of some of the outbursts in stress. Thus we are confident that our results in these cooling runs are not being affected by the box size.




Figure 5.6: Same as Figure 5.5 but for a 256^3 simulation with explicit diffusion coefficients. The peak correlation (red vertical dashed line) occurs when the stress lags behind the pressure by about 5.61 orbits.

To quantify the lag between the stress and the pressure we calculated the cross correlation between the two, which is shown in Figure 5.5 for our 128^3 simulation. The correlation was calculated using data between orbits 56 and 400 in order to omit residual and transient behavior following initialization. We measure a peak correlation coefficient of $r = 0.45$. This peak occurs when the stress lags behind the pressure by around 3.3 orbits. For comparison, the cross correlation for the 256^3 simulation is shown in Figure 5.6 (in this run the data used to calculate the cross correlation was taken between orbits 66 and 193). In this case we measure a peak correlation coefficient of $r = 0.55$, with the peak occurring when the stress lags behind the pressure by around 5.6 orbits. Note that in both simulations the correlation coefficient curve is quite wide, and there is significant overlap between the two curves.

Next we investigate how the stress tracks the pressure when the latter




Figure 5.7: Time-evolution of total stress multiplied by 100 (in black), rolling average of total stress (in red), and total pressure (in blue) from a 128^3 ideal simulation with ‘fast’ cooling timescales of $\tau_H = 30$ orbits and $\tau_C = 4$ orbits. The simulation was restarted (blue vertical dashed line) at orbit 80 from the solution of our ‘intermediate run’ (where we had used $\tau_C = 7.5$ orbits). Note that after the restart the stress is unable to keep up with the rapid variation in the mean pressure.

changes over timescales that are much shorter and much longer than those discussed in the ‘intermediate’ case; we remind the reader that we refer to these simulations as our ‘fast’ and ‘slow’ runs, respectively. To do this we again vary the cooling timescale parameters τ_H and τ_C . To instigate rapid variations in the pressure we require that the pressure increase quickly during the heating-dominated phase (during which the cooling timescale is set by τ_H), and that the pressure decrease quickly during the cooling-dominated phase (during which the cooling timescale is set by τ_C). This can be facilitated by increasing τ_H and decreasing τ_C compared to the intermediate case. Thus in our ‘fast’ simulation we set the cooling timescales to $\tau_H = 30$ orbits




Figure 5.8: Time-evolution of total stress multiplied by 100 (in black), rolling average of total stress (in red), and total pressure (in blue) from a 128^3 inviscid simulation with ‘slow’ cooling timescales of $\tau_H = 20$ orbits and $\tau_C = 10$ orbits. The simulation was restarted (blue vertical dashed line) at orbit 80 from the solution of our ‘intermediate run’ (were we had used $\tau_C = 7.5$ orbits). Note that after the restart the stress closely tracks the gradual variation in the mean pressure.

and $\tau_C = 4$ orbits. This simulation is restarted from orbit 80 of our ‘intermediate’ simulation in order to avoid the initial phase in which the stress does not track the pressure, and was run until orbit 200.

In Figure 5.7 we show the time-evolution of the total stress multiplied by 100, the rolling-averaged total stress, and the pressure from this ‘fast’ run. After the restart position (orbit 80) the mean pressure varies rapidly due to the shorter cooling timescale τ_C . The stress struggles to keep up with the pressure and for the first 2.5 cycles of the pressure there is no correlation at all between the two quantities. After this initial phase, the total stress does occasionally respond weakly to the behavior of the pressure, e.g. there is a




Figure 5.9: Cross-correlation coefficient between stress and pressure from a 128^3 inviscid simulation with ‘slow’ cooling timescales of $\tau_H = 20$ orbits and $\tau_C = 10$ orbits. The peak correlation (red vertical dashed line) occurs when the pressure lags behind the stress by about 0.87 orbits.

rise in the total stress at around orbit 125 that coincides closely with the pressure, but overall the lag is much less evident than in the intermediate cooling timescales case (see Figure 5.3 for comparison).

Finally in our ‘slow’ simulation we look at the other extreme: when the variation in the mean pressure is sufficiently slow that the stress tracks the pressure almost instantly with respect to the timescale of the oscillation (see Figure 5.8). This was achieved by lowering the ‘long’ cooling timescale to $\tau_H = 20$ orbits (resulting in longer rise times of $\langle P \rangle$ during the heating-dominated phase) and increasing the ‘short’ cooling timescale to $\tau_C = 10$ orbits (thus prolonging the decrease of $\langle P \rangle$ during the cooling-dominated phase). In order to obtain a clearer picture of the slowly varying mean pressure we have run this simulation for 420 orbits. We find that the stress closely tracks the pressure after initialization (orbit 80). On short timescales

(< 1 orbit) the stress can even react back on the pressure, and thus appears to lead the pressure slightly. This is confirmed in Figure 5.9 where we plot the cross correlation of stress and pressure, and find that the pressure lags behind the stress by around 0.87 orbits.

5.4 Discussion and Conclusions

We have carried out unstratified MHD shearing box simulations in PLUTO both with and without explicit cooling in order to investigate the relationship between total stress and pressure, with an eye towards applying these results to investigating thermal instability and viscous overstability in future work.

In 256^3 simulations with explicit diffusion coefficients but without cooling, in which the box heats up due to the dissipation of MRI turbulence, we observe two separate power law scaling regimes between stress and pressure, i.e. $\Pi_{xy} \sim P^q$. Shortly after non-linear saturation the stress increases relatively steeply with respect to pressure with $q \approx 2.4$, though we suspect this might be related to transient behavior following the end of the linear phase. More significantly, on longer timescales we find a separate, shallower, dependence of $q \approx 1.45$. This power law behavior continues until the total stress reaches a value of $\Pi_{xy} \approx 0.05 - 0.07$ at which point it plateaus even while the pressure continues to increase, and the simulation enters a box-dominated regime in which the characteristic length scale of turbulent eddies exceeds the vertical box size. These results are broadly consistent with those obtained in similar simulations by Ross et al. (2016), though the exact scaling relationship is likely to change between individual runs as a result of random fluctuations. Ideally we would run several of these simulations and average the results, which we expect would result in a cleaner scaling relation. We plan on doing so in future work.

In our cooling runs we implemented a time-dependent cooling prescription which enabled us to control the mean pressure, and thus measure the resultant response of the stress. Depending on the cooling timescales we chose we found three possible outcomes. For runs in which the cooling timescales

switched between ‘intermediate’ values of $\tau_C = 30$ and $\tau_C = 7.5$ orbits we found a clear lag in stress in both 128^3 ideal and 256^3 non-ideal runs. We quantified this lag by calculating the cross correlation between the volume-averaged stress and the pressure, and found that the peak correlation occurred for a lag of around 3 orbits in the 128^3 run. This is slightly shorter than the lag of around 5 orbits measured in our 256^3 run, although the peaks in correlation are broad in both cases and overlap significantly. For shorter cooling timescales ($\tau_C = 4$ orbits) the stress is unable to keep up with variations in the mean pressure and evolves more or less independently; conversely in a run in which the mean pressure varied only gradually ($\tau_C = 20$ and $\tau_C = 10$ orbits, respectively) the stress was observed to closely track the pressure. In the latter case, fluctuations in the stress were found to react back on the pressure (via dissipative heating). Thus on very short timescales (< 1 orbit) it is the pressure that responds to the stress rather than vice versa.

The main application of these results is to thermal instability, particularly in LMXBs. Earlier work on the linear theory of thermal instability has found that the growth rate of the instability is significantly decreased when there is a lag between the stress and the pressure (Lin et al., 2011). However in those calculations the lag was included as a free parameter, while in our simulations we have varied only the mean pressure in a quasi-periodic manner. Our main result is that, for variations on reasonable timescales, the stress lags behind the pressure by 3-5 orbits. An important caveat to this, however, is that LMXBs are radiation-dominated (at least in their inner regions), whereas our simulations were run with gas pressure only. Nevertheless, our results are indicative that, at least for variations of the mean pressure on timescales over several tens of orbits, thermal instability should be weakened, though perhaps not suppressed entirely.

Finally we note that the response of the stress to the pressure is only one aspect of the turbulent dynamics. More generally, we wish to investigate the dependence of the turbulent stress on the strain. If, as we suspect given the results of this chapter, the former were to depend non-linearly and non-instantaneously on the latter, this would indicate that turbulent stresses

behave in a non-Newtonian fashion as has already been investigated analytically by Ogilvie (2003). Furthermore it would be interesting to determine if there is a time-lag between stress and strain, something that would be particularly applicable to the viscous overstability. If the time-lag between stress and strain is similar to the lag between stress and pressure of a few orbits that we've found in our simulations, we expect the viscous overstability, which has an oscillation period of one orbit, will be suppressed. We plan on exploring these avenues in future work.

Chapter 6

Conclusions

In this thesis we have studied convection in accretion disks, in particular the interaction between convection and the magnetorotational instability, with an emphasis on its consequences for vertical disk structure and angular momentum transport. More generally we have investigated the dynamics of turbulent flows in disks, including how turbulent stresses respond to modulations of the mean pressure.

Convection has been discussed in the context of accretion disks for many years, almost to the beginning of the ‘modern’ era of accretion disk theory (cf. Shakura & Sunyaev (1973); Lynden-Bell & Pringle (1974), and references therein). Identifying the source of the effective viscosity needed to drive accretion in disks was recognized as an outstanding problem early on. Paczynski (1976) and others, almost certainly taking inspiration from the more established theory on the structure of stars, in which convection was thought to play an important role in transporting heat, suggested that convective turbulence might be the source of this viscosity. Since then there has been significant, if intermittent, debate on the role of convection in disks. The discussion has included the relevance of convection to the hysteresis behavior of dwarf novae, whether convection in disks is self-sustaining (in the sense that the viscous dissipation of orbital energy due to turbulent convection is itself sufficient to drive convection), whether convection leads to inward or outward angular momentum transport, and, more recently, whether its in-

teractions with the magnetorotational instability can enhance the transport of angular momentum.

To isolate the generic features of convection in disks, in Chapter 3 we have first undertaken its study in purely hydrodynamic models, first by probing the linear regime both analytically and semi-analytically, and then in fully compressible numerical simulations carried out with the finite-volume code **PLUTO**. Our key result is that we find that hydrodynamic convection can, in general, transport angular momentum outwards in disks, in other words $\alpha > 0$. We have verified this result using **ATHENA**, an alternative finite volume code. This is in agreement with the findings of Lesur & Ogilvie (2010), who found in their Boussinesq simulations that hydrodynamic convection could transport angular momentum outwards at sufficiently high Rayleigh numbers, but is in disagreement with the fully compressible hydro calculations (nearly identical to ours but run with the finite difference code **ZEUS**) carried out by Stone & Balbus (1996) (SB96). Crucial to understanding this discrepancy is that SB96 reported seeing nearly axisymmetric convection in their simulations, which, as they showed analytically, must lead to inward transport. We showed that the diffusivity of the numerical scheme, in particular the use of diffusive Riemann solvers such as HLL, can impose an artificial axisymmetry on the flow field that in turn leads to inward transport.

This issue is likely independent of resolution. In our linear stability analysis of convection we showed that inviscid linear convective modes grow fastest on the shortest possible scales, and will thus lie on the grid in inviscid simulations regardless of the resolution. We ran a suite of simulations with explicit viscosity and thermal diffusivity, which, in addition to mitigating the effects of the grid, allowed us to control the Rayleigh number. We found that the onset of convection occurred above some critical Rayleigh number $10^5 < \text{Ra}_c \leq 10^6$. For Rayleigh numbers just above this critical value we observed inward transport and the flow was largely axisymmetric. For higher Rayleigh numbers ($10^6 < \text{Ra} < 10^9$), we first observed inward transport during the linear phase (during which axisymmetric modes were growing most quickly), but these were eventually attacked by secondary instabilities and a non-axisymmetric flow field, resulting in an outward transport of angular

momentum during the non-linear phase. Finally we have also carried out simulations in which we maintained convection by imposing an artificial heat source. In the non-linear phase convection is manifested as large-scale ($\sim H_0$ in width) cells (perpendicular to the plane of the disk) that are cyclically created and destroyed. These are associated with other coherent structures (in the plane of the disk) such as zonal flows and small-scale ($\sim 0.1H_0$) vortices.

Another important result is that we have not observed self-sustaining convection in any of our simulations, indicating that to drive convection self-consistently in disks the heat source has to come from other processes, such as the dissipation of MRI turbulence. This motivated us to study the interplay between the two instabilities through fully compressible MHD simulations in Chapter 4. To determine the effects of the MRI alone on vertical disk structure and convective instability, we first carried out ideal MHD simulations with open boundaries in the vertical direction, thus enabling the box to cool by advection of thermal energy across the vertical boundaries, but otherwise without an explicit cooling prescription. Here our main result was that MRI heating and box-cooling alone results in the disk settling into a convectively stable equilibrium. Furthermore, MRI turbulence itself is rather adept at moving heat vertically, thus hampering the build-up of heat at the mid-plane. When an explicit cooling prescription was used (with the cooling timescale uniform in space), we recovered the main results of the box-cooled simulation.

Thus in our box-cooled and uniformly cooled simulations we found that the square of the buoyancy frequency was positive everywhere, which rules out the possibility of convection. The onset of convection in a laminar fluid with a microscopic viscosity and thermal diffusivity will occur only if *both* the squared buoyancy frequency is negative *and* the Rayleigh number exceeds some critical value. Likewise, in a turbulent fluid, such as one stirred up by the MRI, there will be a turbulent transport of heat and momentum, which, assuming these behave in a diffusive manner, will result in the fluid being characterized by an effective Rayleigh number. Thus in our MHD simulations the sign of the squared buoyancy frequency is only a necessary, not a sufficient condition, for determining the onset of convection. The latter requires both

a negative squared buoyancy frequency (i.e. an unstable entropy gradient) as well as that the effective Rayleigh number exceed some critical value. This is an important point, which has not been elucidated by previous work investigating convection and the MRI in disks. A corollary of this statement is that the two instabilities might actually not coexist, or at least do not interact in an additive manner, as has been implied in previous work. If the effective Rayleigh number is too low, the MRI is bound to disrupt convective modes before they grow.

The two instabilities can interact in a non-trivial manner, however. We have investigated this by running fully compressible MHD simulations with both *height*-dependent cooling and explicit resistivity. Controlling the cooling timescale enabled us to control the entropy gradient, and by controlling the resistivity we could control the strength of the MRI (relative to ideal MHD simulations) and thus the effective viscosity and thermal diffusivity. We found that at higher magnetic Reynolds numbers ($Rm \geq 4 \times 10^3$) the MRI dominated the flow field, with very few signs of convection despite the negativity of the buoyancy frequency. At lower magnetic Reynolds numbers ($Rm \lesssim 2 \times 10^3$), on the other hand, we found convective/MRI cycles, in which resistivity would first quench the MRI, residual MRI heating would drive convection, and convection could then reseed the MRI. The stresses, pressure, and alpha during the convection-dominated phases were low, but during the MRI outbursts we found that alpha could reach as high as $\alpha \sim 0.08$. Furthermore, during the convection-dominated phases we recovered the large-scale cyclical convective cells that we had observed in our forced hydro simulations, demonstrating that these structures are robust when the heat source for convection is provided in a self-consistent manner, in this case by residual heat from the previous MRI burst.

At the lowest magnetic Reynolds number we investigated ($Rm \sim 200$) we find that resistivity can quench the MRI entirely in the absence of convection. At lower resistivities (higher Rm) the MRI is dampened, rather than quenched, by the resistivity, and in fact we found (in a simulation without explicit cooling) low amplitude cycles even in the absence of convection. Convection, however, can enhance the amplitude of these cycles. It is tempting

to link our results in these more marginal cases to the resistivity driven state changes reported in the isothermal simulations of Simon et al. (2011), but the timescales on which they observed cycles (at similar magnetic Reynolds numbers) are of the order hundreds of orbits, much longer than the tens of orbits or so over which we observe cycles, thus indicating that the two phenomena are likely of different origin.

Finally in Chapter 5 we have investigated the relationship between stress and pressure in unstratified MHD simulations. By implementing a piecewise-in-time linear cooling prescription we were able to control the variation in the mean pressure and thus follow the response of the stress. We found three different outcomes, depending on how rapid the variations in the mean pressure were. Our key result in this chapter is that, for moderate cooling timescales (i.e. such that the mean pressure oscillated with a period of 25-30 orbits), we found that the stress exhibited a clear lag behind the pressure. We quantified this result by calculating the cross-correlation between total stress and pressure. The peak correlation occurs when the stress lags the pressure by around 3 orbits. At double the resolution (256^3 compared to 128^3), the lag is somewhat longer (about 5 orbits), but the peaks of the cross correlation are broad at both resolutions, and there is significant overlap between the two runs. We also investigated the effect of varying the pressure on much shorter timescales (~ 10 orbits). In this case the stress could not keep up with variations in the pressure at all, and the two evolved independently. Finally we explored the opposite regime, i.e. in which the pressure varied only very slowly (i.e. on timescales of around 150 orbits). In this regime the stress not only tracks the pressure very closely (with almost no discernible lag), but it can even react back on the pressure (via dissipative heating). Thus we have found that while stress appears to lag behind pressure on longer timescales, on short timescales ($\lesssim 1$ orbit) it is the pressure that follows the stress. These results are bound to have an impact on thermal instability in LMXBs, and on the viscous overstability, potentially dampening or even mitigating each instability entirely. We plan on exploring these applications in future work.

6.1 Future work

Although we have investigated many aspects of the problem of hydrodynamic convection (linear/non-linear, viscid/inviscid, different numerical schemes, forced and unforced), there remain several avenues for future work. A key result was that, given a sufficiently non-dissipative numerical scheme (in inviscid simulations) or a large enough Rayleigh number (in viscid simulations), hydro convection can transport angular momentum outwards. We interpreted this as meaning that, during the non-linear phase in particular, non-axisymmetric convective modes dominated the flow. Given that our investigations were restricted to the local approximation, it is difficult to get a scope of the global structure of these modes. Thus a natural extension of our work would be to carry out global three-dimensional hydrodynamic simulations. Such simulations have actually been carried out long ago by Klahr et al. (1999), who reported observing non-axisymmetric convective cells in their runs. However, their simulations included an explicit α , and thus the turbulent transport by convection was not modeled in a self-consistent manner. Furthermore, they employed unrealistic reflective boundary conditions at the inner and outer radial disk boundaries. Therefore it would be fruitful to carry out similar global simulations but with a more modern code, more realistic boundary conditions, and employing, for example, the relaxation cooling prescription that we used in our simulations of forced convection in order to sustain convection.

In our forced hydrodynamic simulations we put in the heat source by hand, i.e. by relaxing the thermal equilibrium back to a convectively unstable equilibrium. As we have discussed in Section 1.4.3 there are several realistic heat sources that could drive convection in disks. One of them, the MRI, was investigated in Chapter 4, but there are also purely hydrodynamic heat sources that could drive convection in disks, such as the dissipation of density waves launched by a high mass planet. Previous work on this topic by Lyra et al. (2016) could be extended to include (i) more realistic thermodynamics, (ii) irradiation from the central star, and (iii) dust, with the aim of investigating the interaction of dust and convection and its impact on dust

settling. Details of shock heating and convection could also be investigated using high-resolution simulations in a wide shearing box with an embedded planet, extending previous local simulations (which did not include dust) by Zhu et al. (2016). Finally another potential heat source whose effect on vertical disk structure has not been investigated in great detail is Joule dissipation in mid-plane current sheets, the latter of which have been reported in global simulations of protoplanetary disks by Béthune et al. (2017).

With regards to our work on the MRI and convection, we have found that narrow boxes ($L_x = H_0$) have several limitations and numerical artifacts, leading to bursty behavior, and also suppressing large scale convective cells. It would certainly be worth carrying out radiative transfer MHD simulations similar to those of Hirose et al. (2014) and Scepi et al. (2018a) in wider boxes. Given adequate numerical resources, these runs should be relatively straightforward to carry out.

Another aspect of the problem that we have not touched on in this thesis would be to investigate the effect of our results on the MRI dynamo, which would allow for closer comparison between our work and that of Coleman et al. (2016) and Coleman et al. (2018), who investigated the effect of intermittent and persistent convection on the MRI dynamo in dwarf novae and AM CVn disks, respectively. We already have the spacetime data for the evolution of the magnetic field from our simulations, but the complexity of the problem warrants its own separate investigation to the one we have carried out Chapter 4, where our focus was on understanding how the interplay of convection and the MRI affected the dynamics of the flow rather than the dynamo in particular.

Finally, another aspect of convection in dwarf novae that has not been explored yet is the role of unstable compositional gradients. Although in the introduction we discussed the linear theory of convection driven by *both* entropy and compositional gradients, ultimately all our results were obtained assuming a constant mean molecular weight. In actual dwarf novae, however, there is a transition from a fully ionized mid-plane to a partially ionized atmosphere, and thus we expect the mean molecular weight to increase with vertical distance from the mid-plane. This could potentially set up an un-

stable compositional gradient, in addition to the unstable entropy gradient which is maintained by dissipation of MRI turbulence and radiative cooling. Indeed there is some preliminary evidence in the data from the radiative transfer simulations of Scepi et al. (2018a) that this is the case (Nicolas Scepi; private communication). Thus it would be worth while investigating the role of both entropy and compositional gradients in dwarf novae disks, something that we plan on doing over the course of the next few years.

Appendix A

Numerical issues and convergence studies for MHD

A.1 Numerical issues with $\nabla \cdot \mathbf{B} = 0$ in PLUTO

In this section we briefly report a subtle numerical issue in which a particular numerical algorithm related to the implementation of $\nabla \cdot \mathbf{B} = 0$ in PLUTO can severely mitigate the MRI. This has important implications for studying the interplay between the MRI and another instability (such as convection or gravitational instability), because it is crucial to ascertain whether any enhancement (or mitigation) of MRI turbulence is due to physical or numerical effects.

The Constrained Transport (CT) algorithm is commonly used to enforce $\nabla \cdot \mathbf{B} = 0$ to machine accuracy. Finite-volume codes (such as PLUTO) update cell *centered* fluid variables (which are taken to be volume-averages over the cell). The CT algorithm, on the other hand, works by calculating the electromotive force (emf) at cell *edges*, and then updating the magnetic field (by integrating the induction equation) at cell *faces* (referred to as a ‘staggered’ magnetic field). Cancellation of terms then guarantees $\nabla \cdot \mathbf{B} = 0$. Various averaging procedures are available in PLUTO for calculating the emf at the zone edges, among them UCT-HLL, which employs the HLL Riemann solver. We have discovered that more diffusive averaging procedures




Figure A.1: Time-evolution of volume-averaged magnetic energy (left), and stresses (right) from PLUTO simulations with different algorithms for enforcing $\nabla \cdot \mathbf{B} = 0$: UCT-Contact (solid lines), UCT-HLL (dashed lines). Top row: simulations run with a *non-isothermal* equation of state (box size: $H_0 \times 4H_0 \times 6H_0$, resolution: 32 cells per H_0). Bottom row: simulations run with an *isothermal* equation of state (box size: $2H_0 \times 4H_0 \times 6H_0$; resolution 32 cells / H_0). In all cases, the MRI is mitigated when the UCT-HLL solver is used. The vertical dashed lines indicate the times at we switched from using UCT-Contact to UCT-HLL.

such as UCT-HLL can severely mitigate the MRI.¹ Less diffusive averaging procedures such as UCT-Contact (which restores the contact wave removed by the HLL solver) do not exhibit this effect.²

¹Constrained Transport is also available in ATHENA++ (which is also a finite volume code), but the procedure for calculating the electric field is different to that employed in PLUTO, and does not seem to exhibit the problems the UCT-HLL algorithm does (Jim Stone; private communication).

²The reader should take care to distinguish between the overall Riemann solver used in the code (in our case we always use an HLLD solver), and the Riemann solver employed *only* in calculating the emf at cell edges during the implementation of the Constrained Transport algorithm (in this section either UCT-HLL or UCT-Contact).

We have investigated the dependence of the solution on the averaging-procedure used in the CT algorithm for runs with both an isothermal and a non-isothermal perfect gas equation of state (see Figure A.1).³ The non-isothermal simulations (top row) were run in boxes of size $H_0 \times 4H_0 \times 6H_0$ and at a resolution of 32 cells/ H_0 for a duration of 100 orbits ($629 \Omega^{-1}$). The adiabatic index was set to $\gamma = 5/3$. The isothermal simulations (bottom row) were run in wider boxes size of $2H_0 \times 4H_0 \times 6H_0$ and a resolution of 32 cells/ H_0 for a longer duration of 200 orbits ($1257 \Omega^{-1}$). All simulations were run with a zero net magnetic flux configuration at initialization, included the vertical component of gravity, and employed outflow boundary conditions in the vertical direction (see Section 4.1.1), as well as a mass source term. Explicit cooling was not employed in any of the simulations, i.e. cooling was facilitated only by advection of material across the vertical boundaries. Finally, the UCT-HLL simulations were restarted from the MRI turbulent state of the UCT-Contact simulations from orbit 54 (in the non-isothermal case) and orbit 140 (in the isothermal case).

In Figure A.1, solutions obtained using UCT-Contact are shown with a solid line, and the solutions obtained using UCT-HLL are shown with a dashed line. For both an isothermal and non-isothermal gas equation of state the switch from UCT-Contact to UCT-HLL is followed by a rapid drop in the kinetic and magnetic densities (left-hand panels), and in the stresses (right-hand panels). Furthermore the flow field (not shown) is mostly laminar, i.e. absent of any MRI turbulence, once the UCT-HLL algorithm is used. We conclude that the UCT-HLL algorithm severely mitigates (or even kills) the MRI. Therefore we employ the UCT-Contact algorithm for Constrained Transport in all simulations discussed in the main part of our paper.




Figure A.2: Plot of time-evolution of volume-averaged α for vertically stratified MRI simulations (without explicit cooling or diffusion coefficients) at three different resolutions: $64/H_0$ (blue), $32/H_0$ (red), and $16/H_0$ (black). All three simulations were run in a box of size $4H_0 \times 4H_0 \times 6H_0$.

A.2 MHD convergence studies

A.2.1 Dependence on resolution

Ryan et al. (2017) carried out a convergence study of the inviscid ZNF vertically stratified MRI in isothermal shearing box simulations, and determined that these were not converged with resolution, in agreement with earlier work done on the convergence of the unstratified ZNF MRI (Fromang & Papaloizou, 2007). It is not clear, however, whether the non-convergence of the MRI carries over to simulations with more realistic thermodynamics. Hirose et al. (2014), Coleman et al. (2018), and Scepi et al. (2018a) did not perform convergence studies in their vertically stratified MRI simulations with radiative transfer and opacity tables, presumably due to the large computational

³The equation of state in the isothermal runs was given by $P = c_s^2 \rho$ (with $c_s = 1$), while in the non-isothermal simulations the equation of state was given by $P = (\gamma - 1)\rho e$.

cost of running these simulations. To test the convergence with resolution of the vertically stratified MRI in inviscid simulations with a *perfect gas* equation of state, we have run our fiducial box-cooled simulation (see Section 4.2) at 64, 32, and 16 cells per initial scale height H_0 (all three simulations were run in a box of size $4H_0 \times 4H_0 \times 6H_0$). The simulations, together with time-averaged values for various diagnostics are listed in Table C.1 in Appendix C.

We plot the time-evolution of $\langle\alpha\rangle$ over 100 orbits ($629\Omega^{-1}$) from the three simulations in Figure A.2. The results are generally very similar between the $64/H_0$ and $32/H_0$ simulations, not just in the time-evolution of α but across all volume-averaged quantities (vertical kinetic energy density, magnetic energy density, stresses, etc.). Nor could we detect any appreciable difference in the vertical structure of the disk between the two simulations: the mass and heat fluxes are directed outwards and are of comparable magnitude, and the temperature profile is flat-topped around the mid-plane, dropping rapidly beyond $z \sim \pm 2H_0$. The $16/H_0$ simulation, on the other hand, is noticeably unconverged, with $\langle\langle\alpha\rangle\rangle \sim 0.005$ in this simulation compared to $\langle\langle\alpha\rangle\rangle \sim 0.014$ and $\langle\langle\alpha\rangle\rangle \sim 0.013$ in the $32/H_0$ and $64/H_0$ simulations, respectively.

A.2.2 Dependence on radial box size

In Section 4.2.2 we compared the results of our fiducial box-cooled simulation (run in a box of radial size $L_x = 4H_0$) with those of simulation MRI-S1 in Riols & Latter (2018) which was run with a narrower box of size $L_x = 2H_0$. We commented on the fact that the stresses in MRI-S1 appeared significantly more bursty compared to those measured in our fiducial simulation. In this section we investigate the behavior of the MRI in ZNF vertically stratified boxes (without explicit cooling or explicit dissipation coefficients) as a function of the radial box size. In Table C.2 in Appendix C we list the results of three simulations, differing in radial box size, but otherwise identical. Out-flow boundary conditions were employed in the vertical direction in all three simulations (see Section 4.1.1). The number of cells per scale height is the




Figure A.3: Plot of time-evolution of volume-averaged α for vertically stratified MRI simulations (without explicit cooling or diffusion coefficients) at three different radial box sizes: $L_x = H_0$ (blue), $L_x = 2H_0$ (red), and $L_x = 4H_0$ (black). All three simulations were run with a resolution of 32 cells / H_0 in each direction.

same in all three simulations at $32/H_0$.

In Figure A.3 we plot the time-evolution over 200 orbits ($1257\Omega^{-1}$) of $\langle \alpha \rangle$ for the three different radial box sizes. The time-averaged value of α *increases* as the box radial box size *decreases*: $\langle \langle \alpha \rangle \rangle \sim 0.014, 0.016$, and 0.026 in the $L_x = 4H_0, 2H_0$, and H_0 simulations, respectively. In addition visual inspection of the solutions in Figure A.3 shows that fluctuations in $\langle \alpha \rangle$ also increase markedly as the radial box size is decreased. The difference is particularly noticeable when comparing the $L_x = H_0$ simulation with the other two simulations. This simulation undergoes an outburst between around orbits 80 and 140 in which $\langle \alpha \rangle$ reaches as high as ~ 0.06 , a factor of three larger than the maximum values of $\langle \alpha \rangle \sim 0.03$ and $\langle \alpha \rangle \sim 0.02$ measured over the course of the $L_x = 2H_0$ and $L_x = 4H_0$ simulations, respectively. This appears to confirm our suspicion that narrower boxes are

more ‘bursty’, though the difference is more apparent when comparing the narrowest box size ($L_x = H_0$) to the wider boxes than when comparing the $L_x = 2H_0$ and $L_x = 4H_0$ runs. These results are consistent with those of Simon et al. (2012) who carried out vertically stratified shearing box simulations initialized with a twisted azimuthal magnetic flux tube in boxes of radial size between $L_x = 0.5H_0$ and $L_x = 16H_0$. They reported that turbulent properties such as the volume-averaged α viscosity were only converged for radial box sizes $L_x \gtrsim 2H_0$. Note, however, that our $L_x = 2H_0$ simulation appears significantly less bursty than the $L_x = 2H_0$ simulation MRI-S1 simulation of Riols & Latter (2018) (see Figure 1 in their paper).

The vertical structure of the disk (not shown) also changes with box size, though we find that the buoyancy frequency $\langle \langle N_B^2 \rangle_{xy} \rangle_t$ is positive everywhere in all three simulations. Both the vertical heat and mass fluxes increase as the radial box size is reduced, which is unsurprising given that fluid elements are increasingly constrained to move in the vertical direction as the radial box size is reduced. The narrowest box is much more magnetized than the $L_x = 2H_0$ and $L_x = 4H_0$ boxes with $\langle \langle \beta \rangle_{xy} \rangle_t \sim 32$ at the mid-plane compared to $\langle \langle \beta \rangle_{xy} \rangle_t \sim 55$ in the wider boxes. We conclude that the non-linear MRI is almost certainly not converged in boxes of radial size $L_x = H_0$ (at least for the thermodynamics facilitated by the simple perfect gas equation of state that we have employed in our simulations). In particular narrow boxes seem to give rise to large outbursts that are not observed in wider boxes. Therefore it would be illuminating to investigate whether the MRI/radiative cycles observed in the simulations of Hirose et al. (2014) (which were carried out in boxes of radial size $L_x = H_0$) are less pronounced or even absent in wider boxes.

Appendix B

Tables of hydro simulations

Table B.1: Simulations of unforced compressible hydrodynamic convection in the shearing box carried out in PLUTO. No explicit viscosity or heating/cooling/thermal relaxation were included. The box size in all simulations is $4H \times 4H \times 4H$. The simulations were initialized either with a Gaussian temperature convectively unstable vertical disk profile, or with the convectively unstable vertical disk profile used in Stone & Balbus (1996) (SB96); see Section 3.1.1). $\langle\langle\cdot\rangle\rangle_t$ denotes a time- and volume-averaged quantity. $E_{\text{kin}, z}$ is the vertical kinetic energy density, which we have measured both at non-linear saturation (where it is maximum) and at the end of the simulation. $\Delta\rho$ is the percentage change in density over the course of the simulation, R_{xy} is the Reynolds stress, α is the alpha viscosity parameter, and σ is the linear phase growth rate given in units of Ω . All simulations were run for 38 orbits.

Run	Resolution	Profile	$\max(\langle E_{\text{kin}, z} \rangle)$	$\langle E_{\text{kin}, z} \rangle _{\text{end}}$	$\Delta(\langle\rho\rangle)$ (%)	$\langle\langle R_{xy} \rangle\rangle_t$	$\langle\langle\alpha\rangle\rangle_t$	$\sigma(\Omega)$
NSTR22c01	64^3	Gaussian Temp.	1.7×10^{-4}	2.0×10^{-6}	-1.4	$+4.8 \times 10^{-6}$	$+2.3 \times 10^{-5}$	0.2404
NSTR22c02	128^3	Gaussian Temp.	1.1×10^{-4}	3.7×10^{-6}	-3.4	$+2.7 \times 10^{-6}$	$+1.5 \times 10^{-5}$	0.2707
NSTR22c03	256^3	Gaussian Temp.	7.3×10^{-5}	9.8×10^{-6}	-8.8	$+1.7 \times 10^{-6}$	$+9.6 \times 10^{-6}$	0.2853
NSTR22c04	512^3	Gaussian Temp.	6.4×10^{-5}	1.9×10^{-5}	-12.7	$+1.3 \times 10^{-6}$	$+7.1 \times 10^{-6}$	0.2811
NSTR21c01	64^3	SB96	9.4×10^{-4}	9.0×10^{-6}	-2.1	$+1.9 \times 10^{-5}$	$+2.4 \times 10^{-5}$	0.2059
NSTR21c03	256^3	SB96	3.6×10^{-6}	3.9×10^{-5}	-10.3	$+6.0 \times 10^{-6}$	$+7.7 \times 10^{-6}$	0.2424

Table B.2: Simulations of unforced compressible hydrodynamic convection in the shearing box with different numerical schemes in PLUTO. No explicit viscosity or heating/cooling/thermal relaxation were included. All simulations were run in a box of size $4H \times 4H \times 4H$, and initialized with the Gaussian temperature profile discussed in Section 3.1.1. The time-averages of the volume-averaged Reynolds stress $\langle R_{xy} \rangle$ and of the volume-averaged angular momentum transport parameter $\langle \alpha \rangle$ have been taken over an interval spanning non-linear saturation. All simulations were run for 38 orbits.

Run	Resolution	Solver	Interpolation	Timestepping	$\langle \langle R_{xy} \rangle \rangle_t$	$\langle \langle \alpha \rangle \rangle_t$	Comments
NSTR22e01a	64^3	HLLC	WENO3	RK3	$+4.5 \times 10^{-6}$	$+2.5 \times 10^{-5}$	
NSTR22e01c	256^3	HLLC	WENO3	RK3	$+1.8 \times 10^{-6}$	$+1.0 \times 10^{-5}$	
NSTR22e02a	64^3	HLL	WENO3	RK3	-1.3×10^{-7}	-7.4×10^{-7}	$ \delta \mathbf{u} = \mathcal{O}(10^{-3})$
NSTR22e02c	256^3	HLL	WENO3	RK3	-1.1×10^{-5}	-6.1×10^{-5}	
NSTR22e03a	64^3	TVDLF	WENO3	RK3	-	-	No instability.
NSTR22e11a	64^3	HLLC	Linear TVD	Hancock	$+3.4 \times 10^{-6}$	$+2.2 \times 10^{-5}$	
NSTR22e11i	256^3	HLLC	Linear TVD	Hancock	$+1.5 \times 10^{-6}$	$+9.2 \times 10^{-6}$	UMIST limiter
NSTR22e12a	64^3	HLL	Linear TVD	Hancock	-2.2×10^{-5}	-1.4×10^{-4}	
NSTR22e12b	128^3	HLL	Linear TVD	Hancock	-1.6×10^{-5}	-1.0×10^{-4}	
NSTR22e12c	256^3	HLL	Linear TVD	Hancock	-9.5×10^{-6}	-6.2×10^{-5}	
NSTR22e12i	256^3	HLL	Linear TVD	Hancock	-1.2×10^{-5}	-8.2×10^{-5}	UMIST limiter
NSTR22e13a	64^3	TVDLF	Linear TVD	Hancock	-3.5×10^{-5}	-2.3×10^{-4}	
NSTR22e13b	128^3	TVDLF	Linear TVD	Hancock	-1.4×10^{-5}	-9.4×10^{-5}	
NSTR22e21a	64^3	HLLC	Linear TVD	RK2	$+4.0 \times 10^{-6}$	$+2.6 \times 10^{-5}$	
NSTR22e22a	64^3	HLL	Linear TVD	RK2	-2.2×10^{-5}	-1.3×10^{-4}	
NSTR22e22b	128^3	HLL	Linear TVD	RK2	-2.0×10^{-5}	-1.2×10^{-4}	
NSTR22e23a	64^3	TVDLF	Linear TVD	RK2	-1.9×10^{-5}	-1.2×10^{-4}	
NSTR22e23b	128^3	TVDLF	Linear TVD	RK2	-1.3×10^{-5}	-7.6×10^{-4}	

Table B.3: Simulations of unforced compressible hydrodynamic convection in the shearing box with explicit kinematic viscosity ν and thermal diffusivity χ included. No thermal relaxation was included. All simulations were run in a box of size $4H \times 4H \times 4H$. $\langle\langle\alpha\rangle\rangle_{t|\text{linear}}$ is the volume-averaged alpha viscosity parameter time-averaged over the linear phase, $\langle\langle\alpha\rangle\rangle_{t|\text{NL}}$ is the same quantity time-averaged over the non-linear phase, and $\min(\langle R_{xy} \rangle)$ is the minimum value of the xy -component of the volume-averaged Reynolds stress.

Run	Resolution	Instability?	ν	χ	Ra	Ri	Pr	$\langle\langle\alpha\rangle\rangle_{t \text{linear}}$	$\langle\langle\alpha\rangle\rangle_{t \text{NL}}$	$\min(\langle R_{xy} \rangle)$
NSTR22Ra1	256 ³	N	1.075×10^{-3}	4.300×10^{-4}	10^5	0.05	2.5	-	-	-
NSTR22Ra2	256 ³	Y	3.400×10^{-4}	1.360×10^{-4}	10^6	0.05	2.5	-5.1×10^{-6}	-3.9×10^{-5}	-2.1×10^{-5}
NSTR22Ra3	256 ³	Y	1.075×10^{-4}	4.300×10^{-5}	10^7	0.05	2.5	-3.5×10^{-6}	$+3.9 \times 10^{-6}$	-2.7×10^{-5}
NSTR22Ra4a	256 ³	Y	3.400×10^{-5}	1.360×10^{-5}	10^8	0.05	2.5	-1.0×10^{-6}	$+1.4 \times 10^{-5}$	-1.1×10^{-5}
NSTR22Ra4b	512 ³	Y	3.400×10^{-5}	1.360×10^{-5}	10^8	0.05	2.5	-2.8×10^{-6}	$+7.1 \times 10^{-6}$	-1.4×10^{-5}
NSTR22Ra5a	256 ³	Y	1.075×10^{-5}	4.300×10^{-6}	10^9	0.05	2.5	-3.5×10^{-8}	$+1.5 \times 10^{-5}$	-9.9×10^{-8}
NSTR22Ra5b	512 ³	Y	1.075×10^{-5}	4.300×10^{-6}	10^9	0.05	2.5	-1.6×10^{-7}	$+1.3 \times 10^{-5}$	-3.4×10^{-6}
NSTR22Ra6a	256 ³	Y	3.400×10^{-6}	1.360×10^{-6}	10^{10}	0.05	2.5	$+7.6 \times 10^{-8}$	$+1.5 \times 10^{-5}$	$+2.2 \times 10^{-13}$
NSTR22Ra6b	512 ³	Y	3.400×10^{-6}	1.360×10^{-6}	10^{10}	0.05	2.5	-1.1×10^{-8}	$+1.6 \times 10^{-5}$	-4.9×10^{-8}

Appendix C

Tables of MHD simulations

Table C.1: Vertically stratified MHD simulations: dependence on resolution. No explicit diffusion coefficients or cooling was included. All simulations were run in boxes of size $4H_0 \times 4H_0 \times 6H_0$. $\langle\langle\alpha\rangle\rangle$ is the time- and volume-averaged alpha viscosity parameter. Time-averages were taken from orbit 40 to orbit 100.

Run	Resolution	Grid cells / H_0	$\langle\langle\alpha\rangle\rangle$	$\langle\langle R_{xy} \rangle\rangle$	$\langle\langle M_{xy} \rangle\rangle$
NSTRMRI4e	$64 \times 64 \times 98$	$16/H_0$	—	—	—
NSTRMRI4c	$128 \times 128 \times 196$	$32/H_0$	0.0150	0.0014	0.0064
NSTRMRI4d	$256 \times 256 \times 392$	$64/H_0$	0.0131	0.0013	0.0057

Table C.2: Vertically stratified MHD simulations: dependence on radial box size. No explicit diffusion coefficients or cooling was included. All simulations were run in boxes of size $4H_0$ and $6H_0$ in the azimuthal and vertical directions, respectively. $\langle\langle\alpha\rangle\rangle$ is the time- and volume-averaged alpha viscosity parameter. Time-averages and maxima/minima were taken from orbit 50 to orbit 200.

Run	Radial Box Size	Resolution	$\langle\langle\alpha\rangle\rangle$	$\langle\langle R_{xy} \rangle\rangle$	$\langle\langle M_{xy} \rangle\rangle$	max($\langle\langle\alpha\rangle\rangle$)	min($\langle\langle\alpha\rangle\rangle$)
NSTRMRI3a	H_0	$32 \times 128 \times 196$	0.02618	0.002622	0.012347	0.0645	0.0078
NSTRMRI5	$2H_0$	$64 \times 128 \times 196$	0.01605	0.001577	0.006992	0.0350	0.0064
NSTRMRI4c	$4H_0$	$128 \times 128 \times 196$	0.01404	0.001354	0.006126	0.0234	0.0093

Table C.3: Vertically stratified MHD simulations with height-dependent cooling and explicit resistivity. All simulations were run in a box of size $4H_0 \times 4H_0 \times 6H_0$ and a resolution of $128 \times 128 \times 196$ (32 cells per H_0) except NSTRMC44fIHR which was run with a resolution of $256 \times 256 \times 392$ (64 cells per H_0). The explicit resistivity is denoted by η , τ_c denotes the cooling timescale (in orbits), $\text{Rm} \equiv H_0 c_{s0} / \eta$ is the magnetic Reynolds number, ΔT denotes the duration of the simulation (in orbits), $\max(\bar{F}_{\text{heat}})$ denotes the maximum value of the horizontally- and time-averaged vertical heat flux (rounded to four decimal places), and $\min(N_B^2)$ denotes the minimum value of the horizontally- and time-averaged square of the buoyancy frequency (rounded to two decimal places). Time-averages were typically taken over the last 80 orbits of a simulation.

Run	τ_c	η	Rm	ΔT	Result	$\max(\bar{F}_{\text{heat}})$	$\min(N_B^2)$
NSTRMC44e1	10	5×10^{-3}	200	200	MRI/convective cycles	0.0044	-0.09
NSTRMC44e2	5	5×10^{-3}	200	120	MRI/convective cycles	0.0062	-0.14
NSTRMC44a1	10	1×10^{-3}	10^3	80	MRI/convective cycles	0.0012	-0.04
NSTRMC44a2	5	1×10^{-3}	10^3	120	MRI/convective cycles	0.0018	-0.17
NSTRMC44b1	10	5×10^{-4}	2×10^3	100	MRI/convective cycles	0.0004	-0.07
NSTRMC44b2	5	5×10^{-4}	2×10^3	100	MRI/convective cycles	0.0017	-0.17
NSTRMC44b4	3	5×10^{-4}	2×10^3	100	MRI/convective cycles	0.0003	-0.25
NSTRMC44g1	3	2.5×10^{-4}	4×10^3	100	straight MRI	0.0007	-0.12
NSTRMC44g4	3	2.5×10^{-4}	4×10^3	100	straight MRI	0.0003	-0.30
NSTRMC44g5	2	2.5×10^{-4}	4×10^3	200	MRI/convective cycles	0.0005	-0.43
NSTRMC44g6	1	2.5×10^{-4}	4×10^3	200	MRI/convective cycles	0.0006	-0.51
NSTRMC44c1	10	1×10^{-4}	10^4	100	straight MRI	0.0013	-0.14
NSTRMC44c2	5	1×10^{-4}	10^4	80	straight MRI	0.0020	-0.24
NSTRMC44c4	3	1×10^{-4}	10^4	320	straight MRI	0.0022	-0.34
NSTRMC44c5	2	1×10^{-4}	10^4	100	straight MRI	0.0012	-0.43
NSTRMC44c6	1	1×10^{-1}	10^4	300	MRI/convective cycles	0.0011	-0.59
NSTRMC44c7	0.5	1×10^{-1}	10^4	100	disk disrupted	0.0007	-0.72
NSTRMC44d1	10	5×10^{-5}	2×10^4	100	straight MRI	0.0013	-0.16
NSTRMC44d2	5	5×10^{-5}	2×10^4	100	straight MRI	0.0020	-0.24
NSTRMC44f1	10	1×10^{-5}	10^5	100	straight MRI	0.0018	-0.15
NSTRMC44fIHR	10	1×10^{-5}	10^5	75	straight MRI	0.0029	-0.21

Table C.4: Vertically stratified MHD simulations with height-dependent cooling and explicit resistivity: effect of changing radial box size. The column headers have the same meaning as in Table C.3. The number of bursts refers only to those within the first 100 orbits of initialization.

Run	Radial Box Size	Resolution	τ_c	η	Rm	ΔT	Result	Number of bursts
NSTRMC44e1.1H	H_0	$32 \times 128 \times 196$	10	5×10^{-3}	200	100	MRI/convective cycles	1
NSTRMC44e1.2H	$2H_0$	$64 \times 128 \times 196$	10	5×10^{-3}	200	100	MRI/convective cycles	2
NSTRMC44e1	$4H_0$	$128 \times 128 \times 196$	10	5×10^{-3}	200	200	MRI/convective cycles	3

Table C.5: Vertically stratified MHD simulations without cooling but *with* explicit resistivity. Both simulations were run in a box of size $4H_0 \times 4H_0 \times 6H_0$ and at a resolution of $128 \times 128 \times 196$. The column headers have the same meaning as in Table C.3.

Run	τ_c	η	Rm	ΔT	Result	$\max(F_{\text{heat}})$
NSTRMC44e1NoCool	∞	5×10^{-3}	200	100	No MRI; no convection	0.000136
NSTRMC44b1NoCool	∞	5×10^{-4}	2×10^3	100	MRI; no convection	0.00137

Bibliography

- Aguiar A. C. B., Read P. L., Wordsworth R. D., Salter T., Yamazaki Y. H., 2010, *Icarus*, 206, 755
- Alexiades V., Amiez G., Gremaud P.-A., 1996, *Communications in Numerical Methods in Engineering*, 12, 31
- Aurnou J., Calkins M., Cheng J., Julien K., King E., Nieves D., Soderlund K., Stellmach S., 2015, *Physics of the Earth and Planetary Interiors*, 246, 52
- Bai X.-N., Stone J. M., 2013, *The Astrophysical Journal*, 767, 30
- Balbus S., Hawley J., 1991, *The Astrophysical Journal*, 376, 214
- Barr A. C., McKinnon W. B., 2007, *Geophysical Research Letters*, 34
- Bath G., 1973, *Nature Physical Science*, 246, 84
- Bénard H., 1900, *Rev Gén Sci Pure Appl*, 11, 1261
- Bénard H., 1901, PhD thesis, Collège de France
- Béthune W., Latter H., 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 6103
- Béthune W., Lesur G., Ferreira J., 2017, *Astronomy and Astrophysics*, 600, A75
- Bisnovatyi-Kogan G., Blinnikov S., 1977, *Astronomy and Astrophysics*, 59, 111

- Blaes O., Krolik J. H., Hirose S., Shabaltas N., 2011a, *The Astrophysical Journal*, 733, 110
- Blaes O., Krolik J. H., Hirose S., Shabaltas N., 2011b, *The Astrophysical Journal*, 733, 110
- Blandford R. D., Begelman M. C., 1999, *Monthly Notices of the Royal Astronomical Society*, 303, L1
- Block M. J., 1956, *Nature*, 178, 650
- Bodo G., Cattaneo F., Mignone A., Rossi P., 2012, *The Astrophysical Journal*, 761, 116
- Bodo G., Cattaneo F., Mignone A., Rossi P., 2013a, *The Astrophysical Journal*, 771, L23
- Bodo G., Cattaneo F., Mignone A., Rossi P., 2013b, *The Astrophysical Journal Letters*, 771, L23
- Bodo G., Cattaneo F., Mignone A., Rossi P., 2014, *The Astrophysical Journal Letters*, 787, L13
- Bodo G., Cattaneo F., Mignone A., Rossi P., 2015, *The Astrophysical Journal*, 799, 20
- Boley A. C., Durisen R., 2006, *The Astrophysical Journal*, 641, 534
- Borderies N., Goldreich P., Tremaine S., 1985, *Icarus*, 63, 406
- Boubnov B., Golitsyn G., 1995, *Convection in rotating fluids*. Springer Science & Business Media, Berlin
- Boyd J. P., 2001, *Chebyshev and Fourier spectral methods*. Dover, New York
- Breuer M., Manglik A., Wicht J., Trümper T., Harder H., Hansen U., 2010, *Geophysical Journal International*, 183, 150
- Brogan C., et al., 2015, *The Astrophysical Journal Letters*, 808, L3

- Burrows C. J., et al., 1996, *The Astrophysical Journal*, 473, 437
- Cabot W., 1996, *The Astrophysical Journal*, 465, 874
- Cannizzo J. K., 1993, in Wheeler J. C., ed., *Accretion Disks in Compact Stellar Systems*. World Scientific, pp 6–40, http://www.worldscientific.com/doi/abs/10.1142/9789814350976_0002
- Cannizzo J., Ghosh P., Wheeler J., 1982, *The Astrophysical Journal*, 260, L83
- Cargo P., Gallice G., 1997, *Journal of Computational Physics*, 136, 446
- Chandrasekhar S., 1961, *Hydrodynamic and hydromagnetic stability*. Dover, New York
- Chiang C.-Y., Goto T., Hashimoto T., Kim S. J., Matsuhara H., Oi N., 2019, *Publications of the Astronomical Society of Japan*, 71, 31
- Ciesielski A., Wielgus M., Kluźniak W., Sadowski A., Abramowicz M., Lasota J.-P., Rebusco P., 2012, *Astronomy and Astrophysics*, 538, A148
- Coleman M., Kotko I., Blaes O., Lasota J.-P., Hirose S., 2016, *Monthly Notices of the Royal Astronomical Society*, 462, 3710
- Coleman M. S., Yarger E., Blaes O., Salvesen G., Hirose S., 2017, *Monthly Notices of the Royal Astronomical Society*, 467, 2625
- Coleman M. S., Blaes O., Hirose S., Hauschildt P. H., 2018, *The Astrophysical Journal*, 857, 52
- Courant R., Friedrichs K., Lewy H., 1928, *Mathematische Annalen*, 100, 32
- D’Alessio P., Cantö J., Calvet N., Lizano S., 1998, *The Astrophysical Journal*, 500, 411
- Davidson P. A., 2015, *Turbulence: an introduction for scientists and engineers*. Oxford University Press, Oxford, UK

- Davis S. W., Stone J. M., Pessah M. E., 2010, *The Astrophysical Journal*, 713, 52
- EHT Collaboration 2019, *The Astrophysical Journal Letters*, 875, L1
- Einfeldt B., 1988, *SIAM Journal on Numerical Analysis*, 25, 294
- Emanuel K. A., 1994, *Atmospheric convection*. Oxford University Press, Oxford, UK
- Esposito L. W., 1986, *Icarus*, 67, 345
- Faulkner J., Lin D., Papaloizou J., 1983, *Monthly Notices of the Royal Astronomical Society*, 205, 359
- Favier B., Silvers L. J., Proctor M. R. E., 2014, *Physics of Fluids*, 26, 096605
- Fearn D. R., Loper D. E., 1981, *Nature*, 289, 393
- Flaig M., Kissmann R., Kley W., 2010, *EAS Publications Series*, 44, 117
- Flaig M., Ruoff P., Kley W., Kissmann R., 2012, *Monthly Notices of the Royal Astronomical Society*, 420, 2419
- Fleming T. P., Stone J. M., Hawley J. F., 2000, *The Astrophysical Journal*, 530, 464
- Ford H. C., et al., 1994, *The Astrophysical Journal*, 435, L27
- Freytag B., Ludwig H.-G., Steffen M., 1996, *Astronomy and Astrophysics*, 313, 497
- Fromang S., Papaloizou J., 2007, *Astronomy and Astrophysics*, 476, 1113
- Gammie C., Menou K., 1998, *The Astrophysical Journal Letters*, 492, L75
- Garaud P., 2018, *Annual Review of Fluid Mechanics*, 50, 275
- Ghoreyshi S. M. R., Ghanbari J., Salehi F., 2011a, *Publications of the Astronomical Society of Australia*, 28, 38

- Ghoreyshi S., Ghanbari J., Salehi F., 2011b, Publications of the Astronomical Society of Australia, 28, 290
- Gierasch P., et al., 2000, Nature, 403, 628
- Goldreich P., Lynden-Bell D., 1965, Monthly Notices of the Royal Astronomical Society, 130, 125
- Gough D., 1969, Journal of The Atmospheric Sciences, 26, 448
- Gressel O., 2013, The Astrophysical Journal, 770, 100
- Gu P.-G., Vishniac E. T., Cannizzo J. K., 2000, The Astrophysical Journal, 534, 380
- Guervilly C., Hughes D. W., Jones C. A., 2014, Journal of Fluid Mechanics, 758, 407
- Guillot T., et al., 2018, Nature, 555, 227
- Hadley G., 1735, Philosophical Transactions of the Royal Society of London, 39, 58
- Hameury J.-M., 2019, A review of the disc instability model for dwarf novae, soft X-ray transients and related objects (published in Advances in Space Science Research), <https://www.sciencedirect.com/science/article/pii/S0273117719307586>
- Hameury J.-M., Menou K., Dubus G., Lasota J.-P., Huré J.-M., 1998, Monthly Notices of the Royal Astronomical Society, 298, 1048
- Harten A., Lax P. D., Leer B. v., 1983, SIAM Review, 25, 35
- Harten A., Osher S., Engquist B., Chakravarthy S. R., 1986, Applied Numerical Mathematics, 2, 347
- Hawley J. F., Gammie C. F., Balbus S. A., 1995a, The Astrophysical Journal, 440, 742

- Hawley J. F., Gammie C. F., Balbus S. A., 1995b, *The Astrophysical Journal*, 440, 742
- Heinrich O., 1994, *Astronomy and Astrophysics*, 286, 338
- Held L. E., Latter H. N., 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 4797
- Hellier C., 2001, *Cataclysmic Variable Stars*. Springer Science & Business Media, Berlin
- Hirose S., 2015, *Monthly Notices of the Royal Astronomical Society*, 448, 3105
- Hirose S., Krolik J. H., Stone J. M., 2006, *The Astrophysical Journal*, 640, 901
- Hirose S., Krolik J. H., Blaes O., 2009, *The Astrophysical Journal*, 691, 16
- Hirose S., Blaes O., Krolik J. H., Coleman M. S., Sano T., 2014, *The Astrophysical Journal*, 787, 1
- Horne K., Cook M., 1985, *Monthly Notices of the Royal Astronomical Society*, 214, 307
- Horne K., Marsh T., 1986, *Monthly Notices of the Royal Astronomical Society*, 218, 761
- Hōshi R., 1979, *Progress of Theoretical Physics*, 61, 1307
- Huppert H. E., Turner J. S., 1981, *Journal of Fluid Mechanics*, 106, 299
- Iess L., et al., 2014, *Science*, 344, 78
- Ingersoll A., Gierasch P., Banfield D., Vasavada A., Team G. I., 2000, *Nature*, 403, 630
- Jaffe W., Ford H., Ferrarese L., van den Bosch F., O’Connell R. W., 1996, *The Astrophysical Journal*, 460, 214

- Jiang Y.-F., Stone J. M., Davis S. W., 2013, *The Astrophysical Journal*, 778, 65
- Jones C. A., Kuzanyan K. M., 2009, *Icarus*, 204, 227
- Ju W., Stone J. M., Zhu Z., 2016, *The Astrophysical Journal*, 823, 81
- Julien K., Knobloch E., Rubio A. M., Vasil G. M., 2012, *Physical Review Letters*, 109, 254503
- Kato S., 1978, *Monthly Notices of the Royal Astronomical Society*, 185, 629
- Kenyon S. J., 1990, in Wallerstein G., ed., *Cool Stars, Stellar Systems, and the Sun: Sixth Cambridge Workshop*. Springer, pp 206–216
- King A. R., Pringle J. E., Livio M., 2007, *Monthly Notices of the Royal Astronomical Society*, 376, 1740
- Kippenhahn R., Weigert A., Weiss A., 1990, *Stellar structure and evolution*. Springer, Berlin
- Klahr H., Henning T., Kley W., 1999, *The Astrophysical Journal*, 514, 325
- Kley W., Papaloizou J., Lin D., 1993, *The Astrophysical Journal*, 416, 679
- Knigge C., 2011, in Schmidtbreick M., Schreiber M., Tappert C., eds, *Evolution of Compact Binaries*. American Society of the Pacific, [arXivpreprintarXiv:1108.4716](#)
- Kronbichler M., Heister T., Bangerth W., 2012, *Geophysical Journal International*, 191, 12
- Kurbatov E. P., Bisikalo D. V., Kaygorodov P. V., 2014, *Physics-Uspekhi*, 57, 787
- La Dous C., 1989, *Astronomy and Astrophysics*, 211, 131
- Lada C. J., 2006, *The Astrophysical Journal Letters*, 640, L63

- Landau L. D., Lifshitz E. M., 1987, Fluid mechanics. Pergamon Press, London
- Laney C., 1998, Computational Gasdynamics. Cambridge University Press, Cambridge, UK, <https://books.google.co.uk/books?id=r-bYw-JjKGAC>
- Lappa M., 2009, Thermal convection: patterns, evolution and stability. John Wiley & Sons, New York
- Lasota J.-P., 2001, New Astronomy Reviews, 45, 449
- Latter H. N., Ogilvie G. I., 2008, Icarus, 195, 725
- Latter H. N., Papaloizou J., 2017, Monthly Notices of the Royal Astronomical Society, 472, 1432
- Latter H. N., Ogilvie G. I., Rein H., 2017, in Tiscareno M. S., Murray C. D., eds, Planetary Ring Systems: properties, structure, and evolution. Cambridge University Press, Cambridge, UK, [arXivpreprintarXiv:1701.04312](https://arxiv.org/abs/1701.04312)
- LeVeque R. A. J., 2002, Finite volume methods for hyperbolic problems. Cambridge University Press, Cambridge, UK
- Ledoux P., 1947, The Astrophysical Journal, 105, 305
- Lesur G., Ogilvie G. I., 2010, Monthly Notices of the Royal Astronomical Society: Letters, 404, L64
- Lesur G., Papaloizou J. C., 2009, Astronomy & Astrophysics, 498, 1
- Lesur G., Kunz M. W., Fromang S., 2014, Astronomy and Astrophysics, 566, A56
- Li S., 2005, Journal of Computational Physics, 203, 344
- Lin D., Papaloizou J., 1980, Monthly Notices of the Royal Astronomical Society, 191, 37

- Lin D., Papaloizou J., 1985, in Black D. C., Matthews M. S., eds, *Protostars and Planets II*. The University of Arizona Press, pp 981–1072
- Lin D., Papaloizou J., Kley W., 1993, *The Astrophysical Journal*, 416, 689
- Lin D.-B., Gu W.-M., Lu J.-F., 2011, *Monthly Notices of the Royal Astronomical Society*, 415, 2319
- Lynden-Bell D., Pringle J. E., 1974, *Monthly Notices of the Royal Astronomical Society*, 168, 603
- Lyra W., Richert A. J., Boley A., Turner N., Mac Low M.-M., Okuzumi S., Flock M., 2016, *The Astrophysical Journal*, 817, 102
- Mamajek E. E., 2009, in Usuda T., Tamura M., Ishii M., eds, *Exoplanets and disks: their formation and diversity*. American Institute of Physics, pp 3–10
- Medin Z., Cumming A., 2011, *The Astrophysical Journal*, 730, 97
- Menou K., Hameury J.-M., Lasota J.-P., Narayan R., 2000, *Monthly Notices of the Royal Astronomical Society*, 314, 498
- Menzel D. H., Milne E. A., 1966, *Selected papers on the transfer of radiation*. Dover
- Meyer F., Meyer-Hofmeister E., 1981, *Astronomy and Astrophysics*, 104, L10
- Meyer F., Meyer-Hofmeister E., 1982, *Astronomy and Astrophysics*, 106, 34
- Mignone A., 2014, *Journal of Computational Physics*, 270, 784
- Mignone A., 2015, *PLUTO User’s Guide (version 4.2)*
- Mignone A., Bodo G., Massaglia S., Matsakos T., Tesileanu O., Zanni C., Ferrari A., 2007, *The Astrophysical Journal Supplement Series*, 170, 228
- Mignone A., Flock M., Stute M., Kolb S., Muscianisi G., 2012, *Astronomy and Astrophysics*, 545, A152

- Mineshige S., Osaki Y., 1983, Publications of the Astronomical Society of Japan, 35, 377
- Minoshima T., Hirose S., Sano T., 2015, The Astrophysical Journal, 808, 54
- Miyagoshi T., Kageyama A., Sato T., 2010, Nature, 463, 793
- Miyoshi T., Kusano K., 2005, Journal of Computational Physics, 208, 315
- Moe M., Kratter K. M., Badenes C., 2019, The Astrophysical Journal, 875, 61
- Narayan R., Yi I., 1994, The Astrophysical Journal, 428, L13
- Naumann R., 1999, Technical Report NAS8-97-95, Spacelab Science Results Study, Volume II: Microgravity Science. University of Alabama in Huntsville
- Neufeld J. A., Bryson J. F., Nimmo F., 2019, Journal of Geophysical Research: Planets, 124, 1331
- Nimmo F., 2007, Treatise on Geophysics, 8, 31
- Ogilvie G., 2001, Monthly Notices of the Royal Astronomical Society, 325, 231
- Ogilvie G. I., 2003, Monthly Notices of the Royal Astronomical Society, 340, 969
- Ogilvie G. I., 2016, Journal of Plasma Physics, 82
- Okazaki A. T., 2000, in Smith M., Henrichs H., Fabregat J., eds, Proceedings of the International Astronomical Union Colloquium: the BE phenomenon in early-type stars. Cambridge University Press, pp 409–421
- Osaki Y., 1974, Publications of the Astronomical Society of Japan, 26, 429
- Paczynski B., 1965, Acta Astronomica, 15, 89
- Paczynski B., 1976, Comments on Astrophysics, 6, 95

- Papaloizou J., Lin D., 1988, *The Astrophysical Journal*, 331, 838
- Papaloizou J., Faulkner J., Lin D., 1983, *Monthly Notices of the Royal Astronomical Society*, 205, 487
- Pappalardo R., et al., 1998, *Nature*, 391, 365
- Pfeil T., Klahr H., 2019, *The Astrophysical Journal*, 871, 150
- Piran T., 1978, *The Astrophysical Journal*, 221, 652
- Pontin C. M., Barker A. J., Hollerbach R., André Q., Mathis S., 2020, *Monthly Notices of the Royal Astronomical Society*, 493, 5788
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1988, *Numerical recipes in C*. Cambridge University Press, Cambridge, UK
- Prialnik D., 2000, *An introduction to the theory of stellar structure and evolution*. Cambridge University Press, Cambridge, UK
- Pringle J., 1981, *Annual Review of Astronomy and Astrophysics*, 19, 137
- Quataert E., Gruzinov A., 2000, *The Astrophysical Journal*, 539, 809
- Rahmstorf S., 2006, *Encyclopedia of Quaternary Sciences*, 5
- Rayleigh L., 1916, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 32, 529
- Riley K. F., Hobson M. P., Bence S. J., 2006, *Mathematical methods for physics and engineering: a comprehensive guide*. Cambridge University Press, Cambridge, UK
- Rincon F., 2019, *Journal of Plasma Physics*, 85
- Riols A., Latter H., 2018, *Monthly Notices of the Royal Astronomical Society*, 474, 2212
- Ritter H., Kolb U., 2003, *Astronomy and Astrophysics*, 404, 301

- Roe P. L., 1981, *Journal of Computational Physics*, 43, 357
- Roe P., Pike J., 1984, in *Proceedings of the IMA Conference on Numerical Methods in Fluid Dynamics, Computing Methods in Science and Engineering IV*. Elsevier, Amsterdam
- Ross J., Latter H. N., Guilet J., 2016, *Monthly Notices of the Royal Astronomical Society*, 455, 526
- Ross J., Latter H. N., Tehranchi M., 2017, *Monthly Notices of the Royal Astronomical Society*, 468, 2401
- Rubio A. M., Julien K., Knobloch E., Weiss J. B., 2014, *Physical review letters*, 112, 144501
- Ruden S. P., Papaloizou J., Lin D., 1988, *The Astrophysical Journal*, 329, 739
- Rudiger G., 1987, *Acta Astronomica*, 37, 223
- Ryan B. R., Gammie C. F., Fromang S., Kestener P., 2017, *The Astrophysical Journal*, 840, 6
- Ryu D., Goodman J., 1992, *The Astrophysical Journal*, 388, 438
- Salvesen G., Simon J. B., Armitage P. J., Begelman M. C., 2016, *Monthly Notices of the Royal Astronomical Society*, 457, 857
- Sano T., Inutsuka S.-i., Turner N. J., Stone J. M., 2004, *Progress of Theoretical Physics Supplement*, 155, 409
- Scepi N., Lesur G., Dubus G., Flock M., 2018a, *Astronomy and Astrophysics*, 609, A77
- Scepi N., Lesur G., Dubus G., Flock M., 2018b, *Astronomy and Astrophysics*, 620, A49
- Scepi N., Dubus G., Lesur G., 2019, *Astronomy and Astrophysics*, 626, A116

- Schaefer B. E., 2010, *The Astrophysical Journal Supplement Series*, 187, 275
- Schmit U., Tscharnuter W., 1995, *Icarus*, 115, 304
- Schwarzschild K., 1906, *Nachrichten von der Königlichen Gesellschaft der Wissenschaften zu Göttingen. Math.-phys. Klasse*, 195, 41
- Schwarzschild M., 1958, *Structure and evolution of the stars..* Dover, New York
- Shakura N., Sunyaev R., 1973, *Astronomy and Astrophysics*, 24, 337
- Shakura N., Sunyaev R., 1976, *Monthly Notices of the Royal Astronomical Society*, 175, 613
- Shu C.-W., 1997, Preface to the republication of “Uniformly high order essentially non-oscillatory schemes, III,” by Harten, Engquist, Osher, and Chakravarthy. Academic Press, Cambridge, MA
- Siess L., 2009, *Astronomy and Astrophysics*, 497, 463
- Simon J. B., Hawley J. F., Beckwith K., 2008, *The Astrophysical Journal*, 690, 974
- Simon J. B., Hawley J. F., Beckwith K., 2011, *The Astrophysical Journal*, 730, 94
- Simon J. B., Beckwith K., Armitage P. J., 2012, *Monthly Notices of the Royal Astronomical Society*, 422, 2685
- Smak J., 1971, *Acta Astronomica*, 21, 15
- Smak J., 1982, *Acta Astronomica*, 32, 199
- Smak J., 1984a, *Acta Astronomica*, 34, 93
- Smak J., 1984b, *Acta Astronomica*, 34, 161
- Solheim J.-E., 2010, *Publications of the Astronomical Society of the Pacific*, 122, 1133

- Spiegel E., 1969, *Comments on Astrophysics and Space Physics*, 1, 57
- Spiegel E., 1972, *Annual Review of Astronomy and Astrophysics*, 10, 261
- Spohn T., Schubert G., 2003, *Icarus*, 161, 456
- Stancliffe R. J., Glebbeek E., 2008, *Monthly Notices of the Royal Astronomical Society*, 389, 1828
- Stern M. E., 1960, *Tellus*, 12, 172
- Stevens B., 2005, *Annu. Rev. Earth Planet. Sci.*, 33, 605
- Stone J. M., Balbus S. A., 1996, *The Astrophysical Journal*, 464, 364
- Stone J. M., Gardiner T. A., 2010, *The Astrophysical Journal Supplement Series*, 189, 142
- Stone J. M., Hawley J. F., Gammie C. F., Balbus S. A., 1996, *The Astrophysical Journal*, 463, 656
- Stone J. M., Gardiner T. A., Teuben P., Hawley J. F., Simon J. B., 2008, *The Astrophysical Journal Supplement Series*, 178, 137
- Sun Z.-P., Schubert G., Glatzmaier G. A., 1993, *Science*, 260, 661
- Tayler R., 1980, *Monthly Notices of the Royal Astronomical Society*, 191, 135
- Toro E. F., 2013, *Riemann solvers and numerical methods for fluid dynamics: a practical introduction*. Springer Science & Business Media, Berlin
- Toro E. F., 2019, *Shock Waves*, pp 1–18
- Toro E. F., Spruce M., Speares W., 1994, *Shock Waves*, 4, 25
- Tóth G., Odstrčil D., 1996, *Journal of Computational Physics*, 128, 82
- Turner J., 1974, *Annual Review of Fluid Mechanics*, 6, 37

- Turner N. J., Stone J. M., Sano T., 2002, *The Astrophysical Journal*, 566, 148
- Ulrich R. K., 1972, *The Astrophysical Journal*, 172, 165
- Umurhan O. M., Regev O., 2004, *Astronomy and Astrophysics*, 427, 855
- Vaidya B., Prasad D., Mignone A., Sharma P., Rickler L., 2017, *Monthly Notices of the Royal Astronomical Society*, 472, 3147
- Vauclair S., 2004, *The Astrophysical Journal*, 605, 874
- Wallace J. M., Hobbs P. V., 2006, *Atmospheric science: an introductory survey*. Elsevier, Amsterdam
- Warner B., 1995, *Cataclysmic Variable Stars*. Cambridge University Press, Cambridge, UK
- Weiss N. O., Proctor M., 2014, *Magnetoconvection*. Cambridge University Press, Berlin
- Williams J. P., Cieza L. A., 2011, *Annual Review of Astronomy and Astrophysics*, 49, 67
- Wood J., Horne K., Berriman G., Wade R., O'Donoghue D., Warner B., 1986, *Monthly Notices of the Royal Astronomical Society*, 219, 629
- Yadav R. K., Bloxham J., 2020, *Proceedings of the National Academy of Sciences*
- Yamaleev N. K., Carpenter M. H., 2009, *Journal of Computational Physics*, 228, 4248
- Zhang E.-H., Robinson E. L., 1987, *The Astrophysical Journal*, 321, 813
- Zhang E.-H., Robinson E., Nather R., 1986, *The Astrophysical Journal*, 305, 740
- Zhu Z., Ju W., Stone J. M., 2016, *The Astrophysical Journal*, 832, 193

Zimmer C., Khurana K. K., Kivelson M. G., 2000, *Icarus*, 147, 329

Zingale M., et al., 2002, *The Astrophysical Journal Supplement Series*, 143, 539