



MONASH UNIVERSITY

DOCTORAL THESIS

Dynamics of misaligned black hole accretion flows

Rebecca NEALON

B.Sc. (Astro Hons), B.Eng. (Aero Hons)

A thesis submitted for the degree of

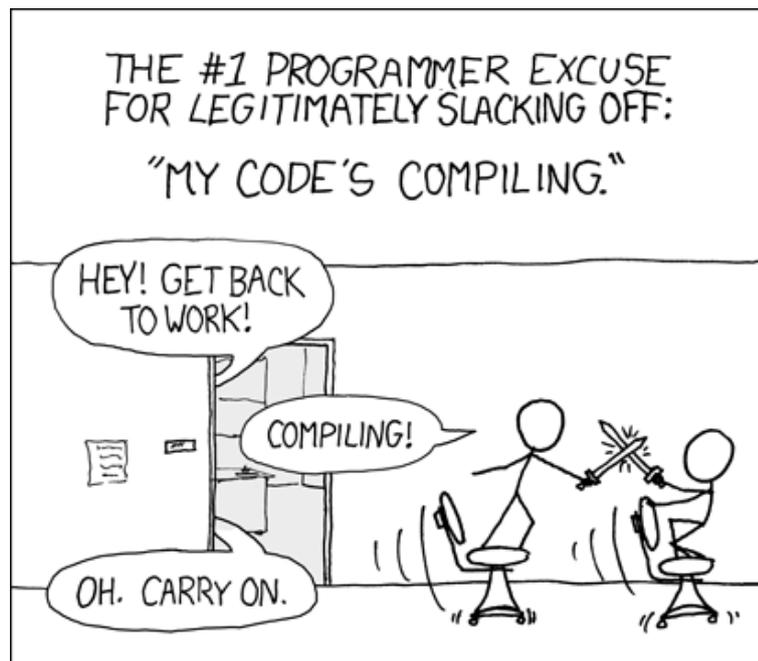
Doctor of Philosophy at the

School of Physics and Astronomy

Monash University, 2017

"Disc modeling is in some sense more of an art than a science ..."

Jim Pringle



'Compiling' - xkcd

Copyright Notice

© Rebecca Nealon (2017)

Under the Copyright Act 1968, this thesis must be used only under the normal conditions of scholarly fair dealing. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the author's permission.

Declaration of Authorship

I hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

This thesis includes two original papers published in peer reviewed journals and one submitted publication. The core theme of the thesis is *Dynamics of misaligned black hole accretion flows*. The ideas, development and writing up of all the papers in the thesis were the principal responsibility of myself, the student, working within the School of Physics and Astronomy under the supervision of Associate Professor Daniel Price.

The inclusion of co-authors reflects the fact that the work came from active collaboration between researchers and acknowledges input into team-based research. The following sections contain published material:

- Sections 3.1-3.5, 2.2.4 and 2.2.5 and Appendix A are presented in Nealon, Price, and Nixon (2015).
- Sections 4.1-4.4 are presented in Nealon, Nixon, Price, and King (2016).
- Sections 1.5.4 and 6.2.1-6.9 has been submitted for publication.

I have renumbered sections of submitted or published papers and notation changed where required in order to generate a consistent presentation within the thesis.

Student signature:



Date: November 20, 2017

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the student's and co-authors' contributions to this work. In instances where I am not the responsible author I have consulted with the responsible author to agree on the respective contributions of the authors.

Main Supervisor signature:



Date: November 20, 2017

Contents

Copyright Notice	iii
Declaration	iv
Acknowledgements	viii
Summary	ix
1 Introduction	1
1.1 The equations of accretion disc evolution	4
1.2 Viscosity in the disc	7
1.2.1 The magnetorotational instability (MRI)	7
1.2.2 The α model	8
1.3 Warped accretion discs	10
1.3.1 Torques in warped discs	11
1.3.2 Viscosity in warped discs	12
1.3.3 Precession of orbits in warped discs	13
1.3.4 Timescales in warped discs	14
1.3.5 Analytical studies	15
1.3.6 Simulations	19
1.3.7 Warped discs summary	26
1.4 Quasi-periodic oscillations	26
1.4.1 Observations of high frequency QPOs	28
1.4.2 Observations of low frequency QPOs	29
1.4.3 Connection to disc states	32
1.4.4 QPO models	33
1.4.5 QPOs summary	37
1.5 Geometrically thick flows	38
1.5.1 Light curves from tidal disruption events	39
1.5.2 Simulations of TDEs	41
1.5.3 The Papaloizou-Pringle instability (PPI)	43
1.5.4 Growth rate of the PPI	44
1.5.5 Simulations of the PPI	45
1.5.6 Tidal disruption events and the Papaloizou-Pringle instability sum- mary	47

2	Numerical methods	49
2.1	Smoothed particle hydrodynamics	50
2.1.1	Discretising properties onto particles	50
2.1.2	General form and identities used in SPH	53
2.1.3	Governing equations in SPH form	54
2.1.4	Conservation properties	55
2.1.5	Errors and derivatives	56
2.1.6	Shock capturing and dissipation terms	57
2.1.7	Physical viscosity	59
2.1.8	Summary of governing equations	62
2.1.9	Equation of state	63
2.2	Phantom	63
2.2.1	Arranging the particles	64
2.2.2	Comparing 3D simulations to 1D solutions	66
2.2.3	Disc viscosity	67
2.2.4	Simulating wave-like accretion discs	69
2.2.5	Modelling Lense-Thirring precession	71
2.3	Limitations	72
3	On the Bardeen-Petterson Effect in black hole accretion discs	75
3.1	Does the Bardeen-Petterson effect hold in the wavelike regime?	76
3.1.1	Is the inner disc aligned?	76
3.1.2	When does the disc break?	78
3.1.3	Can the disc accrete misaligned?	79
3.2	Numerical Method	80
3.2.1	Potentials	80
3.2.2	Precession Frequencies	80
3.2.3	Initial Conditions and Scope	81
3.3	Results	83
3.3.1	Tilt Oscillations	83
3.3.2	When does the disc break?	87
3.3.3	Can the disc accrete misaligned?	94
3.4	Discussion	96
3.5	Conclusion	99
4	Apsidal precession, disc breaking and viscosity	101
4.1	Simulations	102
4.1.1	Disc parameters	102
4.1.2	Initial conditions	103
4.2	Results	105
4.2.1	Apsidal precession neglected	105
4.2.2	Apsidal precession included	107
4.2.3	Breaking discs	107

4.3	Discussion	108
4.4	Conclusion	109
5	Quasi-periodic oscillations	111
5.1	Connecting tearing discs and QPOs	112
5.2	Precession frequencies	113
5.3	Luminosity in the disc	114
5.4	Generating a power density spectrum	115
5.5	Initial conditions	117
5.6	Initial parameter sweeps	118
5.7	Large disc simulations	121
5.7.1	Steady disc simulations	124
5.8	A different approach to measuring the precession	131
5.8.1	Measuring the precession dynamically	132
5.8.2	Luminosity projection with obscuration	133
5.9	High resolution simulations (S7 and S8)	139
5.10	Summary of QPO simulations	142
5.11	Conclusions	143
6	Thick accretion tori	147
6.1	The interaction of a cloud with a rotating black hole	150
6.2	The Papaloizou-Pringle instability	150
6.2.1	Growth of the PPI in a thin torus	151
6.3	Does the Papaloizou-Pringle instability lead to angular momentum transport?	155
6.4	Redistribution of the specific angular momentum	158
6.5	Does the Papaloizou-Pringle instability lead to ballistic accretion?	158
6.6	Are TDE remnants unstable to the PPI?	160
6.7	Which instability drives initial accretion?	162
6.8	Discussion	165
6.9	Conclusions	167
7	Conclusions	169
A	Measuring disc precession	174
B	Simulated power density spectra	176
B.1	PDS Test	176
B.2	Power density spectra	176

Acknowledgements

I sincerely thank my supervisor, Daniel ('the Camel') Price for his patience, guidance and infernal optimism over these last four years. As his student, Daniel has provided me with numerous opportunities and it is with his encouragement that I had the confidence to take them on. I thank him for consistently demanding the best from me, for the friendship he has extended so easily and for his compassion during the difficult times. Daniel is one of the rare breed being both excellent at his job and a most decent person; in him I could not have asked for a better supervisor.

I thank Chris Nixon (my second supervisor in all but name) for his continued support. Chris has given hours of assistance from the other side of the world, at least a million emails and numerous visits, help with my job applications, boosted my confidence at exactly the right times, always been helpful with the difficult problems and has always been a second person for me to talk to (or to listen to). I thank him for all his time over the years, as it has made such a difference for me.

I acknowledge financial support from the Australian Government for my Australian Postgraduate Award. I thank both the School of Mathematics and the School of Physics and Astronomy at Monash University for my Dean's Postgraduate Top Up Scholarship throughout my studies. For financial support while travelling, I thank Michael Morgan and the School of Physics and Astronomy, Monash University and the Astronomical Society of Australia. I additionally thank Chris Nixon and Giuseppe Lodato for their financial support and hospitality. I also acknowledge both gStar and Monarch for the computational time I was allocated during my studies and thank their respective support staff.

For their distraction and enabling of my procrastination, I thank the friendly staff and my office mates at MoCA. From the maths building to the portables to the 'Fish Bowl' to my spot next to the window, I have made far too many friends to mention individually but in particular I thank my PhD siblings Athira, Kate and Ben (and Grace). I thank my overseas friends (especially Giovanni and Alessia) for always showing me a good time and making me feel at home when I am away from home.

I thank my grandparents, Margaret and Geoff, for always being there when I needed them and my mum, Catherine, for giving me the space to be myself. For my sister Amy, one of my greatest achievements in life is having a sister that I am proud of and who is proud of me. I thank her for her ongoing support, right back to when she brought me snacks when I forgot to eat while studying for my exams. I would not have been able to reach as far as I have without the constant support from my extended, loving family. Thank you all.

For the most important person in my life, I thank my partner Marcus. In particular, I thank him for his encouragement of my passions even when it makes life more difficult for him, for making me laugh when I don't want to and for being the most dependable, trustworthy person in my life. I thank him for taking care of me even when I don't know how to take care of myself and for sharing his life, family and happiness with me.

To my dad, Gerry, I wish you were here to share in my success. I miss you and always hope to make you proud.

Summary

Accretion discs are the mechanism that allow rotating compact objects to gradually increase their mass. Most current investigations assume a simplified picture where the disc is smooth, planar, circular and aligned with the spin axis of the central object. Previous research using both analytical and numerical approaches have shown that relaxing these four assumptions reveals physical effects that strongly impact the evolution of accretion discs. This includes the phenomenon of disc ‘tearing’, where the rotation of the compact object is able to tear the misaligned disc into independently precessing rings. In this work we investigate the evolution of an accretion disc where the orbital plane is misaligned to the spin axis of a rotating black hole using three-dimensional numerical simulations.

We find that disc tearing is common to all geometrically thin strongly inclined accretion discs (independent of their disc thickness and viscosity ratio), and hence may be a general phenomenon. In the case where communication in the disc is controlled by pressure forces and the disc is weakly misaligned, our three-dimensional simulations demonstrate that the misalignment of the disc forms an oscillatory tilt profile with a non-zero misalignment at the inner edge — suggesting that any jets that may be launched from inclined discs are not necessarily aligned to the spin axis of the black hole. We demonstrate that general relativistic effects are the critical component needed to reveal this oscillatory profile.

We explore the relevance of tearing discs to observable phenomenon by considering low-frequency quasi-periodic oscillations measured from low mass X-ray binaries. These observational features are commonly explained by a model that invokes a discrete, precessing disc structure similar to those generated naturally in our simulations. We simulate a strongly inclined accretion disc with parameters derived from the low-mass X-ray binary XTE J1550-564, producing structures that precess with a rate comparable to those inferred from observations.

Finally, we investigate the evolution of geometrically thick accretion flows by considering a torus formed during a tidal disruption event. We demonstrate that such a torus may be unstable to the hydrodynamic Papaloizou-Pringle instability, leading to the formation of non-axisymmetric density structures and accretion of the torus. In the case of a torus formed with an initially weak magnetic field, we consider whether this instability may drive accretion before the magnetorotational instability establishes. Scaled to the parameters of a typical tidal disruption event this suggests a super-Eddington mass accretion within the first few hours of disruption. In the context of current observations, this suggests outflows or winds are likely.

List of Abbreviations

BHB	B lack h ole b inary
GR	G eneral r elativity
HFQPO	H igh f requency Q PO
HID	H ardness i ntensity d iagram
HIMS	H ard i ntermediate s tate
HSS	H ard s oft s tate
ISCO	I nnestmost s table c ircular o rbit
KH15	K rolik and H awley (2015)
LFQPO	L ow f requency Q PO
LHS	L ow h ard s tate
LMXB	L ow m ass X -ray b inary
LOP02	L ubow, O gilvie and P ringles (2002)
MHD	M agnetohydrodynamics
MRI	M agnetorotational i nstability
NS	N eutron s tar
PDS	P ower d ensity s pectrum
PPI	P apaloizou- P ringles i nstability
QPO	Q uasi-periodic o scillation
RPM	R elativistic p recession m odel
RXTE	R ossi X -ray T iming E xplorer
SIMS	S oft i ntermediate s tate
SPH	S moothed p article h ydrodynamics

Chapter 1

Introduction

Accretion discs form when gas with angular momentum falls towards a point mass (e.g. a neutron star, white dwarf, black hole or star). As the lowest energy orbit for a particular angular momentum is circular and gas is dissipative, the resultant orbit is naturally circular and is only defined by its distance from the massive object. This gas will have a range of angular momenta, generating a range of orbits at different distances from the object, ultimately forming a disc. Accretion is facilitated through the local and turbulent dissipation of energy and subsequent angular momentum transport, most likely driven by the magnetorotational instability (Balbus and Hawley, 1991). This process reduces the energy of the gas at a particular orbit, driving gas inwards, towards the central object. As the gas moves inwards the associated angular momentum is transported outwards, until most of the mass is close to the central object and a small amount remains at large radius with most of the angular momentum (Pringle, 1981).

Accretion discs are some of the most luminous objects in the Universe due to the large amount of energy they liberate from the accreting material. In young stars, proto-planetary discs lead to planet formation. For more massive objects, accretion discs around black holes are able to launch powerful, high velocity jets (\sim the speed of light). These jets are common across the full range of black hole masses, from low-mass X-ray binaries (with black hole masses $\sim 10 M_{\odot}$) through to active galactic nuclei (with masses $\sim 10^6$ - $10^8 M_{\odot}$) where the jets may extend to spatial scales that are comparable to the size of the host galaxy (for example, Centaurus A). Around smaller black holes, the disc itself may be steadily supplied from a companion star while around larger black holes it may be the result of a transient gravitational interaction like a tidal disruption event, where a star is shredded as it wanders too close to a black hole or the accretion of gas and dust by an active galactic nuclei.

Most investigations into accretion discs assume that they are aligned with the spin of the compact object — that is, the angular momentum vector of the disc is parallel with the spin vector of the compact object. However, consideration of observed X-ray binaries suggests that this should not necessarily be the case, with the degree of misalignment measured by the angle between the black hole spin vector and accretion disc angular momentum vector. In black hole low mass X-ray systems, the orbit of the companion donating material may be misaligned to the black hole spin (Jonker and Nelemans, 2004). As the angular momentum of the accretion disc is shared with the misaligned star, this would generate an accretion

disc that is misaligned to the spin of the black hole. Close to the black hole, where general relativistic effects are strongest, space-time is distorted (‘frame-dragging’) by the rotation of this massive object. This induces precession into orbits of nearby material, known as Lense-Thirring precession (Lense and Thirring, 1918). This effect acts on the misaligned disc, warping it such that the misalignment angle varies with distance from the black hole (Bardeen and Petterson, 1975). Theoretical studies also suggest that warping of misaligned accreting material can be driven by stars, when the magnetic field of the star is misaligned to the stellar spin axis (Lai, 1999) or from tidal forces from a companion star (Papaloizou and Terquem, 1995). Additionally, material that is aligned may be warped through radiation warping (Pringle, 1996).

The evolution of the disc and ultimately how long it takes to accrete is controlled by the rate of angular momentum transfer. This process and the profile of the disc as a function of radius is strongly affected by the presence and magnitude of a warp (e.g. Pringle, 1992). For small inclination warps ($\lesssim 15^\circ$), previous analytical work found that the profile of an inclined disc around a rotating black hole is oscillatory, with a misaligned inner edge (Ivanov and Illarionov, 1997; Lubow, Ogilvie, and Pringle, 2002). This finding has strong implications for any jets that may be launched from the innermost edge as they would not necessarily be aligned with the black hole spin — although alignment is often assumed by observational techniques used to measure the black hole spin. At higher misalignments (e.g. $\gtrsim 30^\circ$), analytical models can no longer be used and three-dimensional simulations must be employed. Recent simulations of strongly misaligned discs demonstrate that the disc may break into distinct, independently precessing structures rather than maintaining a continuous warped disc (Nixon et al., 2012). These ‘breaking’ and ‘tearing’ discs currently provide the only mechanism in accretion disc theory to generate discrete, precessing disc structures.

Numerous observations confirm both disc warping and precession occurs in nature. For example, the binary star system HK Tau has protoplanetary discs around both stars and both are misaligned with respect to the binary orbit (McCabe et al., 2011). Recent observations of the protoplanetary system TW Hydrae shows evidence of a shadow on the outer disc, with geometric constraints and timing of this shadow suggesting it comes from the precession of an inner misaligned disc section (Debes et al., 2017). Additionally, CO observations of HD 142527 imply the disc may be broken into two sections that are strongly misaligned with each other ($\gtrsim 70^\circ$, Casassus et al., 2015). Disc precession (and hence a misaligned disc) is also used to explain long X-ray periods in low mass X-ray binary systems like Her X-1 (e.g. Larwood, 1998). Both neutron star and black hole low mass X-ray binaries exhibit quasi-periodic oscillations, complex phenomena that have long been associated with Lense-Thirring precession (e.g. in neutron stars Stella and Vietri 1998 and type-C low-frequency quasi-periodic oscillations Ingram, Done, and Fragile 2009). Finally, observations of water maser emission also demonstrate the presence of warped discs — e.g. observations from V778 Cyg are consistent with an edge on warped disc (Babkovskaia et al., 2006).

Relaxing the assumption that the accretion disc and compact object spin are aligned reveals interesting physics and raises questions about the evolution of such discs. Are warped discs a transient or steady state? Do three-dimensional simulations agree with the current analytical predictions of disc evolution? Do the numerical methods that are used to simulate the evolution of warped discs agree? Is it possible to continuously generate the precessing structures that are made when a disc at large misalignments tears? Are these precessing structures related to observational features? If they are, can this be confirmed with simulations? Do these findings extend to geometrically thick accretion flows?

In this work we constrain ourselves to the case of a misaligned accretion disc around a rotating black hole. We begin by reviewing the relevant literature on misaligned accretion discs with both analytical and numerical approaches. We introduce the observational history of quasi-periodic oscillations and tidal disruption events, both of which will be considered in later chapters. This is followed by a summary of the numerical technique used and a brief description of the particular features that we employ from our numerical code PHANTOM. Future directions and a summary chapter complete this work. We consider four broad topics to better understand misaligned accretion flows:

In Chapter 3 the behaviour of a geometrically thin, inclined disc around a black hole is investigated at a variety of inclinations. At small inclinations, we compare to previous analytical results, validating our method. At higher inclinations, we confirm the phenomenon of disc ‘breaking’ and ‘tearing’, where the disc is observed to separate into distinct, precessing planes. We consider these results in the context of previous three-dimensional simulations that use a different numerical technique.

In Chapter 4 we consider the evolution of a warped disc around a rotating black hole in direct comparison to a previously published simulation that uses a different numerical technique. We compare the evolution of the disc when the techniques are modelling the same physics (i.e. a code comparison) as well as when more accurate physics is included. We compare our results using a similar visualisation method to the previously published work to better demonstrate our comparison. We additionally show that although a disc with the previously published parameters is not expected to tear, one at suitably higher inclination will.

In Chapter 5 we investigate the connection between the precessing structures found in tearing disc simulations with observations. We focus on quasi-periodic oscillations observed in low mass X-ray binaries where the compact object is a black hole. Observations of these features suggest that not only do they come from the innermost region of the accretion disc (where the precessing features are found in our simulations) but they are also likely to be strongly related to the precession of at least one discrete disc structure. We conduct simulations that have parameters based on the low mass X-ray binary J1550-564 and focus on simulating physically motivated tearing behaviour. A power density spectrum analysis is used in an attempt to compare directly to observations.

In Chapter 6 we finally consider the limit of a geometrically thick accretion in the form of a torus. Around a rotating black hole, these are likely to form after a black hole and star gravitationally interact in a tidal disruption event when the gas is able to circularise but not cool efficiently. We show that an idealised version of such a torus is unstable to the hydrodynamical Papaloizou-Pringle Instability and as such, should display non-axisymmetric perturbations. With three-dimensional simulations we consider the potential effect on the light-curve that is observed from such events.

1.1 The equations of accretion disc evolution

We begin by deriving the fundamental equations to describe the evolution of accretion discs. For convenience, we use cylindrical-polar coordinates with radius R , azimuthal angle ϕ , height z , using r to refer to the spherical radius. We assume that the disc is geometrically thin, equivalent to defining that the properties in the disc may be integrated in the z direction such that properties per unit surface can be considered rather than per unit density and assume axisymmetry. For example, the surface density, $\Sigma(R, t)$, may be expressed in terms of the density ρ with

$$\Sigma = \frac{1}{2\pi} \int_0^{2\pi} \int_{-\infty}^{+\infty} \rho dz d\phi. \quad (1.1)$$

This allows the continuity equation to be expressed in terms of Σ as (Pringle, 1981)

$$\frac{\partial \Sigma}{\partial t} + \frac{1}{R} \frac{\partial}{\partial R} (R \Sigma v_r) = 0, \quad (1.2)$$

where v_r is the radial component of the velocity vector \mathbf{v} . The Navier-Stokes equation can be written as

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\frac{1}{\rho} (\nabla P - \nabla \cdot \sigma) - \nabla \Phi, \quad (1.3)$$

where σ is the stress tensor, P the pressure and Φ the gravitational potential. The left side of this equation represents the momentum that is advected into the fluid through velocity gradients, and the right side the collaborative effect of pressure forces, viscous stress forces and gravity. Here the viscosity is assumed to be a shear viscosity such that it only acts in the $R\phi$ direction. Following Lodato (2008), we consider the three components of Equation 1.3 in turn, beginning with the centrifugal balance in the disc (the \mathbf{r} component of Equation 1.3). Using the cylindrical definition of ∇ and noting that the gradient of σ has no component in the r direction, this reduces to

$$\frac{\partial v_R}{\partial t} - \frac{v_\phi^2}{r} = -\frac{1}{\rho} (\nabla P - \nabla \Phi) |_{R}. \quad (1.4)$$

The radial velocity is much smaller than the azimuthal component, so this is discarded in a first order approximation. For the right hand side, the sound speed, $c_s(R)$, of the gas is defined as

$$c_s^2 \equiv \frac{dP}{d\rho}. \quad (1.5)$$

The sound speed is assumed to be much slower than the azimuthal velocity such that the effects of radial pressure gradients may be neglected — in practice, these effects are included in the initial conditions of our simulations to ensure the disc is started as close to equilibrium as possible. Additionally, for gas around a point with mass M the Newtonian gravitational potential is described by $\Phi(r) = -GM/r$. Here we assume that the gas is not self-gravitating, such that the mass of the disc is much less than the mass of the central mass M . With these two definitions, Equation 1.4 may be solved for the azimuthal velocity of the gas, giving

$$v_\phi = \sqrt{\frac{GM}{R}}. \quad (1.6)$$

Using the cylindrical R here recognises our assumption that the disc is geometrically thin, such that $R \approx r$ for most of the gas. Thus the velocity of material around a central potential of mass M is Keplerian ($v_\phi = v_k$). From the above definition it also follows that the Keplerian angular velocity is $\Omega_k = \sqrt{GM/R^3}$ and the angular momentum per unit area is $L(R) = \Sigma\sqrt{GMR}$.

The hydrostatic balance of the gas can be considered with the z component of Equation 1.3. Here v_z is also considered to be negligible so that pressure forces are purely balanced by the z component of the potential:

$$\frac{1}{\rho} \frac{\partial P}{\partial z} = -\frac{d\Phi}{dz}. \quad (1.7)$$

As before, the z position above the mid-plane of the of the gas in the disc is much smaller than the distance from the central potential. Using trigonometry, the right hand side of the above equation may be re-written as $d\Phi/dz = GMz/R^3$ (see Figure 2, Lodato, 2008). Assuming vertically isothermal, barotropic gas (i.e. $P = c_s^2(R)\rho$), the hydrostatic balance may thus be expressed as

$$\frac{c_s^2}{\rho} \frac{\partial \rho}{\partial z} = -\Omega_k^2 z. \quad (1.8)$$

Assuming that the sound speed of the gas is independent of z , the above equation may be integrated through the thickness of the disc (in the z direction) to give the description of the density in the disc

$$\rho(z) = \rho_0 e^{(-z^2/2H^2)}, \quad (1.9)$$

where ρ_0 is an integration co-efficient and we have introduced the scale height H of the gas, defined by $H \equiv c_s/\Omega_k$. Because this definition may be re-written as $H/R \equiv c_s/v_\phi$,

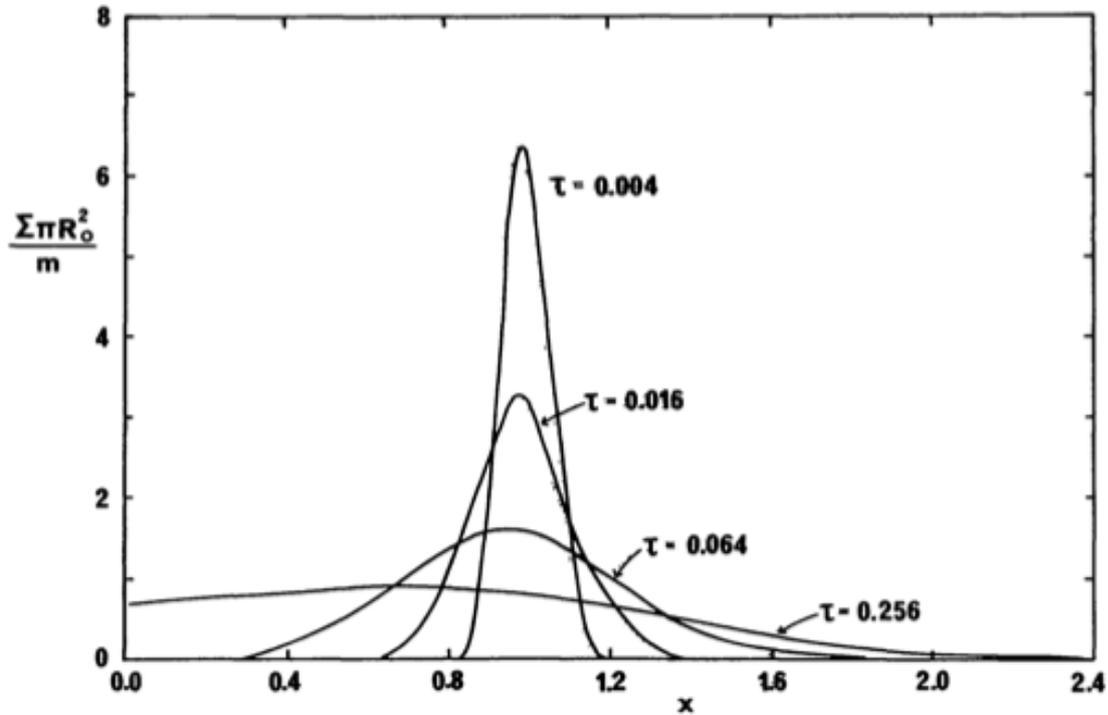


FIGURE 1.1: Viscous evolution of a ring of matter with mass m . Mass is transported inwards as angular momentum is transported outwards. Here $x = R/R_0$ where R_0 is the initial location of the ring and $\tau = 12\nu t/R_0^2$ (Figure 1, Pringle, 1981).

the assumption that the disc is geometrically thin is equivalent to stating that the sound speed is much slower than the azimuthal speed (which we assumed earlier).

Finally, we consider the azimuthal component of Equation 1.3. Integrating in the z direction, this may be expressed as

$$\Sigma \left(\frac{\partial v_\phi}{\partial t} + \frac{v_R v_\phi}{R} + v_R \frac{\partial v_\phi}{\partial R} \right) = \frac{1}{R^2} \frac{\partial}{\partial R} (R^2 T_{R\phi}). \quad (1.10)$$

In this form, the right hand side of the above equation represents the torque exerted by viscous forces, where $T_{R\phi}$ is the integral of $\sigma_{R\phi}$ in the z direction. These stresses may be re-written using $\sigma_{R\phi} = \rho \nu \Omega'$, where ν is the kinematic viscosity of the gas and Ω' is the derivative of Ω with respect to the radius. Assuming that the viscosity is independent of z ,

$$T_{R\phi} = \int \sigma_{R\phi} dz = \nu R \Sigma \frac{d\Omega}{dR}. \quad (1.11)$$

If a Keplerian disc is assumed with v_k and Ω_k as defined above, Equations 1.2 and 1.10 can be combined to produce a single equation describing the evolution of the surface density profile in the disc.

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

Importantly, Equation 1.12 shows that the evolution of the surface density profile is controlled exclusively by the present surface density and the kinematic viscosity ν . The nature of this description further implies that when a discontinuity is present in Σ it will generate viscous torques in radially neighbouring regions. If the disc is discretised radially into annuli, the internal torque of an inner annulus on the adjacent outer annulus is expressed as

$$G(R) = -2\pi\nu\Sigma R^3 \frac{d\Omega}{dR}. \quad (1.13)$$

This description has two implications. First, any discontinuities in Σ tend to be smoothed out as the torques in the disc are proportional to the radial gradient of the angular momentum. Second, because accretion discs have a decreasing Ω profile with radius, $G > 0$ and angular momentum is transported outwards as the surface density evolves. Figure 1.1 demonstrates this behaviour with the canonical ring spreading test (Lynden-Bell and Pringle, 1974; Pringle, 1981). Starting with a narrow Σ profile with mass m , the surface density is evolved with a chosen kinematic viscosity ν . As angular momentum is transported outwards mass is transported inwards, resulting in accretion of the matter and an increase of Σ at the inner edge. As Equation 1.12 demonstrates, the kinematic viscosity entirely governs the evolution of the surface density. As such, accurately representing the viscosity is necessary to model the evolution of accretion discs.

1.2 Viscosity in the disc

Turbulence in the disc is thought to drive these viscous torques and thus angular momentum transfer, but what drives the turbulence? For accretion discs around black holes, hydrodynamic and self-gravitating instabilities have been explored to try and explain this (Balbus, 2003; Lodato, 2008), but these do not generate sustained accretion on the timescales required (Balbus, 2003). In black hole accretion discs, the most likely underlying mechanism for the viscous transport of angular momentum is turbulence generated through the magnetorotational instability (MRI, Balbus and Hawley, 1991). In accretion discs not considered here, angular momentum transport may be primarily driven by other mechanisms including (but not limited to) tidally induced spiral shocks due to a companion (e.g. Livio and Spruit, 1991) and magnetically induced protoplanetary disc winds (e.g. Hasegawa et al., 2017).

1.2.1 The magnetorotational instability (MRI)

Continuous accretion discs naturally meet the two criteria to be unstable to the MRI (e.g. Balbus and Hawley, 1998);

$$\frac{d\Omega^2}{d\ln R} < 0 \quad \text{and} \quad \mathbf{B} \neq 0, \quad (1.14)$$

where Ω is the angular velocity profile, R the radius and \mathbf{B} the magnetic field in the disc. The first of these represents a decreasing angular velocity profile as the radius increases — naturally met by Keplerian discs that are differentially rotating. The second is a non-zero magnetic field both in the plane of the gas motion (the azimuthal direction) and the plane perpendicular to this (the z direction).

The instability can be conceptualised by considering two parcels of gas, situated adjacent to each other in the disc at slightly different radii as shown in Figure 1.2. The presence of the magnetic field means that these gas particles are connected by a magnetic tension. As the gas parcels orbit throughout the disc the inner one will do so faster due to the angular momentum profile, thus increasing the magnetic tension between them. The tension pulls backwards on the inner mass and forwards on the outer mass, transferring angular momentum from the inner to the outer mass. This transfer moves the inner mass onto a smaller orbit (decreasing the radius), the outer mass onto a larger orbit and increases the tension between the two masses. Increasing the magnetic tension then repeats the angular momentum transfer, moving the masses and increasing the tension again — hence the runaway nature of the MRI (Balbus and Hawley, 1991; Balbus and Hawley, 1998). Because this process requires the masses to be able to be moved further apart, the initial magnetic field must be weak rather than strong. The growth rate of the MRI is independent of this initial field strength, found to be 0.75 times the orbital frequency (Balbus and Hawley, 1991). As the viscosity in the disc is driven by the MRI, this mechanism is ultimately responsible for accretion of the material in the disc and the evolution of the surface density profile and so should be modelled accurately in numerical simulations.

1.2.2 The α model

Modelling the full magnetohydrodynamics (MHD) that generates turbulence which drives viscosity is difficult both analytically and numerically. Out of this difficulty, the α prescription was developed using dimensional analysis (Shakura and Sunyaev, 1973). The α parameter represents the internal stresses in units of local pressure, equivalent to dictating the rate at which angular momentum is transported. Because the transport of angular momentum results in accretion, α thus controls the evolution time scale of the disc. It is related to the kinematic viscosity averaged perpendicular to the plane of the disc, ν , by

$$\nu = \alpha c_s H, \quad (1.15)$$

where c_s is the sound speed in the disc and H the scale-height as defined previously. In this formulation, $0 \leq \alpha \leq 1$ and is often assumed to be constant throughout the disc (although it is not required to be, see King, Pringle, and Livio, 2007; Lodato, 2008). The development of the α parametrisation implies two features about the nature of the turbulence: first that the largest turbulent eddies are smaller than the disc thickness H and second that there are no shocks (i.e. the viscous flow speed is not supersonic, Pringle, 1981). The advantage of this parametrisation is that it allows simulations to be conducted that can mimic the

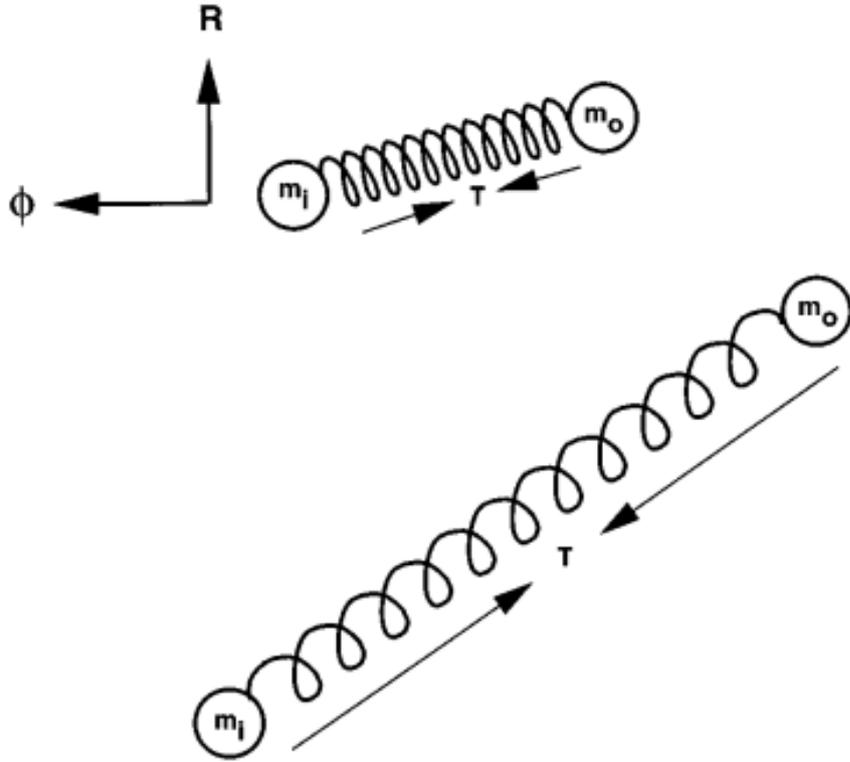


FIGURE 1.2: Conceptual representation of the MRI with spring tensions. m_i and m_o represent gas parcels and an inner and outer radius (r_i and r_o). Differential rotation increases tension T in the spring, transferring angular momentum from the inner to the outer mass and decreasing the inner while increasing the outer radius. This movement of the gas increases the tension, resulting in a run-away process (Figure 16, Balbus and Hawley, 1998).

effect of viscosity without actually simulating the full MHD turbulence thought to generate viscosity. The disadvantage of this approach is that it does not reveal anything about the nature of viscosity in accretion discs.

The actual value of α for a given disc can be estimated through observations (as it dictates the evolution timescale of the disc) and from simulations that include effects of the MRI. However, the observationally inferred values of $\alpha \approx 0.1 - 0.4$ are more than an order of magnitude higher than those measured from simulations (King, Pringle, and Livio, 2007). The discrepancy between these estimates was likely to be due to scaling issues, vertical boundary condition problems and grid size (resolution) issues (King, Pringle, and Livio, 2007). While it is likely that these problems can be solved using global accretion disc simulations (particularly those with a domain large enough that any boundary effects do not compromise the disc evolution) these are computationally challenging.

Until this point we have considered discs that are aligned to the black hole spin vector, however a similar analysis of warped discs demonstrates that the discrepancy between the estimate of α by observations and numerical simulations exists for these discs as well (King et al., 2013).

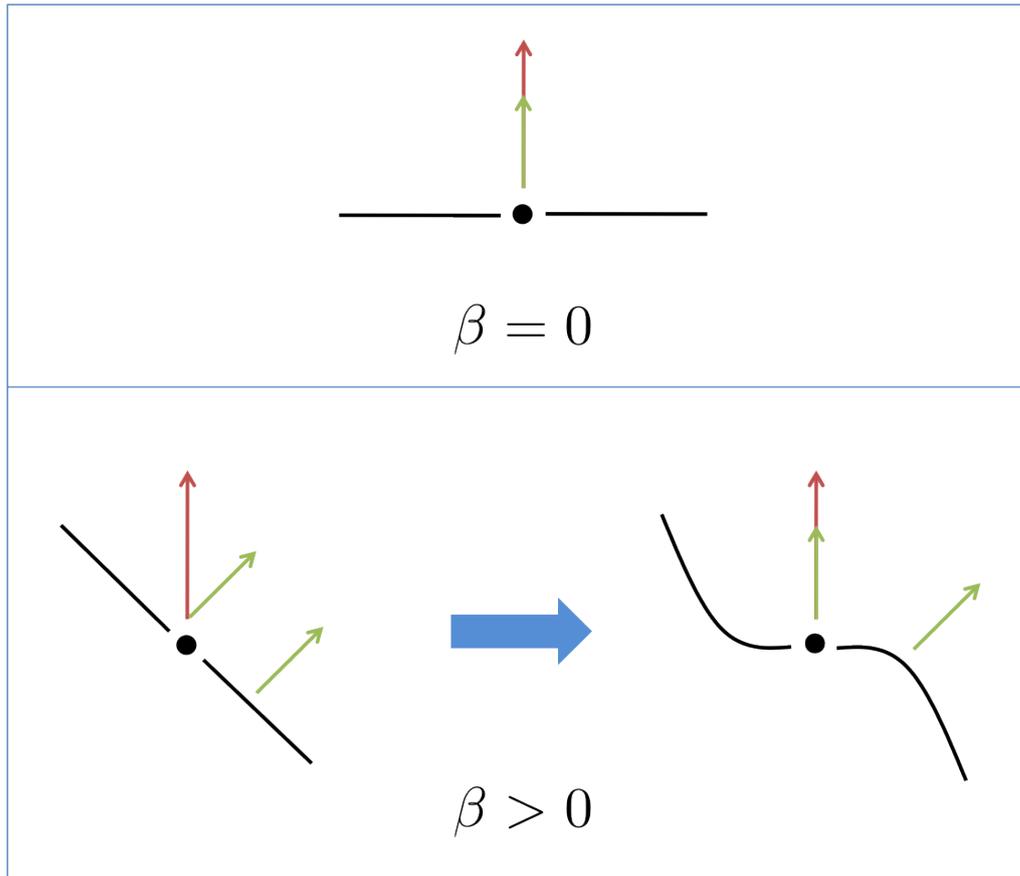


FIGURE 1.3: A side on depiction of an aligned (upper) versus misaligned (lower) accretion disc. Here the red vector represents the black hole spin and the green the local angular momentum vector of the accretion disc. The right-most disc represents the predicted steady state profile of a misaligned disc, known as the Bardeen-Petterson Effect (Bardeen and Petterson, 1975).

1.3 Warped accretion discs

The equations presented in § 1.1 inherently assume that angular momentum is transferred strictly in the radial direction. In the case where the central potential is a rotating black hole, this is equivalent to assuming that the spin vector of the black hole is parallel to the angular momentum vector of the accretion disc (which is perpendicular to the surface of the disc). The upper panel of Figure 1.3 demonstrates this configuration, which we refer to as an aligned disc. In the case where the disc is misaligned instead, the angular momentum transfer is not restricted to the plane of the black hole spin and the evolution of the disc becomes more complex.

The formation of a disc tilted to the black hole spin may occur naturally during low mass X-ray binary (LMXB) formation. As the larger star undergoes a supernova and forms a black hole or neutron star, an angular momentum kick can be given to the less massive companion (Jonker and Nelemans, 2004). This angular momentum kick inclines the orbit of the companion relative to the initial alignment, assumed to be in the spin plane of the more massive object. As the accretion disc is formed from material donated from

the companion, inclining its orbit forces the outer edge of the disc to be similarly inclined. Population synthesis models have shown that two-thirds of these systems should form with a small inclination ($\lesssim 10^\circ$), but that the rest have larger inclinations, with $\gtrsim 90^\circ$ possible but rarer (Fragos et al., 2010). Martin, Tout, and Pringle (2010) consider this formation scenario for the micro-quasar GRO J1655-40 which has an observed misalignment, finding that a natal kick of only a few tens of km/s would be required.

Maccarone (2002) found in systems such as LMXBs that the alignment timescale for a disc that is inclined to the black hole spin is at least on the order of the lifetime of the system itself (as angular momentum transferred through the disc is much smaller than the angular momentum of the black hole). This suggests that when a misaligned disc is formed it should be observable and indeed, misaligned discs have already been invoked to try and explain some features of LMXBs. For example, in binaries like GRO J 1655-40 and SAX J 1819-2525, misalignment between the disc and black hole has been offered as a potential explanation for the misalignment between the outer disc and the observed jets (Hjellming and Rupen, 1995; Greene, Bailyn, and Orosz, 2001; Fragile, Mathews, and Wilson, 2001). Importantly, stronger misalignments have been observed, with an inclination of $\gtrsim 55^\circ$ found for V4641 Sgr (Orosz et al., 2001).

Observational measurements of black hole spin often assume that there is no inclination between the jet and the inner disc plane (that launches the jet): the X-ray continuum spectral fitting method assumes this (e.g. Shafee et al., 2006) and the Fe K spectral line method would be more accurate with an inclination measure (e.g. Brenneman and Reynolds, 2006). The position of the inner disc here is inferred assuming alignment with the outer disc (e.g. upper panel of Figure 1.3), as there is currently no method to measure the position of the inner disc independently. However, the Bardeen-Petterson Effect would invalidate this when the outer disc is misaligned to the black hole (lower panel of Figure 1.3). Although the assumptions made during observations are justified and current techniques are not able to isolate the position of the inner disc, this may be able to be taken into account in future observations.

1.3.1 Torques in warped discs

In the case of a misaligned accretion disc around a rotating black hole, the rotation introduces the Lense-Thirring torque, defined per unit area as

$$\mathbf{T} = \mathbf{L} \times \boldsymbol{\Omega}_p, \quad (1.16)$$

where \mathbf{L} is the angular momentum vector of the disc per unit area with $\mathbf{L} = \Sigma R^2 \Omega \mathbf{l}$ and \mathbf{l} is the local angular momentum unit vector. The Lense-Thirring frequency is represented by $\boldsymbol{\Omega}_p$ with a magnitude defined by the black hole spin parameter a as $|\boldsymbol{\Omega}_p| = 2G^2 M^2 a / (c^3 R^3)$. Here M is the mass of the compact object, c is the speed of light and G is the gravitational constant. The spin vector and angular momentum vectors are indicated in red and green in

Figure 1.3, respectively. The misalignment angle of the disc is defined as the angle $\beta(R, t)$ formed between these two vectors (also referred to as the inclination angle). Inferring from Figure 1.3, the Lense-Thirring torque is zero in the case of an aligned disc where $\beta = 0$. However for a misaligned disc, the torque is non-zero and thus misaligned flows will respond to the rotation of a black hole.

The right side of the lower panel of Figure 1.3 shows the predicted disc profile where the inclination angle varies as a function of radius. This steady warped disc profile was introduced by Bardeen and Petterson (1975) by considering the competition between the Lense-Thirring torque and the viscous torques holding the disc as one continuous structure. At small radii, where the Lense-Thirring torque is strongest ($\mathbf{T} \propto R^{-3}$) the disc aligns with the black hole spin. At large radii, where the torque is weakest the disc maintains its original misalignment. This picture is completed by connecting the two regions with a smooth, continuous disc as in Figure 1.4 (Bardeen and Petterson, 1975). Additionally, as Figure 1.4 also describes the twist of the disc, $\gamma(R, t)$, defined as the angle made by the angular momentum vector away from the spin vector axis in a plane perpendicular to the tilt vector. The disc is most twisted in the innermost radii as the Lense-Thirring precession is fastest in this region. With these Euler angles, the local angular momentum vector can be described as

$$\mathbf{l}(R, t) = (\cos \gamma \sin \beta, \sin \gamma \sin \beta, \cos \beta). \quad (1.17)$$

This profile of an aligned inner region, smooth transition and a misaligned outer region has formed the foundation of the understanding of tilted accretion discs. Subsequent studies found that the equations employed by Bardeen and Petterson (1975) did not conserve angular momentum but found that the general profile is still consistent (Papaloizou and Pringle, 1983; Kumar and Pringle, 1985).

The strongest implication of the Bardeen-Petterson profile applies to the direction of jets that may be launched from the inner disc. As the jet is launched it will share the angular momentum of the disc, meaning that it will be launched perpendicular to its surface. In the case of a disc with an aligned inner edge, this coincides with the direction of the black hole spin *and* the outer disc (as is often assumed in observations). In the case of a misaligned disc and assuming the jet is still launched from the inner region, the jet direction will still be parallel to the spin (as the inner edge is aligned through the Bardeen-Petterson Effect) but will not necessarily align with the outer disc. This is in line with current observational findings, where the direction of jets in Seyfert galaxies does not appear to correlate to the inclination of the disc (Kinney et al., 2000).

1.3.2 Viscosity in warped discs

For aligned discs, viscosity was restricted to the plane of the black hole spin. However, warping of the disc complicates this because the viscosity in the disc must also seek to smooth the warp. In this case, the viscous forces that seek to flatten the disc generate

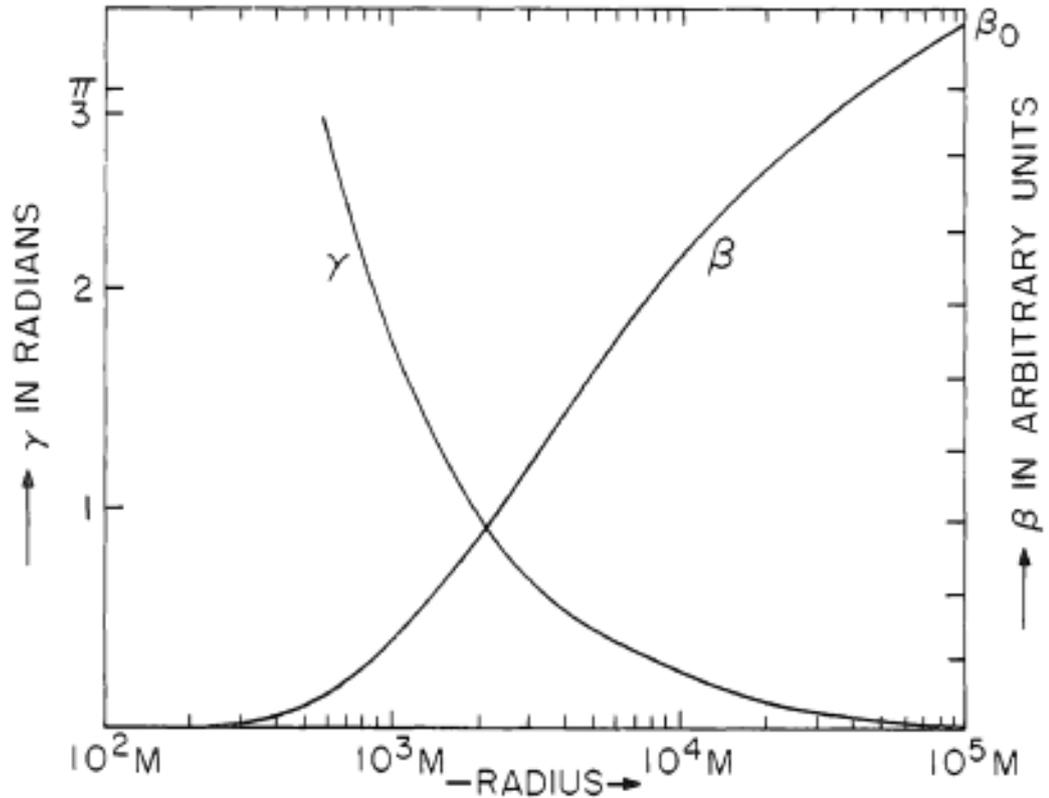


FIGURE 1.4: Relaxing of the inner disc into the black hole spin plane as predicted by Bardeen and Petterson (Figure 3, 1975). The misalignment angle β is measured between the black hole spin and disc angular momentum vectors, with the twist angle γ measured between an axis perpendicular to those and the angular momentum vector.

torques that induce local precession of the gas (Papaloizou and Pringle, 1983). This results in a disc that is smoothed faster than the viscous timescale, such that the warp will evolve faster than the disc is accreted.

Viscous effects in a warped disc can be treated in the same way as the traditional viscosity, ν , by introducing directional viscosity components for the radial and vertical communication of angular momentum. Here, ν_1 represents the component parallel to the local radial plane and ν_2 the component perpendicular to the local radial plane. A third term ν_3 is introduced due to precessional torques. In the case of an aligned disc, $\nu_2 = \nu_3 = 0$ and ν_1 reduces to ν . As before, these components can be related to the Shakura and Sunyaev (1973) parameters using the formulation in Equation 1.15 (i.e. $\nu_1 = \alpha_1 c_s H$, $\nu_2 = \alpha_2 c_s H$ and $\nu_3 = \alpha_3 c_s H$).

1.3.3 Precession of orbits in warped discs

Inclining the disc also complicates the orbits of gas in the inner region. Due to general relativistic effects, rather than traversing a purely circular (or elliptical) orbit, material orbiting around a rotating compact object traces out small epicycles. These epicycles can be decomposed into the radial and vertical direction and described by the radial and

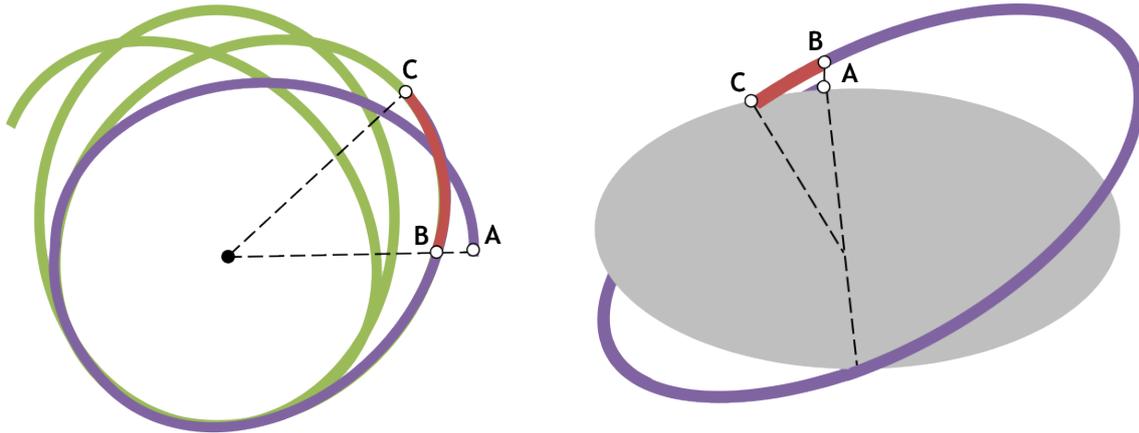


FIGURE 1.5: Apsidal (left) and nodal (right) precession across one orbit. In both cases, traversing from A to B is one orbital cycle; traversing from A to C is an epicycle and B to C is the associated precession (adapted from Figure 1 of Belloni and Stella, 2014).

vertical epicyclic frequencies, respectively. Figure 1.5 describes the relation of these to the orbital frequency, ν_ϕ (time taken to complete an orbit of 2π in the orbital plane). In the plane, the radial epicyclic frequency, ν_r , is determined from the time taken to return to the radius at apocenter (e.g. A→C). But the orbit precesses during this time (B→C, left panel of Figure 1.5) with the apsidal precession frequency $\nu_\phi - \nu_r$. Apsidal precession of the orbit is independent of its inclination or the spin of the compact object, so also occurs in aligned discs around non-rotating black holes.

Even in misaligned discs, apsidal precession is restricted to the plane of the orbit. Nodal precession of the orbits is similar (right panel, Figure 1.5) but only occurs for inclined orbits. Out of the orbital plane the vertical epicyclic frequency is found from the time taken to cross the line of nodes again (this is the same as crossing the projection of the orbital plane, A→C in the right panel). The vertical epicyclic frequency, ν_θ , determines the nodal precession frequency $\nu_\phi - \nu_\theta$. As we shall see in Chapters 3 and 4, accounting for both nodal and apsidal precession in simulations is critical to correctly recovering the evolution of warped accretion discs around rotating black holes.

1.3.4 Timescales in warped discs

In an aligned disc there are three important timescales, each defined as a function of the radius R . The first is the dynamical time scale, constrained by the time it takes for material to complete an orbit:

$$t_{\text{dynamical}} = \frac{1}{\Omega(R)}. \quad (1.18)$$

The second is the sound-crossing timescale, defined by the time taken for material to communicate over the distance R as (e.g. Papaloizou and Lin, 1995)

$$t_s \equiv \frac{R}{c_s(R)}. \quad (1.19)$$

The timescale over which material accretes in the disc is governed by the rate of angular momentum transfer that is controlled by ν . For a warped disc ν is treated in three components, however the accretion rate is only dependent on ν_1 , acting in the radial direction. Thus the accretion timescale can be written as

$$t_{\text{accretion}} \sim \frac{R^2}{\nu_1}, \quad (1.20)$$

reducing to R^2/ν for an aligned disc. In the case of warped discs, there are three additional timescales to consider. The timescale for a warp to diffuse in the disc is controlled by the angular momentum transfer perpendicular to the radial plane (ν_2) as

$$t_{\text{warp diffusion}} \sim \frac{R^2}{\nu_2}. \quad (1.21)$$

The timescale to damp such a warp (i.e. the diffusion time across the thickness of the disc) is can be expressed as (Lubow and Ogilvie, 2000)

$$t_{\text{damp}} = \frac{1}{\alpha_1 \Omega}. \quad (1.22)$$

Misaligned discs will also experience precession, either due precession of the individual orbits or from external forces (e.g. Lense-Thirring precession). The magnitude of the precessional torque is proportional to the magnitude of the Lense-Thirring frequency, given as $\Omega_p = |\mathbf{\Omega}_p|$. The precession timescale may be expressed as

$$t_{\text{precession}} \sim \frac{1}{\Omega_p(R)}. \quad (1.23)$$

The behaviour of an accretion disc can be informed by careful consideration of these timescales and by comparisons between them — for example, when the accretion timescale acts faster than the warp diffusion or sound crossing timescales, we would expect a misaligned disc to accrete material before it is able to align with the spin of the black hole.

1.3.5 Analytical studies

Figure 1.6 shows the cross section of fluid (represented by the grey boxes) orbiting in a warped disc. In this representation, the neighbouring rings of fluid are most strongly misaligned at this instant of the orbit (Lodato and Pringle, 2007). After another quarter of an orbit however the neighbouring rings of fluid will almost be aligned (i.e. the relative warp is weakest when the disc crosses the $z = 0$ plane). Thus as the material in a warped disc orbits, it cycles through regions of strong and weak misalignment with its neighbouring region. The arrows in Figure 1.6 show the radial pressure gradient that occurs when neighbouring rings of fluid are misaligned. The strength of this pressure gradient is dependent on the misalignment of the neighbouring rings — thus as material orbits around

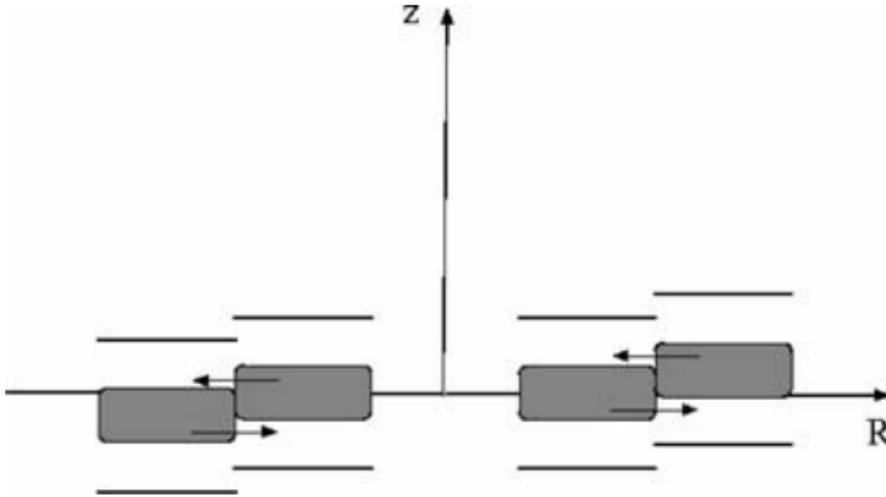


FIGURE 1.6: Horizontal pressure gradients experienced by gas during its orbit. The shaded area has higher pressure, with arrows showing the direction of horizontal pressure gradients induced by the warp. As the gas orbits it feels a strong radial pressure gradient at the regions shown here, but none where the disc crosses the $z = 0$ plane. Thus the gas experiences an oscillating radial pressure gradient which induces epicyclic motions in the gas (Figure 10, Lodato and Pringle, 2007).

a warped disc, it experiences an oscillating radial pressure gradient that is dependent on the strength of the warp (Papaloizou and Pringle, 1983). This description assumes that the warp is linear (as represented in Figure 1.6), that the disc is near-Keplerian and that there are no large scale magnetic fields that will affect the evolution of the disc.

The result of this oscillating radial pressure gradient is the wave propagation of the warp. The radial pressure gradient drives epicyclic motions in the gas, generating a torque in the disc and launching a wave that communicates the presence of a warp. The turbulence in the disc (from the MRI) that is required for accretion damps this wave propagation. Two regimes of warp propagation can be described through comparison of the warp propagation timescale (Equation 1.19, which communicates the warp) and the damping timescale (Equation 1.22, which slows down communication of the warp). By using the α parameter to describe the viscosity in the disc, these two regimes may equivalently be described by comparison of α to H/R . In the case where $\alpha \lesssim H/R$ the warp propagation is described as ‘wave-like’ and in the case where $\alpha \gtrsim H/R$ the warp propagation is ‘diffusive’.

Diffusive regime: $\alpha \gtrsim H/R$

In thin discs communication of warps is governed by diffusion (Papaloizou and Pringle, 1983). This means that disturbances in the disc are smoothed out (diffused) over time. The evolution of the angular momentum in the disc is described by (Pringle, 1992; Ogilvie,

1999; Nixon and King, 2012):

$$\begin{aligned}
\frac{\partial \mathbf{L}}{\partial t} = & \frac{1}{R} \frac{\partial}{\partial R} \left(\frac{(\partial/\partial R)[\nu_1 \Sigma R^3 (-\Omega')]}{\Sigma(\partial/\partial R)(R^2 \Omega)} \mathbf{L} \right) \\
& + \frac{1}{R} \frac{\partial}{\partial R} \left(\frac{1}{2} \nu_2 R |\mathbf{L}| \frac{\partial \mathbf{L}}{\partial R} \right) \\
& + \frac{1}{R} \frac{\partial}{\partial R} \left(\left[\frac{\frac{1}{2} \nu_2 R^3 \Omega |\partial \mathbf{L} / \partial R|^2}{(\partial/\partial R)(R^2 \Omega)} + \nu_1 \left(\frac{R \Omega'}{\Omega} \right) \right] \mathbf{L} \right) \\
& + \frac{1}{R} \frac{\partial}{\partial R} \left(\nu_3 R |\mathbf{L}| \mathbf{1} \times \frac{\partial \mathbf{L}}{\partial R} \right) \\
& + \boldsymbol{\Omega}_p \times \mathbf{L}.
\end{aligned} \tag{1.24}$$

Here the first term represents the viscous diffusion of mass. The second term, involving ν_2 , represents the diffusion of the tilt of the disc. The third term describes the advection of angular momentum (inwards or outwards, depending on its sign), thus representing the advective torque. The first cross term represents the precessional torque which causes successive rings of material in the disc to precess and only occurs when the disc is tilted (Ogilvie, 1999). The final cross term represents the external torque from the black hole rotation, and causes the angular momentum vector for each discrete ring in the disc to precess around the spin vector of the black hole.

Simulations that assume properties of the disc can be described as a function of the radius only (a ‘1D’ formulation) were conducted in this regime most recently by Nixon and King (2012). In comparison to previous studies, they included an external torque provided by Lense-Thirring precession and the ν_3 term in Equation 1.24 above as well as non-linear fluid effects specified in Ogilvie (1999). In doing so, they found that it is possible for the disc to break into distinct planes, identified by a steepening of the tilt profile. Breaking was observed most clearly at $\beta \approx 45^\circ$, suggesting that it was due to non-linear effects. As the misalignment in systems that are likely to be in the diffusive regime (such as LMXBs) is a long lasting property of the system, they concluded that disc breaking should be considered when interpreting observations and considering the accretion rate of these systems (Nixon and King, 2012). Finally, when jets are launched from such a disc (not taken into account in these simulations), they are unlikely to be aligned with the outer disc.

In this regime, the relationship between the radial and perpendicular viscosity components for small warp amplitudes was found by Ogilvie, 1999 (and confirmed by Lodato and Price, 2010) to be

$$\frac{\nu_2}{\nu_1} = \frac{1}{2\alpha^2} \frac{4(1+7\alpha^2)}{4+\alpha}, \tag{1.25}$$

for linear warps. Additionally, Ogilvie and Dubus (2001) found that the α component due to precessional torques can be related by

$$\alpha_3 = \frac{3(1-2\alpha^2)}{2(4+\alpha^2)}. \tag{1.26}$$

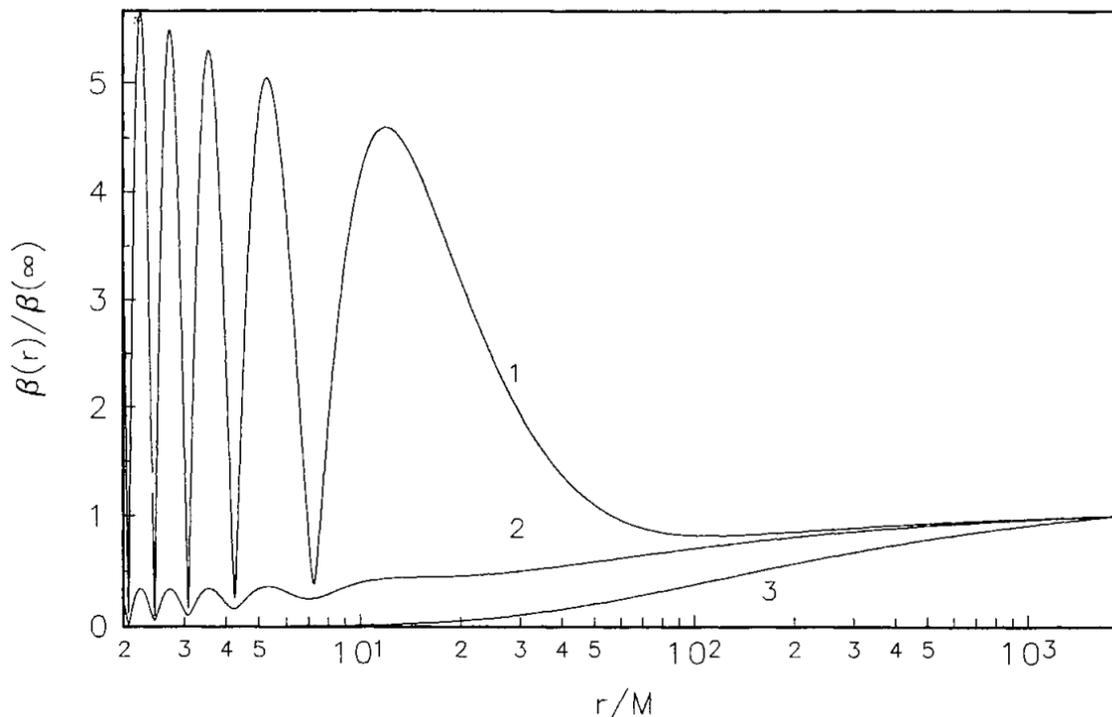


FIGURE 1.7: Oscillations in the inclination of the disc as a function of the radius, predicted by Ivanov and Illarionov (Figure 2, 1997). The different lines represent the effect of different viscosity values, with lines 1-3 corresponding to $\alpha = 0.01, 0.1, 1.0$, respectively.

Bending wave regime: $\alpha \lesssim H/R$

In thicker (but still geometrically thin) discs the communication of warps is governed by pressure forces, travelling with the half sound speed of the disc as a bending wave (Papaloizou and Pringle, 1983). In these discs, disturbances are transmitted as waves, separating into two travelling waves (one moving outwards and one moving inwards in the disc). In this regime, the conservation of horizontal angular momentum (i.e. in the plane of the unwarped disc) suggests (Lubow and Ogilvie, 2000):

$$\Sigma R^2 \Omega \frac{\partial \mathbf{l}}{\partial t} = \frac{1}{R} \frac{\partial \mathbf{G}}{\partial R} + \mathbf{T}, \quad (1.27)$$

where T represents the magnitude of the Lense-Thirring torque. The internal viscous stresses in the disc evolve via

$$\frac{\partial \mathbf{G}}{\partial t} - \left(\frac{\Omega^2 - \kappa^2}{2\Omega} \right) \mathbf{l} \times \mathbf{G} + \alpha \Omega \mathbf{G} = \frac{PR^3 \Omega}{4} \frac{\partial \mathbf{l}}{\partial R}. \quad (1.28)$$

Here the κ and Ω refer to the horizontal epicyclic and orbital frequencies of the gas in the disc. Further simplifications of these equations assume that the inclination is small enough to be considered linear. Ivanov and Illarionov (1997) found under these conditions that the disc should have an oscillatory profile close to the black hole, and that the innermost edge of the accretion disc should not be aligned to the black hole spin. This is shown in

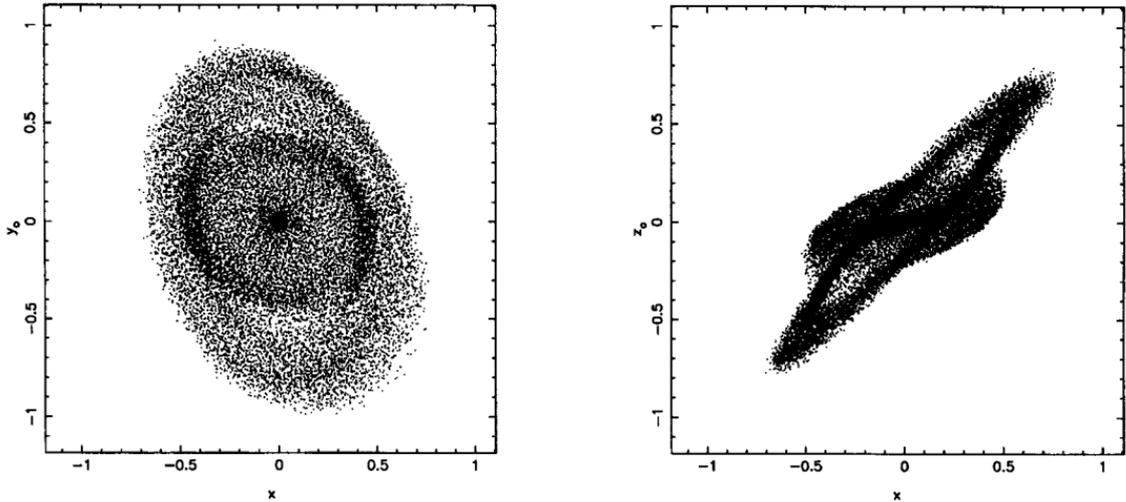


FIGURE 1.8: First identification of ‘disc breaking’, simulated by Larwood et al. (1996, Figure 7). Initially inclined at 45° , despite this SPH simulation using 17500 particles it is still able to show a final state where the disc separated into two distinct planes.

Figure 1.7 and was also studied in depth by Lubow, Ogilvie, and Pringle (2002). These oscillations are dependent on the opposing signs of the precession frequencies and so are limited to the case around a prograde black hole — for example, the precession frequencies that describe the behaviour around a binary or a retrograde black hole have opposing signs, and so this effect will not occur in these systems. The linear theory presented here cannot be extended beyond small inclination angles, although Ogilvie (2006) has extended this to the weakly non-linear regime.

As in the diffusive case, the resultant disc profile has implications for jets that are launched from the disc. As jets are launched along the angular momentum vector of the disc at the innermost edge, if this is not aligned with the black hole spin the jet may be misaligned to the disc and the black hole spin. The analytical approach thus suggests that depending on the regime, jets may be misaligned to either the outer disc (as in the case of diffusive discs) or to both the black hole spin and the outer disc (in the case of wave-like discs). This is wholly contradictory to an aligned disc, which effectively assumes that the outer and inner disc are both aligned with the jet of the black hole. Current observations are restricted to assume that jets are parallel to the spin, regardless of the disc alignment.

1.3.6 Simulations

Two main numerical techniques are used to investigate accretion disc evolution in 3D; grid based and particle methods. As each method is useful in a particular disc regime, simulations are naturally grouped by their techniques. Broadly speaking, grid based simulations are able to include magnetic fields allowing viscosity to be driven through the MRI but are restricted to lower inclinations. By contrast, smoothed particle hydrodynamics (SPH)

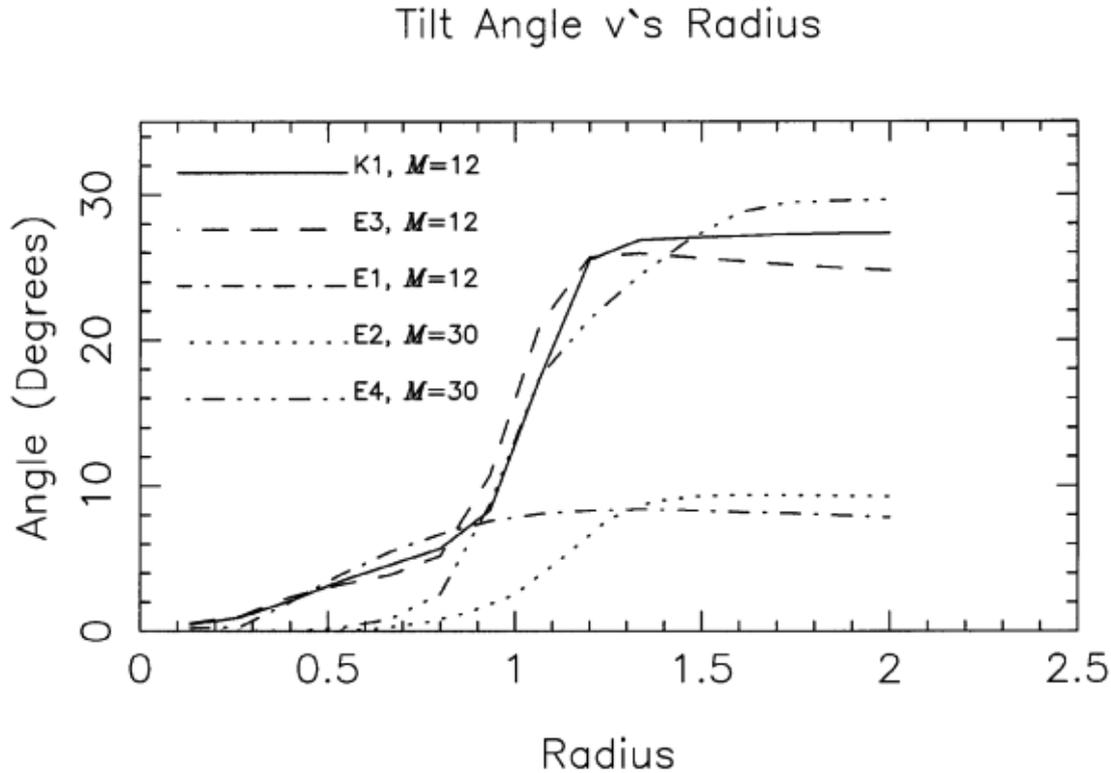


FIGURE 1.9: First simulations looking at larger disc inclinations by Nelson and Papaloizou (2000, Figure 12). The profile of aligned inner region, smooth transition and misaligned outer region at smaller inclinations confirms the Bardeen-Petterson effect (see Figure 1.4). The rapid increase in the tilt at $R = 1$ is suggestive of disc breaking in the high inclination simulations.

simulations use the α model to represent viscosity in the disc and are able to explore large disc/black hole misalignments.

Particle method

Due to its formalism SPH is ideal for investigating flows with complex geometries and can be used for accretion disc simulations in either the bending-wave or diffusive regimes. Using this technique, the effect of large inclinations in the disc was first studied by Larwood et al. (1996). Their simulations made use of 17500 particles and considered disc inclinations up to 45° . Despite the relatively low resolution, solid body precession (where the disc precesses as a whole) was found for most of their simulations. Additionally, as shown in Figure 1.8, they found that the disc may be broken into distinct planes defined by their tilt due to differential precession. Although they considered an accretion disc around a binary, this behaviour is similar to what would be expected in black hole discs.

Investigations into the Bardeen-Petterson Effect using SPH simulations continued with Nelson and Papaloizou (2000). They conducted simulations with inclinations of 10° and 30° , starting with the whole disc fully inclined to the black hole spin and using a post-Newtonian approximation. For a variety of disc thicknesses and α values they found behaviour that was consistent with the Bardeen-Petterson description. Figure 1.9 shows

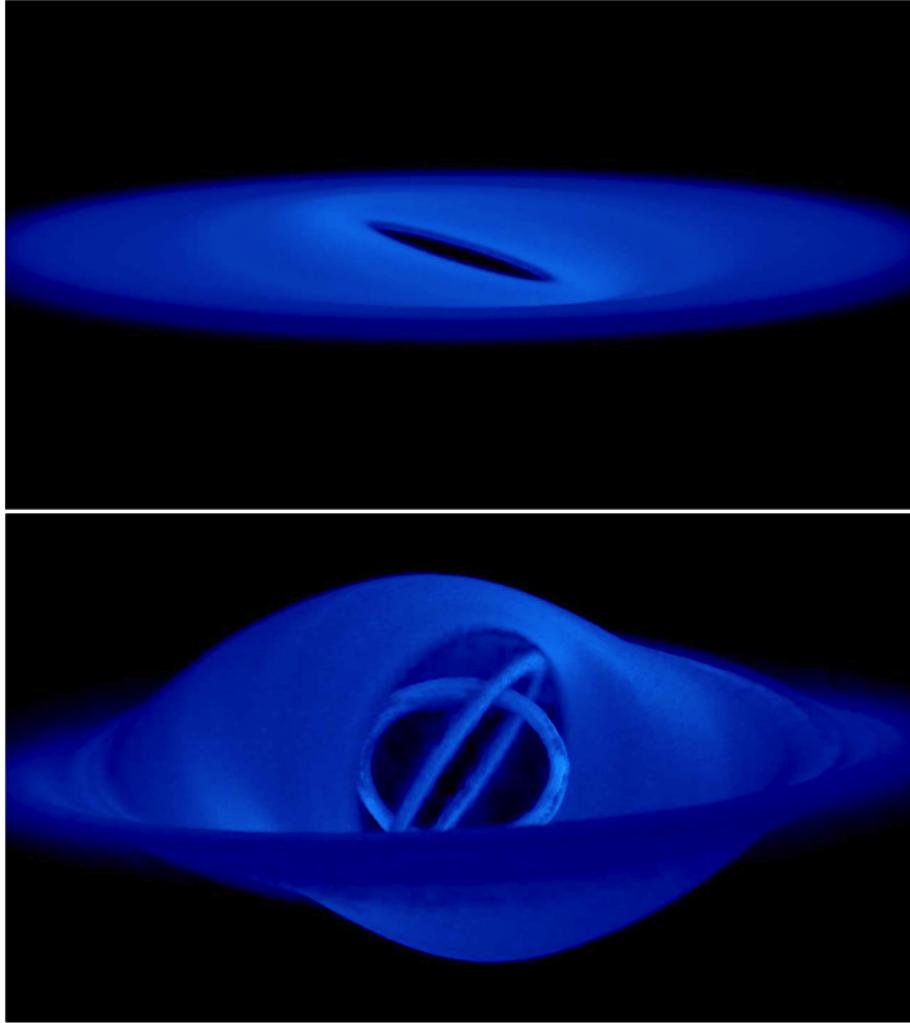


FIGURE 1.10: Density rendering of the first three-dimensional simulation exhibiting disc ‘tearing’ using the SPH code PHANTOM (Figures 1 and 2, Nixon et al., 2012).

their general findings for initial inclinations of 10° and 30° : inner regions are aligned to the black hole spin, the outer regions maintained their original misalignment and there exists a smooth transition between the two. Their simulations, even at low inclinations, showed no evidence of the oscillatory behaviour predicted by Ivanov and Illarionov (1997). While the source of this discrepancy is unclear, Nelson and Papaloizou (2000) suggested that ‘non-linear effects lead to the damping of these short wavelength features’, however later work by Lubow, Ogilvie, and Pringle (2002) found that the wavelengths of these oscillatory modes are comparable to the disc scale height. Specifically in the case of the larger inclination simulations ($\beta = 30^\circ$) Nelson and Papaloizou (2000) noted that:

“... a disc which is forced to maintain a non-linear warp due to severe misalignment will tend to break into two or more disconnected pieces, rather than maintain a smoothly warped structure ...”

Despite identifying a steepening of the disc tilt, Nelson and Papaloizou (2000) were unable to find clear evidence for disc breaking as in Larwood et al. (1996). By comparison, the

simulations by Nelson and Papaloizou (2000) made use of at most 2×10^5 particles (an order of magnitude larger than Larwood et al., 1996).

Higher resolution simulations of accretion discs at high inclination have been conducted using 2×10^6 particles by Nixon et al. (2012). In this case, the disc was initially inclined at 60° . Figure 1.10 shows two of these simulations; the upper panel with an initially low inclination (10°) shows results consistent with previous work. The lower panel shows an initially strongly misaligned disc (60°), where the disc was observed to break multiple times — the disc ‘tears’ into precessing rings of gas. These rings were generated in the same way as disc breaking; the Lense-Thirring torque generated by the rotation of the black hole overcomes the viscous torques holding the disc together. In the case of highly inclined discs, Nixon et al. (2012) found that this is sufficient to ‘tear’ successive rings of material off the disc (from the inner edge outwards). These rings precessed effectively independently, given that the precession rate is a strong function of radius. As they precessed, neighbouring rings developed opposing angular momenta and thus directly cancelled angular momentum, leading to a mass accretion rate that was an order of magnitude higher than for an aligned disc (Nixon et al., 2012). Disc breaking has subsequently been found in simulations of circumbinary discs (Nixon, King, and Price, 2013; Facchini, Lodato, and Price, 2013). Nixon et al. (2012) additionally explored enhanced accretion due to counter-rotating disc structures in the context of active galactic nuclei.

The location of where the disc breaks was estimated by comparing the Lense-Thirring torque and the viscous stresses that maintain the disc as a continuous, smooth structure. Nixon et al. (2012) derived that this should occur, for diffusive discs, around

$$R_{\text{break}} \lesssim \left(\frac{4}{3} |\sin \beta| \frac{a R}{\alpha H} \right)^{2/3} R_g, \quad (1.29)$$

where $R_g = GM/c^2$ is the gravitational radius. The above expression implies that the disc is more likely to break at a lower viscosity, high spin and large misalignments. However, Equation 1.29 assumes that the α parameter here describes all three components in the viscosity formulation — neglecting α_2 and α_3 . Doğan et al. (2015) examined the validity of this approximation in the context of a binary system. They found that it was less accurate for small α (as α_2 becomes significant). Taking into account the contributions from the vertical viscosity, they found agreement between the breaking radius in simulations and the full analytical expectation. Thus Equation 1.29 should only be used as a guide for the expected break radius.

Grid-based method

Numerical simulations that use grid based techniques have been restricted to small inclinations ($\lesssim 15^\circ$) and (in general) thicker discs. However they are able to generate viscosity through the MRI and include general relativistic effects. Fragile and Anninos (2005) and Fragile et al. (2007) conducted some of the first of these simulations, constructing a disc

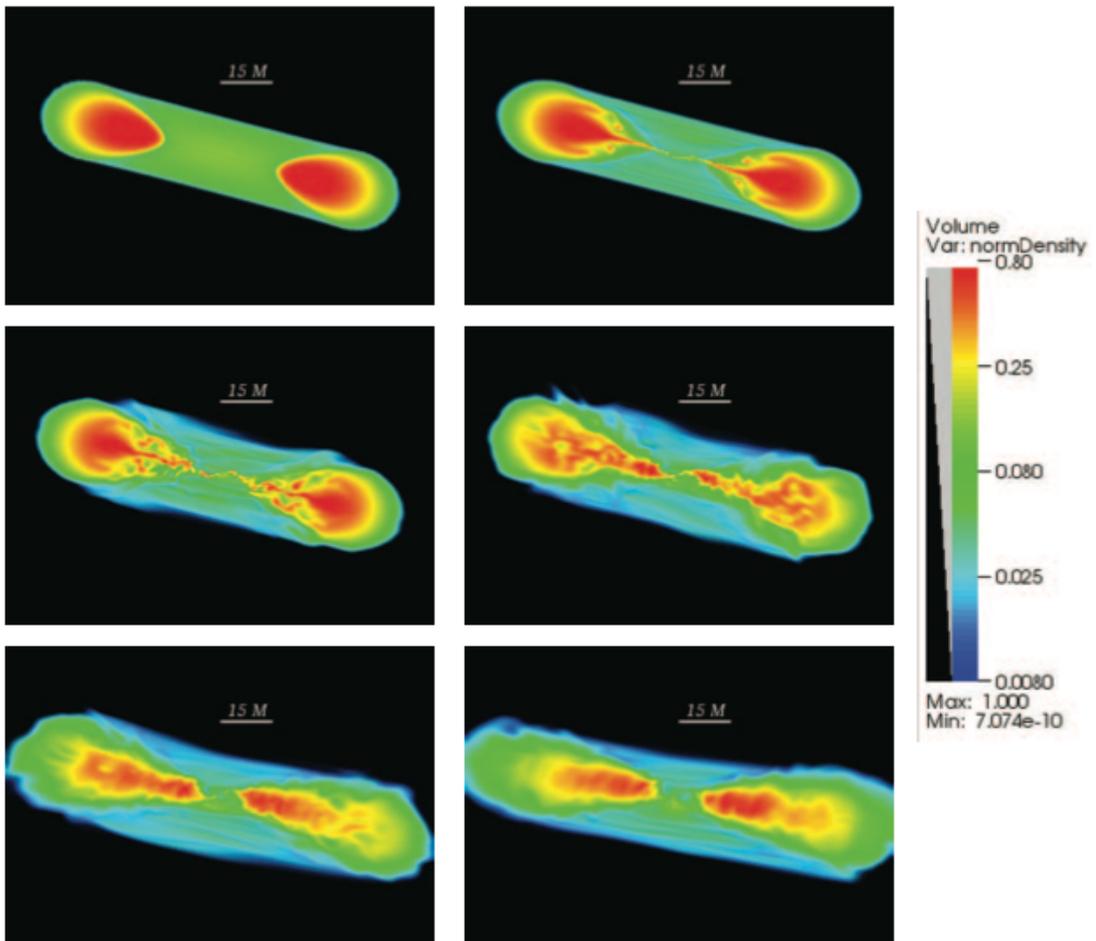


FIGURE 1.11: Volume visualisation of the density at 0, 1, 2, 4, 7 and 10 orbits of a torus tilted at 15° . This disc was simulated with an MHD grid-based code and showed evidence of an oscillatory inner region (Figure 2, Fragile et al., 2007).

inclined at 15° with aspect ratio $H/R \sim 0.1$, nested with poloidal magnetic field loops. Figure 1.11 represents the typical disc evolution of such a simulation, identifying a non-zero tilt at the inner edge. Fragile et al. (2007) also suggested that there is evidence of an oscillatory disc profile (see Figure 1.11, fourth panel) however this was not a steady profile as predicted by Lubow, Ogilvie, and Pringle (2002) and these simulations did not show evidence of the Bardeen-Petterson alignment in the inner regions. Importantly, their disc also displayed solid body precession (where the entire disc precesses with a shared precession frequency), behaviour that served as inspiration for geometric models of quasi-periodic oscillations (discussed in § 1.4).

In contrast to this, Sorathia, Krolik, and Hawley (2013) found evidence of the Bardeen-Petterson Effect. They conducted two simulations to allow a comparison, where the only difference was that their first simulation was hydrodynamical and the second was fully magnetohydrodynamical. To simplify the more complicated MRI driven simulation, a full GR treatment was avoided and a post-Newtonian approximation was employed. In both simulations there was evidence of alignment in the inner regions, but stronger and faster

alignment occurred for the MRI simulation. These simulations suggested that alignment was independent of the viscosity prescription used, but that modelling of the general relativistic terms (which is approximated when using a post-Newtonian potential but exact in the previous works) may be responsible for the different profiles.

The most recent simulations that employ full GR and MRI have confirmed the results that have not found Bardeen-Petterson alignment: Zhuravlev et al. (2014) and Morales Teixeira et al. (2014) conducted separate simulations of tilted accretion discs, but neither found evidence of the Bardeen-Petterson Effect in the prograde case. In line with analytical predictions (e.g. Figure 6 of Lubow, Ogilvie, and Pringle, 2002), there was evidence of (at least partial) alignment in the retrograde case. The most recent of these simulations, by Morales Teixeira et al. (2014), was noted for being the thinnest simulation of its type to date, with an aspect ratio of $H/R \sim 0.08$ that was comparable to the effective α generated in the simulation. The authors noted that although there was no evidence of Bardeen-Petterson alignment, there were indications of solid body precession.

A common theme in the discussion presented by Sorathia, Krolik, and Hawley (2013) and Morales Teixeira et al. (2014) was that because the stresses generated by the MRI were not isotropic, they cannot be consummate with the ‘isotropic α ’ description. The lack of isotropy in α was demonstrated from their simulations by showing the horizontal and vertical α components were not equal (Sorathia, Krolik, and Hawley, 2013), and measuring the local value of α as a function of position in the disc (Figure 14 Morales Teixeira et al., 2014). As discussed by Nixon (2015), as α is an averaged quantity the significance of these measurements on local scales is not clear. Additionally, it was possible that the simulations by Sorathia, Krolik, and Hawley (2013) may not have fully resolved the MRI (see section 2.2 of Morales Teixeira et al., 2014). This entire discussion also rests on the definition of ‘isotropic viscosity’, a phrase which appears to be interpreted ambiguously in the literature. From its original definition, isotropic viscosity references the *isotropic nature of the rate of shear* — i.e. isotropic viscosity does not imply isotropic torques (see the discussion in Nixon and King, 2015). From this definition it follows that the viscosity components $\nu_{1,2,3}$ are not the same ($\nu_1 \neq \nu_2 \neq \nu_3$). However, this appears to be the hypothesis investigated by Sorathia, Krolik, and Hawley (2013), Morales Teixeira et al. (2014) and others.

A similar study considering the evolution of an inclined disc was conducted by Krolik and Hawley (2015). As with previous simulations, the viscosity in their simulation was derived entirely from the MRI driven turbulence. To reduce the computational cost of their simulation, they chose to make use of a post-Newtonian approximation to model the black hole rotation effects. As discussed in Sorathia, Krolik, and Hawley (2013), this allowed them the freedom to use a black hole spin parameter that was as high as they liked because these effects scale linearly. Krolik and Hawley (2015) thus choose $a \geq 1.05$ so that the effects of precession could be observed over a short simulation time (25 orbits at $R = 10$). This made their simulation almost identical to Sorathia, Krolik, and Hawley (2013), except that they had more mass in the outer regions to combat difficulties found by the previous

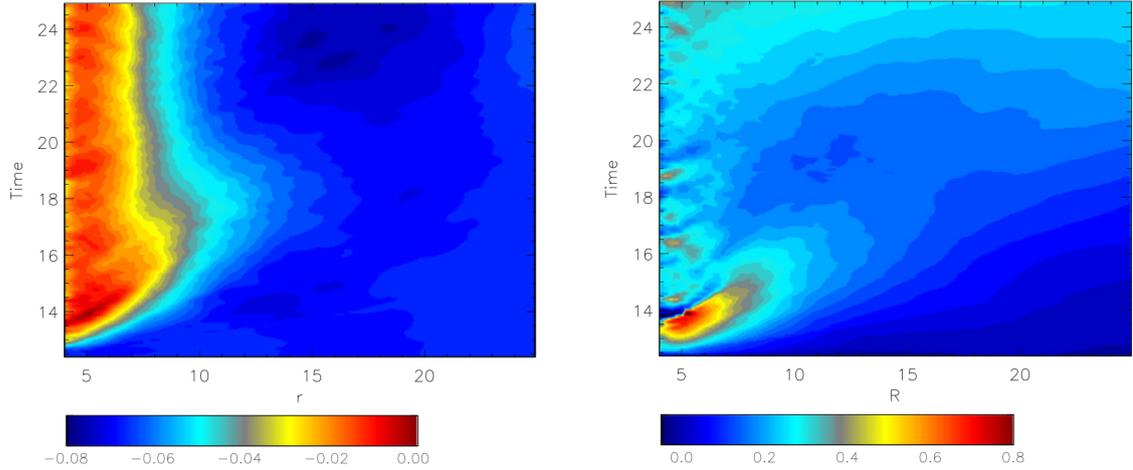


FIGURE 1.12: Left: the evolution of the disc tilt as a function of radius, scaled in radian/ π (the negative reflects an arbitrary definition). Right: the evolution of the disc twist as a function of radius, scaled in radian/ π . In both panels, the disc property is scaled by colour across the figure and evolves from bottom to top. Adapted from Figures 2 and 3 of Krolik and Hawley (2015).

authors. Shown in Figure 1.12, their simulation showed strong evidence of the Bardeen-Petterson Effect, with the inner disc aligned and connected by a smooth transition to the outer, misaligned disc. Although inclined at only 12° , the authors estimated that the disc should break by considering the diffusive analytical disc breaking criteria (Equation 1.29). This discrepancy about whether the disc should break or not was used to strengthen the argument that the α viscosity used in SPH simulations is not equivalent to a full, MHD turbulence driven accretion disc description. This discrepancy is addressed in detail in Chapter 4.

Liska et al. (2017) conducted the most recent numerical investigation of a misaligned disc around a black hole. These grid simulations were notable because of their extremely high resolution (up to ~ 1 billion cells) despite being restricted to ~ 110 orbits at the inner edge at $12.5R_g$. Each simulation started with a torus inclined at 30° around a black hole with $a = 0.9375$ with either a weak ($\beta = 100$) or strong magnetic field. In either case, Liska et al. (2017) identified a jet launched perpendicular to the inner torus, expansion of the outer edge and solid body precession of the torus. In the weak field simulations, the inner edge of the torus remained misaligned to the black hole spin such that the jet was also misaligned and precessed with the torus. In the strong field case, the inner edge of the torus was pulled into alignment with the black hole spin and was thus parallel with the relativistic jet that was produced. Although the torus in these simulations was shown to be precessing (and indeed, the behaviour observed was invoked as a method of generating quasi-periodic oscillations, see § 1.4), the resolved simulations show swift alignment and the rate of precession decreases significantly after ~ 150 - 200 orbits at the inner edge of their moderately resolved simulations (their Figure 1, panels e and f).

1.3.7 Warped discs summary

Although many authors have contributed to the investigation of warped accretion discs, discrepancies in the literature remain. We summarise these into three main points:

First, discrepancies exist between the analytical and numerical studies concerning the Bardeen-Petterson Effect. At moderate inclinations, some simulations clearly found the Bardeen-Petterson Effect (e.g. Krolik and Hawley, 2015) while some do not (e.g. Zhuravlev et al., 2014) despite using similar methods to simulate the disc. At smaller inclinations, oscillations in the tilt profile of a wave-like disc have been predicted by both Ivanov and Illarionov (1997) and Lubow, Ogilvie, and Pringle (2002), but are yet to be shown in numerical simulations (e.g. Nelson and Papaloizou, 2000).

Second, higher inclination simulations have revealed disc tearing behaviour but only in the diffusive regime. Is it possible for wave-like discs to also show this behaviour? If so, do the rings have similar properties (e.g. thickness) or are these properties constrained by the disc properties? Is enhanced accretion from tearing discs generic to all accretion discs?

Finally, the validity of the α model has been challenged with grid based simulations by Sorathia, Krolik, and Hawley (2013), Morales Teixeira et al. (2014) and Krolik and Hawley (2015). While their simulations showed important small scale behaviour of the disc viscosity, this challenge appears to be based on an ambiguity in the literature. If correct, this may call into question the purely hydrodynamical approach that is commonly adopted by particle methods.

1.4 Quasi-periodic oscillations

Direct observational evidence of tearing accretion discs around black holes is currently not possible (but see Casassus et al., 2015, for observations of a potentially strongly inclined, broken proto-planetary disc). Thus to understand the connection between tearing discs in nature and the theoretical work previously discussed, indirect observational features must be considered. Black hole low mass X-ray binaries (LMXBs) are an ideal candidate for this as there is observational and theoretical motivation to suggest that the accretion disc in these systems are tilted (discussed in § 1.3). Hence, we should expect that the Lense-Thirring precession can influence the evolution of the disc in these systems. Additionally, LMXBs are part of a group of objects that generate quasi-periodic oscillations (QPOs). As discussed below, observations and theoretical work on QPOs strongly suggest that Lense-Thirring precession is responsible for driving these currently unexplained phenomena.

QPOs are detected as oscillations in the X-ray, optical and infrared lightcurves from a number of astrophysical objects, including LMXBs, active galactic nuclei, cataclysmic variables and ultra-luminous X-ray sources (Gribsin, Feldman, and Plagemann, 1970; Patterson, Robinson, and Nather, 1977; Joss, Avni, and Rappaport, 1978; van der Klis et al., 1985; Strohmayer et al., 2003; Gierliński et al., 2008). They represent rapid oscillations in the

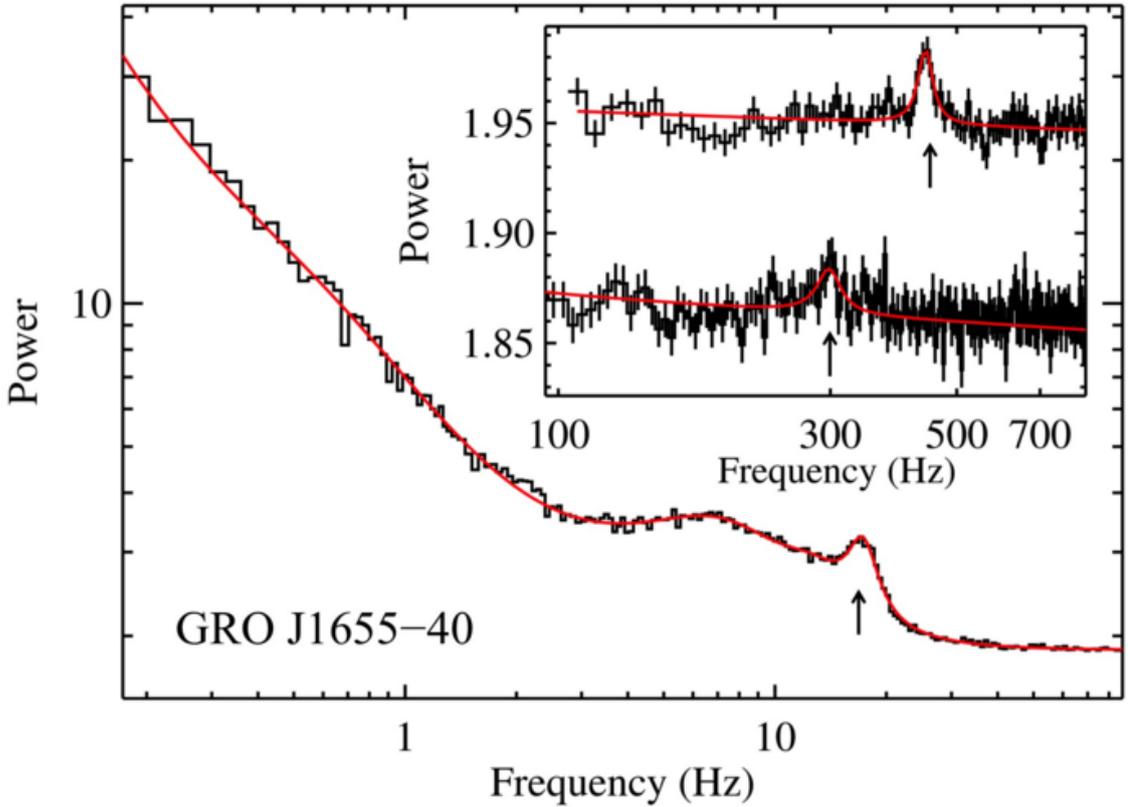


FIGURE 1.13: Power density spectra of GRO J1655-40 identifying a type-C low-frequency QPO at 18Hz and a high-frequency pairing at 300Hz, 450Hz (Figure 9 of Belloni and Stella 2014, see also Motta et al. 2014b).

flux with an amplitude of up to 20% (in type-C QPOs, but can be as low as 10% in type-B). Given this nature, they are found by generating a power density spectrum (PDS) from the observed lightcurve and often multiple lightcurves are stacked together. An example is shown in Figure 1.13, with observed QPOs identified with arrows by their characteristic ‘bump’ in the PDS.

Two clues as to the origin of QPOs are fairly well established. First, because their frequency is so rapid these oscillations must be occurring within $60R_g$ of the compact object (Axelsson, Hjalmarsdotter, and Done, 2013). Second, as QPOs are observed in such a wide variety of sources, they must be generated from the accretion disc itself as this is the only common feature. This then implies that the mechanism that drive QPOs must originate in the innermost part of an accretion flow, where relativistic effects are strongest. Additionally, QPOs are strongest around stellar mass black holes in X-rays. Although QPOs observed in non-relativistic objects (such as those in cataclysmic variables) are observationally similar to those from black hole and neutron star binaries, they are likely to be caused by a different mechanism (Kluźniak, W. et al., 2005).

The nature of QPOs is further complicated by observations indicating that they fall into distinct types. Broadly, QPOs are split into high- and low-frequency with subcategories within the low-frequency group (type-A, -B and -C). In any given LMXB, both or either of

these types may be observed individually or simultaneously. Figure 1.14 shows a summary of the QPOs observed in the LMXB J1550-564 from 1999 to 2015, indicating some of the possible combinations of simultaneous QPOs. Because high- and low-frequency QPOs have such distinct observational characteristics, they are often assumed to originate from different mechanisms and are thus treated separately. In black hole LMXBs, only type-C QPOs are currently associated with Lense-Thirring precession.

The measurement of a precise frequency (e.g. for each QPO in Figure 1.13, the width of the bump) has fundamental implications for the structure of an accretion disc. Current disc theory models the accretion disc as a smooth, continuous whole. But how does a smooth, flat structure produce a rapid variation in the flux? How does it do so at two (or three) different frequencies simultaneously? The implication is that there must be some inhomogeneity in the region of the disc from where QPOs are sourced. Indeed, some of the first explanations of QPOs considered ‘blobs’ of material above and below the disc plane (in neutron star LMXBs, Stella and Vietri, 1998). We will show that disc tearing provides a natural method to prevent the continuous, smooth nature predicted in discs as gas inclined to the rotation of the black hole spin breaks into discrete structures due to the rotation induced torque. We briefly summarise the relevant points of current observations to better understand the potential connection between tearing discs and QPOs, focusing on those from LMXBs which have a black hole as the compact object.

1.4.1 Observations of high frequency QPOs

High-frequency QPOs (HFQPOs) in LMXBs occur with frequencies $\gtrsim 60$ Hz and have some analogies with HFQPOs from accreting neutron stars (where they were first detected, van der Klis, 2006). When observed, they do not always take the same frequency (e.g. Figure 1.14) but sometimes do — this suggests that the preferred frequencies are set by some fundamental parameters of the system (Belloni, 2010). Owing to their high frequency, HFQPOs are often associated with the orbital frequency near the inner edge of the accretion disc. HFQPOs are difficult to detect in black hole LMXBs, and as such have been found in only a few sources (Belloni, Sanna, and Méndez, 2012). Deeper comparison of these QPOs with their neutron star (NS) counterparts show that they are probably not the same as in NS sources (Belloni and Motta, 2016; Motta et al., 2017).

In black hole binary systems, HFQPOs can also be seen in pairs with particular ratios between the observed frequencies. Remillard et al. (2002b) found that these ratios are constant to within about 15% as the source changes flux (but this is not true for neutron star LMXBs, Remillard and McClintock, 2006). Often reported is the ‘3:2’ ratio, however Belloni, Méndez, and Homan (2005) show that this ratio sits somewhere between ≈ 1.3 -3.0 and other ratios are reported (e.g. GRS1905+105, Morgan, Remillard, and Greiner, 1997; Strohmayer, 2001). Although HFQPOs can be observed simultaneously with their low-frequency counterparts, (Remillard and McClintock, 2006) suggest that it is difficult to demonstrate correlation between the two. By contrast, Stella and Vietri (1998) and Stella

and Vietri (1999) have shown a strong connection between type-C and HFQPOs. Using this relationship, Motta et al. (2014b) and Motta et al. (2014a) were able to use simultaneously observed QPOs to estimate properties of the black hole. More recently, Belloni, Sanna, and Méndez (2012) have discussed the possible connections between HFQPOs and the different types of LFQPOs.

1.4.2 Observations of low frequency QPOs

Low-frequency QPOs (LFQPOs) are found at frequencies generally less than 30 Hz. These QPOs are observed to be stable and to vary on timescales of minutes (Muno et al., 2001). As depicted in Figure 1.13, they have a large amplitude and can be narrow (with a coherence of $Q > 10$, where Q is measured as the ratio of the frequency and the variance of the QPO Remillard and McClintock, 2006). Phase lags of LFQPOs suggest that analogous to neutron star systems, there are three subtypes; type-A, -B and -C (Casella, Belloni, and Stella, 2005). These three subtypes are further identified by their timing properties, coherence, centroid frequency, energy bands and where they are observed in the hardness-intensity diagram (discussed in § 1.4.3). We note that there is observational suggestions of further, anomalous QPOs (e.g. Remillard et al., 2002a) but do not discuss them here.

Type-A LFQPOs

Perhaps the most poorly understood of the LFQPOs, type-A were originally classified by Wijnands, Homan, and van der Klis (1999) and have since been observed in only a handful of systems (as of 2016 about 10 have been found; Motta, 2016). They are found between ≈ 6.5 -8Hz (Belloni and Stella, 2014). As they are broader and weaker than the other QPOs, they require the data to be averaged over a long time period to be identified (e.g. Homan et al., 2001).

Type-B LFQPOs

Type-B LFQPOs are observed during the transition from the hard to soft state and are associated with relativistic jets (e.g. Motta et al., 2015). They are often interpreted as an indication of a transition in the disc as they follow type-C LFQPOs by less than a second (Motta, 2016). In cases where transitions between the subtypes of LFQPOs occur, type-B are always involved (e.g. B \rightarrow C or B \rightarrow A, Nespoli et al., 2003). Additionally, while the disc goes from hard to soft (and vice versa) there are often multiple transitions from and to type-B QPOs involving both type-A and -C QPOs. In contrast to the other LFQPOs, observations of type-B are fitted with a Gaussian shape (because they exhibit jitter, Belloni, 2010). These LFQPOs are currently restricted to ≈ 0.8 -6.4Hz (Belloni and Stella, 2014) but are only observed between ≈ 4 -6Hz at high flux (Belloni, 2010). They are characterised most strongly by where they appear in the hardness-intensity diagram and their strong peak (see Figures 1.15 and 1.16). Finally, the strength of type-B QPOs shows a tight

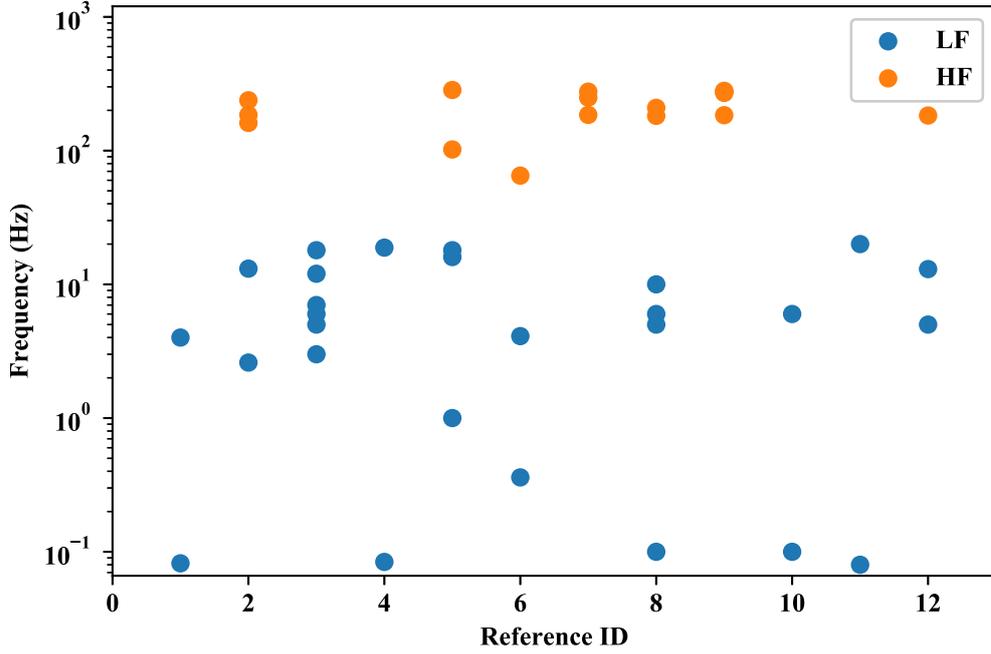


FIGURE 1.14: Summary of QPO observations using RXTE data for J1550-564. The red points represent low-frequency and the blue high-frequency (simultaneous observations are not indicated). The Reference ID and authors are found below in Table 1.1.

TABLE 1.1: Corresponding references for Figure 1.14.

Reference ID	Authors	Frequencies (Hz)	Type
1	Cui et al. (1999)	82mHz - 4 Hz	Low
2	Remillard et al. (1999)	185, 161, 238	High
	Remillard et al. (1999)	2.6 - 13.1	Low
3	Wijnands, Homan, and van der Klis (1999)	5 - 7	Type-A
	Wijnands, Homan, and van der Klis (1999)	3, 6, 12, 18	Type-B
4	Sobczak et al. (2000)	0.084 - 18.81	Low
5	Homan et al. (2001)	1 - 18, 16 - 18	Low
	Homan et al. (2001)	102 - 284	High
6	Kalemci et al., 2001	65	High
	Kalemci et al. (2001)	0.36 - 4.1	Low
7	Miller et al. (2001)	185, 249 - 276	High
8	Remillard et al. (2002a)	6	Type-A
	Remillard et al. (2002a)	5 - 6	Type-B
	Remillard et al. (2002a)	0.1 - 10	Type-C
	Remillard et al. (2002a)	182 - 209	High
9	Remillard et al. (2002b)	184, 276	High
	Remillard et al. (2002b)	280, 270	High
10	Rodriguez et al. (2004)	0.1 - 6	Type-C
11	Li et al. (2013)	0.08 - 20	Low
12	Motta et al. (2014a)	5, 13	Low
	Motta et al. (2014a)	183	High

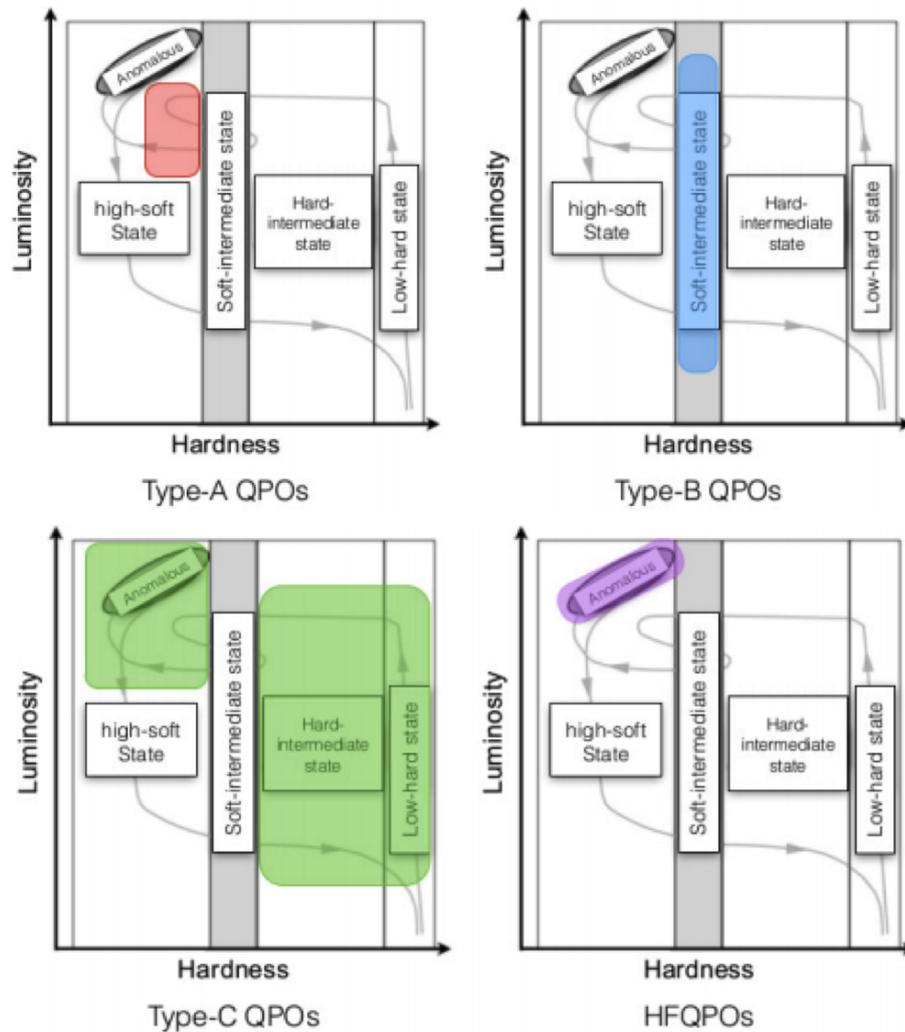


FIGURE 1.15: Evolution of LMXBs on the hardness-intensity diagram with the locations of each particular type of QPO (Figures 1 and 7 of Motta, 2016; Belloni, 2010, respectively).

dependence on viewing inclination that is different to the dependence on type-C QPOs (Motta et al., 2015).

Type-C LFQPOs

Type-C LFQPOs are the most commonly observed and well documented of the LFQPOs and were first defined by Wijnands, Homan, and van der Klis (1999). They are observed up to ≈ 30 Hz and are more similar to type-A than to type-B LFQPOs (e.g. the disc state they can be observed in, Figure 1.15). They are found by their strong and narrow peaks and are seen to drift in frequency and strength on timescales of a few days. In contrast to the other LFQPOs, type-C are observed in multiple states including the low hard state, the high soft state and the ultra-luminous state, discussed below (Motta, 2016).

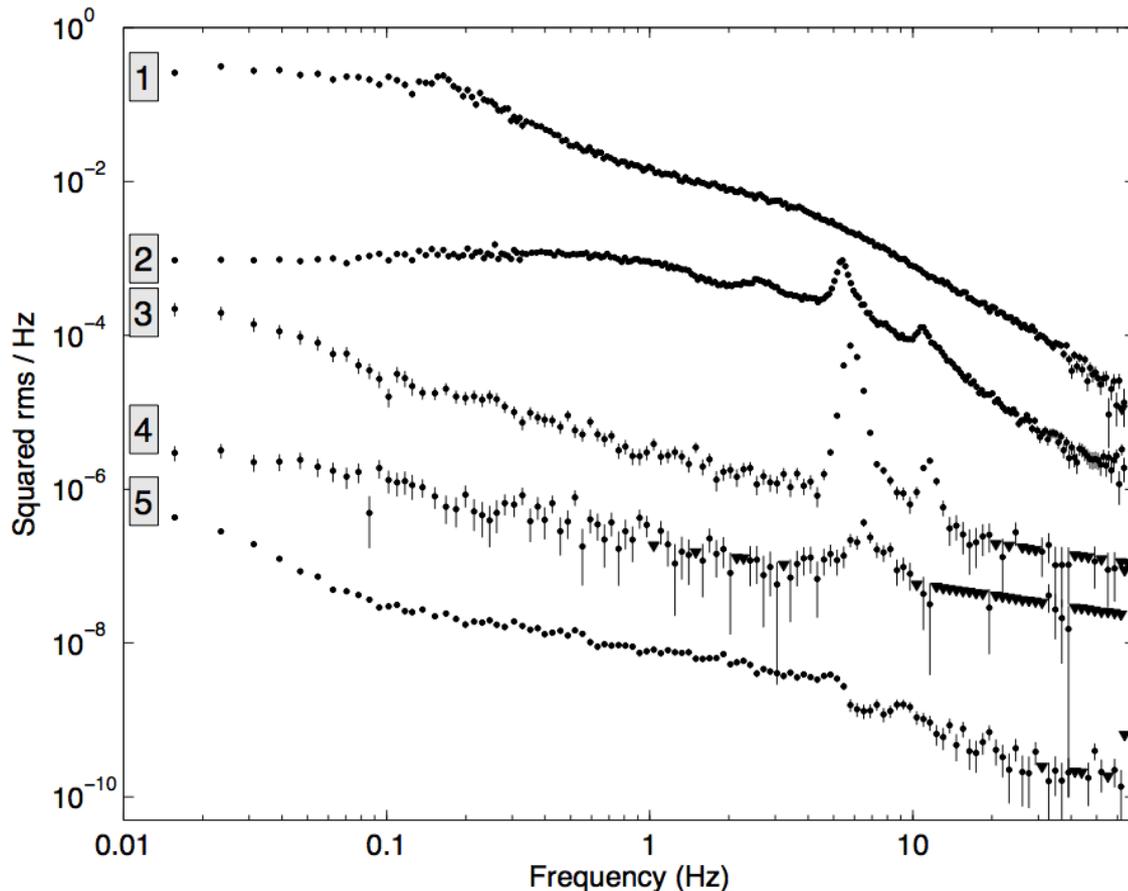


FIGURE 1.16: Example PDS showing the different LFQPOs and power law slopes. PDS 1 corresponds to the LHS, PDS 2 to the high flux LHS with type-C QPOs present, PDS 3 to the SIMS state with type-B QPOs, PDS 4 to the HSS with type-A QPOs present and PDS 5 to the softest state. These data are taken from the 2002/2003 outburst of GX 339-4 and the PDS are shifted in power so that they can be distinguished (Figure 3, Belloni, 2010).

1.4.3 Connection to disc states

The relationship between different types of LFQPOs is further understood by their relationship to the disc state, indicated by the position on the hardness-intensity diagram (HID, also known as a ‘q’ or ‘turtle-head’ diagram, Homan et al., 2001). The ‘hardness’ of the system is measured as the ratio between counts in two particular energy bands, and when plotted against the total counts from the system reveals a particular evolutionary path over the course of a several months. Figure 1.15 shows a schematic of the evolution of a LMXB and the right hand panel of Figure 1.18 an observed HID. In the HID, the high-soft state exhibits a soft spectrum with thermal emission from the accretion disc and the low-hard state a hard spectrum driven by Comptonised or synchrotron self-Compton emission (Done, Gierliński, and Kubota, 2007). We note that although the generic path through the HID is the same for many objects, the actual path taken differs between objects and sometimes even for the same object during different outbursts.

Transition through the different states is understood to be constrained by the mass accretion rate and is defined mostly by the position in the HID (while the HIMS to SIMS and SIMS to HSS transitions are defined by features in the PDS). The cycle begins with the object in quiescence in the low hard state (LHS) in the lower right region of the HID. The luminosity increases, tracking the object up the right hand side of the HID. This movement marks the highest mass accretion rate and is followed by the transition into the hard/soft intermediate states (HIMS/SIMS), travelling to the left on the HID. In the softest, most luminous state (hard soft state, HSS) the hardness is observed to fluctuate, with the object zig-zagging across the left hand side of the HID. As the luminosity then decreases the hardness increases, evolving back through the HIMS at a lower luminosity, then returning to the LHS. The difference in PDS as the object evolves is documented in Figure 1.16, where each different LFQPO has a distinct observational signature. As indicated by the coloured regions in Figure 1.15, the different sub-types of QPOs are only observed in particular places in the evolution.

Type-C QPOs are predominantly observed in the LHS, HIMS and the softest, most luminous part of the HSS (indicated by the green regions in Figure 1.15). As the flux increases in the LHS, the characteristic frequencies of the LFQPOs also increase (Belloni and Stella, 2014). Although it is not clear what initiates the decrease in hardness and movement into the HIMS, it is here that emission from the thermal disc component is observable (Belloni, 2010). Although the HIMS is often seen as an extension of the LHS, due to the blurred transition between the two and similar LFQPOs, they are distinct.

Type-B LFQPOs are seen exclusively in the narrow SIMS region (indicated by the blue region in Figure 1.15). This state is characterised by lower variability in the HID and only a small increase in the softness in the spectrum. This transition also roughly corresponds to the launching of relativistic jets (crossing the ‘jet line’, Belloni and Motta, 2016).

Type-A QPOs are observed in a state slightly softer than the SIMS (indicated by the coral region in Figure 1.15). As the object zig-zags through this region transitions are observed between the three types of LFQPOs but are thought to pass through type-B. Additionally, most HFQPOs are observed around this hardness at higher luminosity.

1.4.4 QPO models

Models attempting to explain the driving mechanism of QPOs are broadly grouped into geometric and instability models. The instability models invoke a range of oscillatory behaviours (e.g. see Tagger and Pellat, 1999; Titarchuk and Fiorito, 2004; Varnière, Tagger, and Rodriguez, 2012; Cabanac et al., 2010, for a few examples). However, these models are largely ruled out because they do not predict the observed dependence on viewing inclination of LFQPOs (Motta et al., 2015). Here the two main geometric models are introduced followed by an interpretation using tearing discs.

The relativistic precession model

This model associates the fastest of the paired HFQPO with the orbital frequency, the slower HFQPO with the periastron frequency and the LF type-C QPO with the nodal precession at a particular radius (Stella and Vietri, 1998; Stella and Vietri, 1999). With these relations, this model effectively assumes that each type of QPO is sourced from a self luminous blob of material at a known radius. This is modelled physically as a test particle at a characteristic radius (Stella and Vietri, 1998). As a result, if three frequencies are simultaneously observed they can be used to solve for the spin and mass of the black hole as well as the radius that the QPOs are sourced from — as used by Motta et al. (2014b) to precisely measure the mass and spin of the black hole GRO J1566-40. In the case where only two QPOs are observed, this model may be combined with dynamical mass estimates to constrain the spin of the black hole (Motta et al., 2014a) or even to place limits on the mass and spin (Ingram and Motta, 2014). Notably, the results of this model appear to predict a lower spin than the values that are observed using other methods (although these have their own issues, for example the Fe $K\alpha$ method assumes that the disc and the black hole are aligned). This model has recently been extended by Franchini, Motta, and Lodato (2016) to assume that the LFQPO is produced at the innermost edge of a hot, precessing inner flow (similar but distinct from the rigid disc model, see below). This interpretation assumes that the rate of precession of this extended flow can be approximated by the rate of precession at the inner edge (true for a narrow flow with no radial extent).

The rigid precession model

The rigid precession model focuses on the mechanism driving type-C LFQPOs and is built on the framework introduced by the truncated disc model (concerning disc states, transitions and correlations between spectral and timing properties). It was introduced by Done, Gierliński, and Kubota (2007) and built on by Ingram, Done, and Fragile (2009) and Ingram and Done (2011) (but see Ichimaru, 1977; Esin, McClintock, and Narayan, 1997; Poutanen, Krolik, and Ryde, 1997). This model is distinct from the relativistic precession model because it assumes that the frequency driving the LFQPO is generated from the precession of a radially extended structure with a shared precession frequency, weighted across the structure. Figure 1.17 depicts the inner and outer disc structure invoked in this model. Both regions are misaligned to the black hole spin, with the inner region exhibiting solid body precession. As the inner structure precesses, it alternately casts bright and dark patches on the outer disc (the rainbow region in Figure 1.17). The patches appear to orbit on the outer disc with the precession rate of the inner structure, projecting to an observer a *bright-dark-bright-dark* pattern in the flux. The rainbow colour scheme on Figure 1.17 indicates the expected red- and blue-shift that an observer will see due to the rotation of the outer disc (Ingram and Done, 2011). As the size of this inner precessing structure decreases, the frequency drifts to higher values (Ingram, Done, and Fragile, 2009). Within the context of the truncated disc model, Ingram and van der Klis (2013) describe the

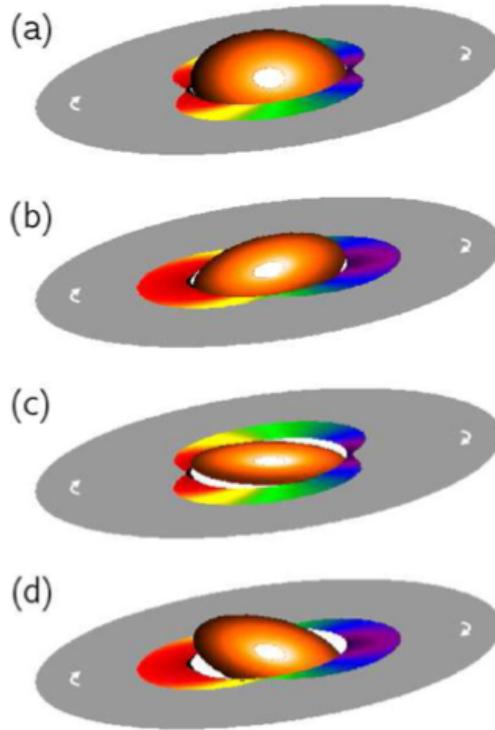


FIGURE 1.17: Description of the double peaked lightcurve generated using the truncated disc model (Ingram et al., 2016, Figure 11). The grey outer disc is alternately brightened/darkened due to geometric obscuration from the inner, red disc. The rainbow colour scheme indicates the blue-/red-shift experienced by the observer. Here, the orange inner disc corresponds to the grey ‘X’s in Figure 1.18.

source of the broad-band noise as changes in the radiation from the inner region which is ultimately controlled by changes in the mass accretion rate from the outer disc.

The evolution of the truncation radius of the thin disc is closely tied to the evolution through the HID, and ultimately controls the size of the rigidly precessing inner thick (hot) disc. In the LHS it is implied that the disc is cool, geometrically thin and is truncated at some radius larger than the innermost stable orbit (ISCO). Within this disc exists a hot, thick flow that is able to precess as a solid body (Ingram, Done, and Fragile, 2009). As the object evolves up the right-hand side of the HID, the mass accretion rate increases while the truncation radius moves inwards (Done, Gierliński, and Kubota, 2007). This leads to increased illumination of the inner hot flow by cooler photons from the disc, leading to steeper spectra (Ingram and van der Klis, 2015). The transition from hard to soft occurs when the truncation radius of the outer disc is commensurate with the ISCO (Done, Gierliński, and Kubota, 2007). Additional confirmation of this model has been recently achieved by Ingram et al. (2016) using Fe-line modelling: continuum photons that are reflected off the outer disc are distorted by relativistic effects, as one side of the disc has material approaching the observer and the other side receding from the observer. This should lead to rocking of the Fe-line between blue- and red-shifted over a characteristic timescale. Ingram et al. (2016) found this rocking and showed that i) it must be driven by

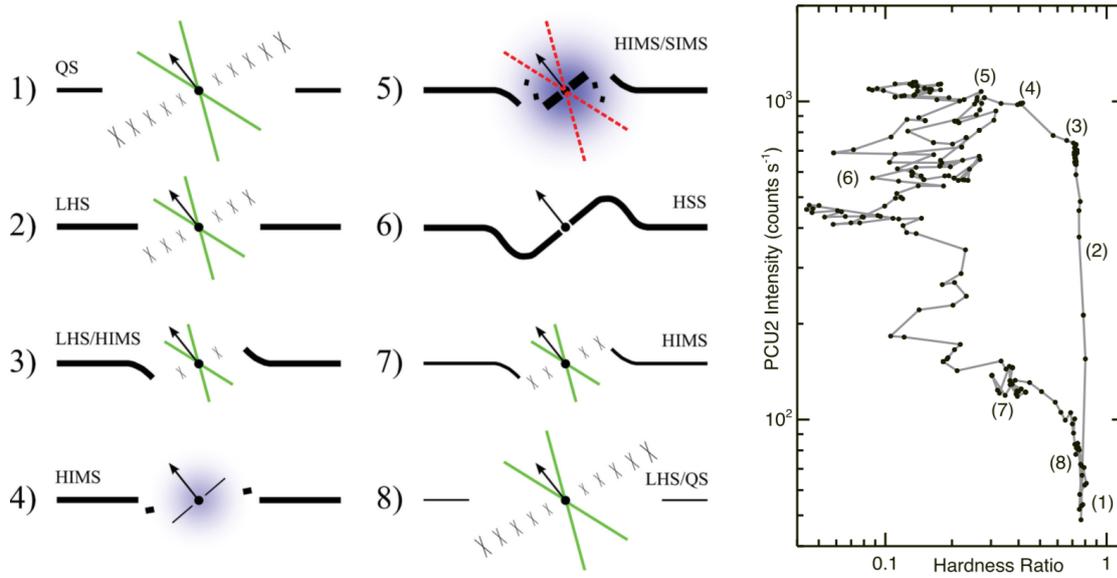


FIGURE 1.18: Evolution in the hardness intensity diagram (right) explained using tearing processes in the disc (left) (Figure 1, Nixon and Salvesen, 2014). Here the thickness of the black line describes the mass accretion rate experienced through the disc (not the disc thickness), with the continuous lines showing a Shakura-Sunyaev disc and the discontinuous a torn disc. The green lines indicate a steady jet, the red a transient jet and the black circle with an arrow in the direction of its spin. The inner components evolve from hot, low density gas (grey ‘X’s) to a corona (purple shading). Here, the grey ‘X’s correspond to the orange inner disc in Figure 1.17.

Lense-Thirring precession, ii) any changes in the line profile must be caused by geometric variations in the structures generating this feature and iii) this feature was created by two distinct oscillating lightcurve signals, confirming the geometry proposed in Figure 1.17. Additionally, as this model depends entirely on the geometry of the system, it naturally explains the observed strong dependence on viewing inclination (Schnittman, Homan, and Miller, 2006; Heil, Uttley, and Klein-Wolt, 2015; Motta et al., 2015).

Despite agreement with observations, this model has difficulties due to its phenomenological origin: there is no mechanism described to generate this inner region that is disconnected from the outer disc or to explain the dynamic evolution of its outer radius. Additionally, it is not clear how the transition from a thin, cool accretion disc to a hot thick inner region co-exist in such a short radial extent. As depicted in Figure 1.17, the inner region must be sufficiently thin to allow for two bright patches on the outer disc but their model predicts it should be hot and thick (or there must be ‘a very large misalignment between the disc and inner flow’, Ingram et al., 2016). A further difficulty is that the history of the truncated model dates back to a single simulation that is only conducted for one quarter of a precession (see Figure 1.11 and Figure 16 of Fragile et al., 2007), and so the long term evolution of this model is yet to be validated by numerical simulations.

A disc tearing interpretation

Although not strictly a geometric model, disc tearing has also been used to interpret the development and evolution of LFQPOs. Figure 1.18 shows a potential evolution pathway through the HID motivated by a misaligned disc undergoing tearing, proposed by Nixon and Salvesen (2014). In this model, type-C QPOs are generated in the LHS and LHS/HIMS by the warping of the inner edge of the inwardly accreting outer disc. Transition into the HIMS is represented by tearing from the outer disc and the formation of a hot corona. As some material is able to cool and form an inner disc, gas is able to form structures down to the ISCO, launching jets and marking the transition into the SIMS. Over time, the disc is able to cool completely and form a warped, connected disc as it moves into the HSS. The system returns through the HIMS as this disc accretes, leaving behind a low density, hot gas component aligned with the black hole spin that is ready for another cycle.

1.4.5 QPOs summary

Despite having been observed for more than 45 years, the mechanism that drives QPOs is still not fully understood. Careful analysis of these observations has revealed a number of features that must be addressed by any model that seeks to explain how they arise:

- The relationship between the mechanisms that drive different types of QPOs: observations suggest that they are caused by different phenomena. However, they are all sourced from a similar region in the disc and in the case of LFQPOs, transitions occur between each type. Although this suggests different phenomena for each type of LFQPO, it may also be possible that they are caused by related processes.
- The particular locations where each sub-type of QPO appears in the HID and hence how they relate to the current disc state: e.g. type-C LFQPOs may be found when the disc is very hard or soft, but do not appear in the SIMS. What does this imply about the disc structure in these regions and its dynamic evolution?
- The analytical relationship between the observed frequencies and fundamental properties like the black hole mass and spin: the relativistic precession model sometimes allows an analytic approach to be used to measure these properties, but the measurements are not often in agreement with other observations. Contrastingly, the rigid precession model assumes a different relationship but to date this model has not yet been used to measure properties of the disc. As QPOs are sourced from the innermost region of the disc, there must be a strong relationship between the frequencies observed and the properties of the black hole.
- The observed frequencies imply discrete disc structures: current disc models predict that the disc should be continuous from the outer edge through to the inner edge,

and it is currently unclear how a continuous disc is able to produce multiple observable frequencies. QPOs fundamentally imply that the inner region of the disc is inhomogeneous rather than continuous and smooth.

Of the models that have been put forward to consider LFQPOs, we consider the relativistic precession model and the truncated disc model. Each of these models associates type-C LFQPOs with the nodal precession due to frame dragging around a spinning black hole. While the relativistic precession model associates this with a characteristic radius (which also drives the high frequency QPO pairing), the truncated disc model connects this with the averaged precession of an extended torus. While both models have their own advantages, neither are able to address all of the above questions.

1.5 Geometrically thick flows

Until this point we have focused on geometrically thin accretion flows around low mass black holes. Thick flows or tori may be expected in scenarios like tidal disruption events around supermassive black holes, where the relative inclination of the encounter to the black hole spin has a significant effect on the subsequent evolution. Tidal disruption events occur when the trajectory of a star (typically modelled as a solar mass type star) takes it close enough to a supermassive black hole for the gravity of the black hole to dominate the local gravity of the star. The star is disrupted and shredded into a stream of material as it passes through pericenter, with approximately half the mass of the star remaining bound to the black hole. The debris stream continues to orbit, returning to the black hole at a predictable rate. Through (presumably) shock interactions, eventually the debris gas circularises into a disc/torus (depending on the cooling rate) and accretes onto the black hole.

The inclination of the orbit of the incoming star relative to the black hole spin is not predetermined. As a result, both apsidal and nodal precession are relevant to the circularisation process. Discussed below, including the effects of apsidal precession causes the head of the stream to intersect its own tail, leading to faster circularisation. However, including nodal precession causes the orbits to precess around, preventing the self crossing of the stream and delaying circularisation. We thus focus on the scenario where the original orbit of the star has a small inclination to the black hole spin such that nodal precession is not significant. We also note that the Papaloizou-Pringle instability (introduced below) is independent of any inclination to the black hole spin.

1.5.1 Light curves from tidal disruption events

Tidal disruption events occur when a star enters the tidal radius, R_t of the black hole. This radius is defined by the star (M_*) and black hole masses as (Phinney, 1989)

$$R_t = R_*(M/M_*)^{1/3}, \quad (1.30)$$

with the depth of the stars penetration measured with the parameter $\beta = R_t/R_p$ (such that tidal disruption events occur when $\beta > 1$), where R_p is the pericentre radius. In the case that a star is disrupted, half of the star mass remains bound to the black hole (Lacy, Townes, and Hollenbach, 1982; Rees, 1988). Phinney (1989) found the expected mass accretion rate onto the black hole at late times by considering the specific energy of the star. Before disruption, the range of specific energies of the gas is narrow and does not vary greatly from the specific energy at the centre of mass of the star. During the disruption, the orbits of the gas are ‘squeezed’ together (Lodato, King, and Pringle, 2009), destroying the hydrostatic balance of the star and broadening the range of specific energy of the gas. As a result, after the disruption that material that has negative specific energies is bound to the black hole. Using the Keplerian period of the gas T , the energy of this gas can be expressed as (e.g. Lodato, King, and Pringle, 2009)

$$E = -\frac{1}{2} \left(\frac{2\pi GM}{T} \right)^{2/3}. \quad (1.31)$$

Assuming that the returning debris loses both angular momentum and energy faster than it orbits the black hole, the material accretes quickly and observational features like flares are expected (and observed, Gezari et al., 2012), with initial accretion rates predicted to be super-Eddington (Evans and Kochanek, 1989; Loeb and Ulmer, 1997). Following Lodato, King, and Pringle (2009), the mass accretion rate \dot{M} may be estimated as (Rees, 1988; Phinney, 1989)

$$\dot{M} = \frac{dM}{dE} \frac{dE}{dT} = \frac{(2\pi GM)^{2/3}}{3} \frac{dM}{dE} T^{-5/3}. \quad (1.32)$$

When the star has a uniform energy distribution this reduces to a $t^{-5/3}$ dependence on the mass accretion rate, and if the mass accretion rate controls the luminosity the same dependence is expected in the light curve.

The usefulness of this prediction is demonstrated in Figure 1.19, where the grey line shows the $t^{-5/3}$ prediction as compared to the observed light curve in multiple optical bands for the disruption PS1-10jh. Amongst others, Gezari et al. (2012) have used this profile as a fit to identify tidal disruption events. Figure 1.19 also demonstrates that this $t^{-5/3}$ profile is only true at late times, with the luminosity initially rising to a peak. Subsequent to the original derivation presented above, Lodato, King, and Pringle (2009) showed that the predicted steepness of the curve also depends on the internal structure of the star before it is disrupted, namely the nature of the energy distribution. Guillochon and Ramirez-Ruiz

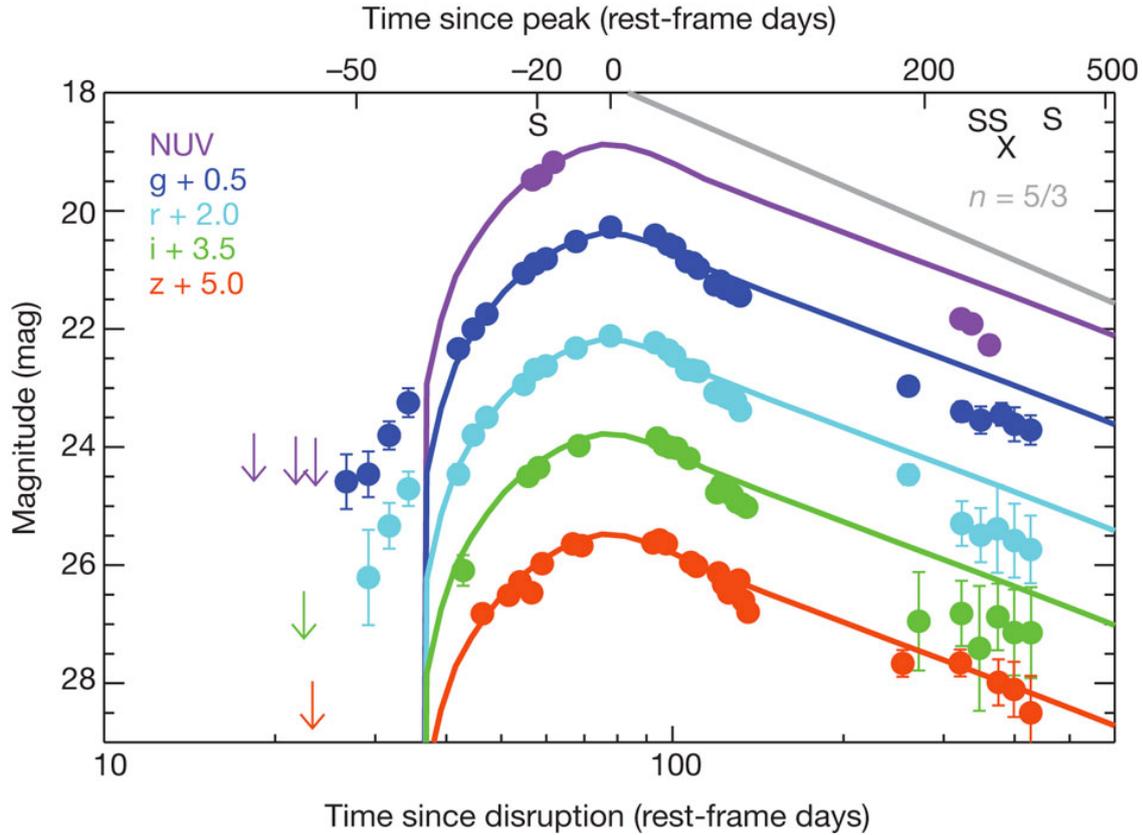


FIGURE 1.19: Optical light curves of the tidal disruption event PS1-10jh. Here the discrete points represent the observations, the solid lines a least squares best fit of a numerical model of a tidal disruption event to each band and the grey line represents the $t^{-5/3}$ decay prediction (Figure 2 from Gezari et al., 2012).

(2013) further demonstrated that the bound core of the star is more likely to survive during shallower encounters (i.e. with a shallower impact parameter β), so the penetration factor is also important to determining the subsequent evolution.

The above derivation inherently assumes that the stellar debris around the black hole is able to circularise, form a disc (requiring that it efficiently cools) and viscously accrete before the returning stellar debris arrives. This may be a difficult task, as it requires the timescales between the circularisation, viscous accretion and radiative cooling processes to co-operate (Evans and Kochanek, 1989). For example, a disc will form if the time taken to circularise is less than the time for the gas to accrete. Of the circularised gas, a thin disc will only form when the cooling timescale is also shorter than the circularisation timescale (as assumed by Cannizzo and Gehrels, 2009; Shen and Matzner, 2014). If instead the gas cannot cool faster than it circularises, a geometrically thick torus evolves (e.g. Loeb and Ulmer, 1997; Coughlin and Begelman, 2014). In either case, the rate of accretion is governed by the development of the MRI. As stars have an initially weak magnetic field, this may take quite a few orbits to become established. For a ratio of magnetic field and gas pressure of $\beta \sim 100$, we estimate it should take almost 3 orbits for parameters given by Bonnerot et al. (2016) (see § 6.7, Equation 6.9).

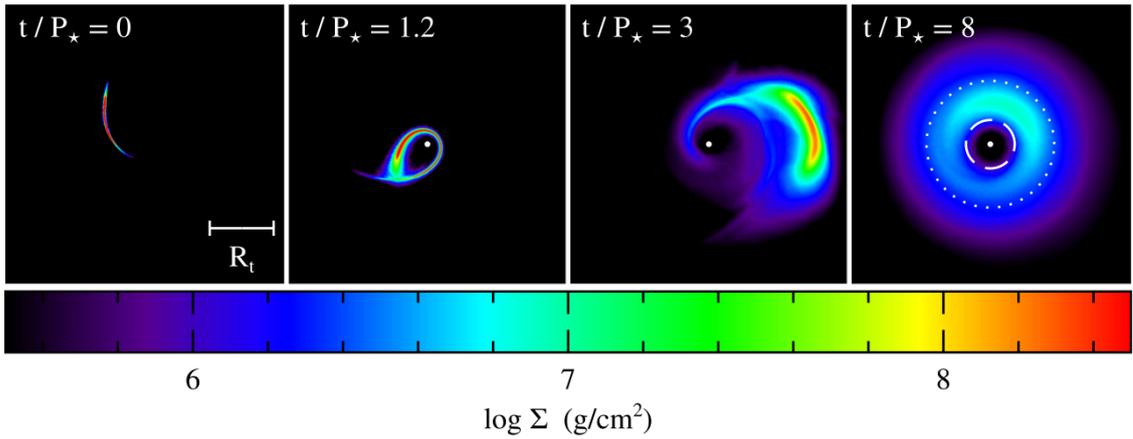


FIGURE 1.20: Surface density evolution of a TDE where $P_{star}=2.8\text{hr}$ and the gas is not allowed to cool efficiently. The white point denotes the black hole, the dashed circle is the circularisation radius and the dotted circle is semi-major axis of star (Figure 5, Bonnerot et al., 2016).

Piran et al. (2015) summarised one of the major difficulties when considering tidal disruption events: ‘why do we see only 1% of the expected energy?’. Predictions suggested that the luminosity from these encounters should be super-Eddington, where the luminosity is driven by the accretion of the remnant (e.g. Evans and Kochanek, 1989; Komossa, 2015). However recent observations of optical disruption events demonstrated the observed temperature and bolometric luminosity are much lower than this (Piran et al., 2015). Additionally, if shocks are the dominant mechanism used to circularise the gas during a disruption, they must be capable of dissipating a lot of energy and should thus be quite luminous themselves. Piran et al. (2015) argued that the shocks during the circularisation process rather than the accretion of the remnant itself were the major source of the luminosity, with a model that fits an observed peak luminosity. By contrast, Ayal, Livio, and Piran (2000) suggested that the shocks may heat gas enough to unbind it, resulting in a less massive structure around the black hole after the disruption. Poutanen et al. (2007) proposed that wind loss may reduce the mass accretion rate at any time during the remnant’s evolution, perhaps explaining the difference between the prediction and observation. Finally, Shen and Matzner (2014) suggested that if a disc forms it is able to expand beyond the disruption radius in order to conserve angular momentum, with the outer material eventually accreting faster than the stellar feedback rate. The difference between the observed and predicted luminosity may also be an indication of optically thick outflows; as material accretes at a super-Eddington rate, winds or jets may be powered, reducing the luminosity observed — consistent with observations that show jets (e.g. Gezari et al., 2012). These jets may carry energy away from the black hole (Komossa, 2015).

1.5.2 Simulations of TDEs

The spatial scales and mass ratios involved in tidal disruption events make them particularly challenging to simulate. In terms of spatial scales, the star initially is compact but

after disruption is stretched into an isolated, low density stream. This range of length-scales is well suited to SPH, so many simulations have been conducted using this method. However the typical mass ratios (say a $1M_{\odot}$ star encountering a 10^6M_{\odot} black hole) combined with a parabolic orbit adds more difficulty; these ratios lead to very small, fast orbits around the black hole which are computationally expensive. In an attempt to mitigate these difficulties, often the stars initial orbit is altered to be eccentric (e.g. Bonnerot et al., 2016; Hayasaki, Stone, and Loeb, 2013) or the mass ratio between the star and the black hole is decreased. In the former case, more debris naturally ends up on a bound orbit than in the parabolic case, such that it is easier for the star to circularise (it has to lose less energy) and the disc that forms is more massive than would be expected.

Recent simulations by Bonnerot et al. (2016), Hayasaki, Stone, and Loeb (2013) and Shiokawa et al. (2015) explored the circularisation of the torus including the effects of relativistic precession from the black hole rotation. Figure 1.20 displays a representative simulation demonstrating the effect of apsidal precession, where the gas was additionally not able to cool efficiently. As the star was disrupted around the black hole, apsidal precession of the orbits lead to self-crossing of the streams (second panel). Shocks that were induced at the stream crossings heat the gas, increasing its thermal energy and leading to faster circularisation of the gas. The resulting torus was ‘thick’ (in the sense that the aspect ratio $H/R \sim 1$), with a cross section shown in Figure 1.21. Additionally the material was highly eccentric. Using a post-Newtonian description of the relativistic effects, Shiokawa et al. (2015) found that a relatively small difference in the apsidal precession angle of 10° is enough to lead to stream intersections of 90° . Although including apsidal precession led to faster circularisation of the gas, if the incoming gas was on a misaligned orbit nodal precession may prevent self-crossing of the streams entirely and thus delay circularisation (Hayasaki, Stone, and Loeb, 2013). Guillochon and Ramirez-Ruiz (2015) suggested that the effect of nodal precession would even delay any luminous signal from a TDE for years in the case of a 10^7M_{\odot} black hole. Recent work by Franchini, Lodato, and Facchini (2016) assumed that in the case that a misaligned disc can form, it would globally precess and may produce a modulation in the light curve. Using a 1D code (similar to that used in Lubow, Ogilvie, and Pringle, 2002) they simulated the evolution of a wave-like disc with typical TDE parameters; e.g. the disc extends from the innermost stable orbit to the circularisation radius. Their simulations confirmed that such discs can globally precess but over time align due to viscous dissipation of the induced warp. While the initial precession of such discs may cause quasi-periodic features in the light curves from TDEs, such a signal has only been tentatively found in one TDE (Swift J1644, Reis et al., 2012).

Recent simulations have also focused on the evolution of the stream after the initial disruption. Coughlin et al. (2016) showed that as the stream moved away from the black hole it was able to collapse into clumps due to self-gravity. As the debris returned the mass accretion rate was punctuated by these clumps, a potentially observable phenomenon (e.g. the G2 cloud, Saitoh et al., 2012). Guillochon and McCourt (2017) conducted the first simulations of tidal disruption events including magnetic fields for a variety of encounters.

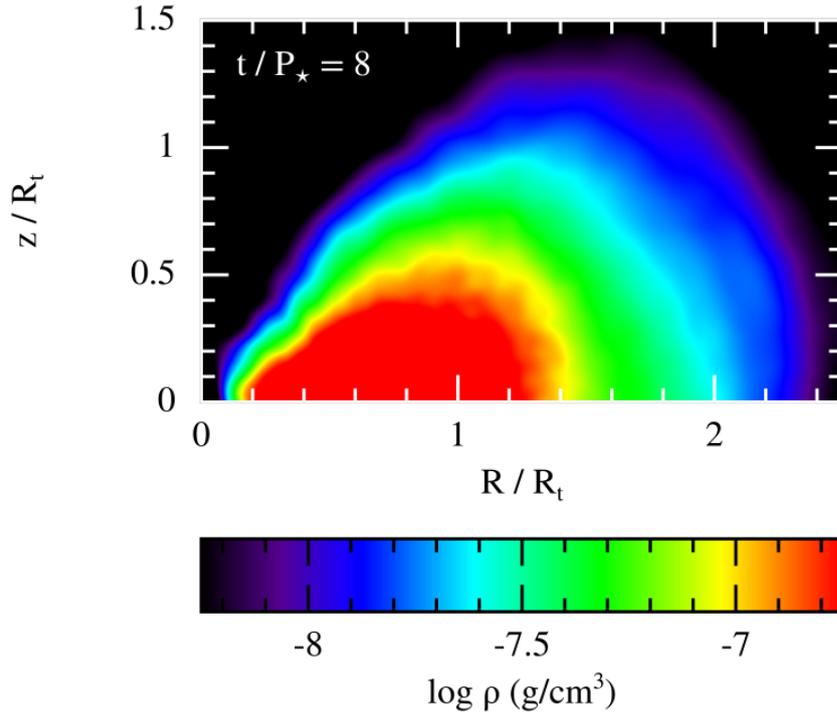


FIGURE 1.21: Cross section of a thick torus that forms after a TDE, including apsidal precession and when the gas is not allowed to cool efficiently ($R_t = 0.5\text{AU}$, Figure 8, Bonnerot et al., 2016).

They found that the stellar debris that falls back is mostly strongly magnetised by the time it reaches the black hole. Additional simulations of magnetised disruptions by Bonnerot, Rossi, and Lodato (2017) have focussed on the stream evolution, where they showed that the presence of strong magnetic stresses could cause gas to accrete ballistically rather than circularise into a disc. In the case of gas that inefficiently cools, they again found the formation of a geometrically thick torus.

1.5.3 The Papaloizou-Pringle instability (PPI)

Loeb and Ulmer (1997), Coughlin and Begelman (2014), Ulmer (1999), and Strubbe and Quataert (2009) argued that a geometrically thick structure like that shown in Figure 1.21 is the fate of a tidal disruption remnant, rather than a geometrically thin accretion disc. However, the subsequent evolution of the torus may be complicated by the Papaloizou-Pringle instability (PPI, Papaloizou and Pringle, 1984). This hydrodynamic instability requires i) the torus have well defined inner and outer boundaries (i.e. it is pressure supported and not accreting from the inner edge) and ii) that the profile of the specific angular momentum is shallow. As we will show in Chapter 6, it is possible that tori formed from tidal disruption events meet these criteria and are thus are unstable to the PPI.

Goldreich, Goodman, and Narayan (1986) found that the PPI is caused by communication between waves generated inside the torus from its boundaries: both the inner and outer

edges generate two surface waves, one travelling in the direction of the fluid flow and one travelling in the opposing direction. When the reverse wave from the inner edge has the same phase velocity as the forward wave from the outer edge, they are able to communicate angular momentum and energy. This communication ultimately results in a decrease of angular momentum at the inner boundary and an increase of angular momentum at the outer boundary. This redistribution of angular momentum is a run-away effect, with the growth of this instability occurring on the orbital timescale at the density maximum of the torus (Papaloizou and Pringle, 1984).

Damping of the PPI is easily achieved by accretion through the inner boundary (Blaes, 1987). As material accretes onto the central object through this boundary, angular momentum is advected out of the inner region. When the rate at which this occurs is faster than angular momentum can be communicated outwards in the torus the instability damps (Blaes, 1987). Evidence of the PPI is thus not expected in simulations investigating the MRI in thick discs ($H/R \sim 1$, effectively tori), as they are accreting and hence prevent the PPI. The PPI can also be avoided by specific choices of angular momentum profiles (see Papaloizou and Pringle, 1985; Glatzel, 1987).

1.5.4 Growth rate of the PPI

Here we summarise the method used to determine how quickly the PPI develops in a torus with a given radial extent. As we shall see, the growth rate of this instability is determined purely by the location of the inner and outer boundaries of the torus. The dimensions of the torus are in turn a product of how the torus is formed. As we are focussing on a torus formed after a tidal disruption event, we also calculate the growth rate of the PPI in a torus with the same physical parameters as the torus shown in Figure 1.20. A comparison of this growth rate to the growth rate of the magneto-rotational instability motivates Chapter 6.

A full analytical description of the PPI can be found in Blaes and Glatzel, 1986 with a simpler analysis presented by Pringle and King (2014). They assume a cylindrical flow of incompressible fluid, neglect any self gravity or z dependence and assume an angular velocity profile of the form

$$\Omega(R) = \Omega_0 \left(\frac{R_0}{R} \right)^2, \quad (1.33)$$

where R is the cylindrical radius and Ω_0 is the angular velocity at the radius of maximum density R_0 . The frequency ω of a mode with azimuthal wavenumber m is then a solution of the following equation (e.g. Blaes and Glatzel, 1986),

$$\frac{(\omega + m\Omega(R_-))^2 + mg(R_-)/R_-}{(\omega + m\Omega(R_+))^2 + mg(R_+)/R_+} = \left(\frac{R_+}{R_-} \right)^{2m} \frac{(\omega - m\Omega(R_-))^2 + mg(R_-)/R_-}{(\omega - m\Omega(R_+))^2 + mg(R_+)/R_+}, \quad (1.34)$$

where the effective gravity is defined according to

$$g(R) = \frac{GM R_0}{R^3} \left[1 - \frac{R}{R_0} \right], \quad (1.35)$$

with M as the mass of the central object and R_+ and R_- denoting the outer and inner edge of the torus, respectively.

Recalling that this instability is driven by interactions from the boundaries, it follows that the solutions to Equation 1.34 are determined purely by the choice of the inner and outer radii. Two of the four solutions to Equation 1.34 are real, representing stable modes (Blaes and Glatzel, 1986). The two remaining solutions correspond to unstable modes, one growing and one decaying. For the growing unstable mode, the growth rate is found from the imaginary component of the frequency, $Im(\omega)$, and depends on the wavenumber m .

For the torus shown in Figure 1.20, the radial extent corresponds to a dominant mode of $m = 1$ and the associated growth rate is $\omega \gtrsim 0.5\Omega$. Although this is smaller than the growth rate for the MRI (0.75Ω , Balbus and Hawley, 1998), the PPI develops and saturates within ten or so orbits at R_0 . In the scenario of a tidal disruption event the initial magnetic field in the torus is quite weak, requiring ≈ 3 orbits to establish accretion through the MRI. In this case then, it may be possible for the PPI to grow alongside the MRI and develop over-densities before the MRI has established itself.

1.5.5 Simulations of the PPI

As the PPI manifests itself in radially narrow and wide tori differently, simulations are naturally grouped by the type of torus simulated. In a thin torus, the PPI is demonstrated by the formation of over-densities (historically referred to as ‘planets’) and in a wide torus with a spiral pressure wave. Initial progress was made by Hawley (1987) simulating a radially slender torus with a grid based code. He demonstrated the growth of isolated modes in such a torus and the structure of the mode was as predicted by theory (Goldreich, Goodman, and Narayan, 1986). Further to this, he discussed how the instability transports angular momentum through the formation of over-densities, causing the torus to spread radially. In each case, his simulations ended when the torus was broken up by the presence of these planets.

In the case of radially wide tori, saturation (when the density perturbations stop growing further) occurs in the form of spiral pressure waves (Blaes and Hawley, 1988; Hawley, 1991). Zurek and Benz (1986) investigated the redistribution of specific angular momenta in wider tori using SPH, parameterising the specific angular momentum profile with $l(r) \propto r^q$. They showed that the damping criteria of accretion for the PPI is equivalent to redistributing the angular momentum profile. Figure 1.22 summarises their findings in terms of the evolution of the specific angular momentum profile; tori that are stable to the PPI show no significant evolution while those that are unstable show a dramatic change. Here a stable torus is represented with $q \gtrsim 0.25$ and an unstable torus with $q \lesssim 0.25$ (in Figure 1.22 the stable

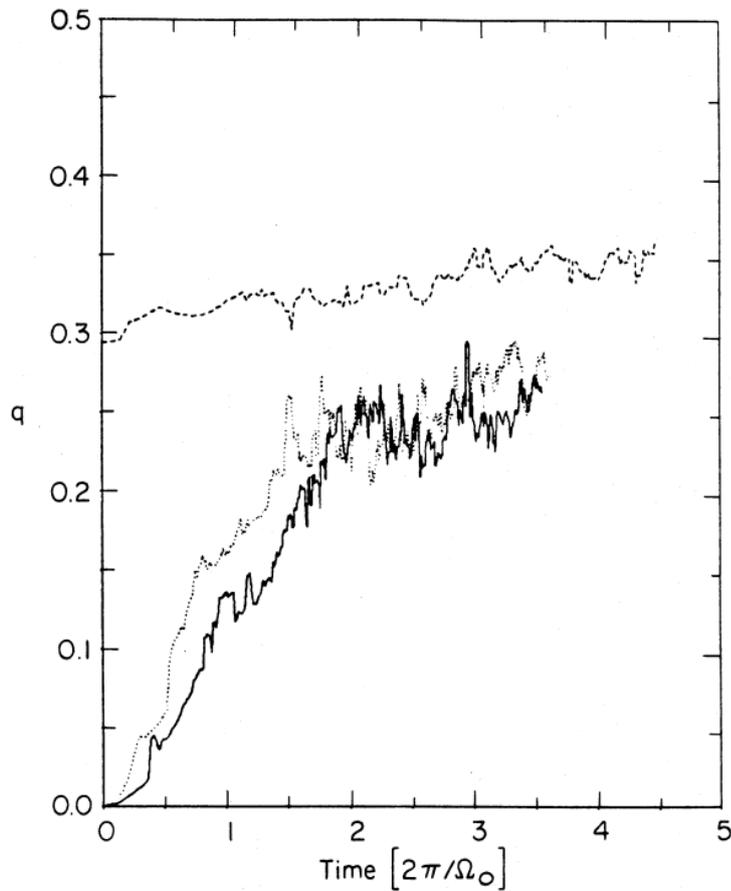


FIGURE 1.22: Power law-fit (where $l = Br^q$) of the specific angular momentum profile. Here the solid and dotted lines are tori started with $l \propto r^0$ and the dashed line has $l \propto r^{0.3}$ (Figure 3 of Zurek and Benz, 1986).

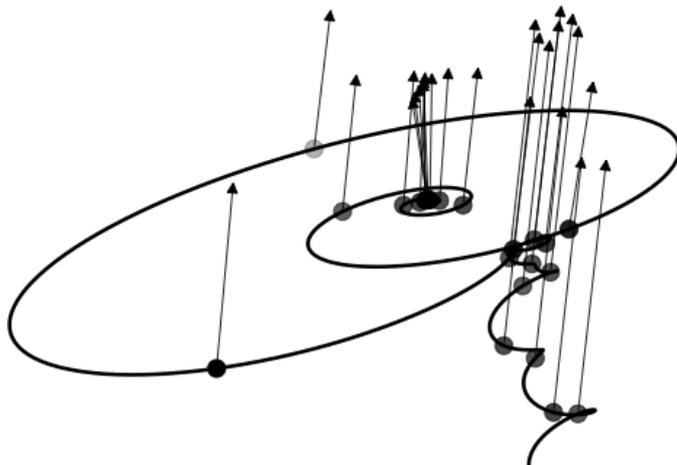


FIGURE 1.23: Figure 14 of Mewes et al. (2016), demonstrating the possible movement of a central black hole in response to the development of the PPI in a tilted torus. At each timestep the dot represents the location of the black hole and the scaled vector the direction of its spin vector. Each timestep is equally spaced to show that the movement of the black hole is much greater at the end of the simulation, when the PPI is saturating.

torus has $q \approx 0.3$ and the unstable with $q = 0$). The tori that start in a configuration that is unstable to the PPI tend to a stable configuration within about two orbits. These simulations confirmed the original specific angular momentum stability criteria ($q = 0.25$, Papaloizou and Pringle, 1984).

The PPI has also been revived in the context of gravitational waves and even QPOs. Mewes et al. (2016) conducted numerical simulations of the tilted torus that formed after the merger of two black holes. Immediately after this merger the torus was able to have a mass up to 16% of the new black hole mass and is regarded to be geometrically thick. In the case that such a torus was unstable to the PPI, its development of over-densities would cause the black hole to move in response to the PPI. An example of this possible movement is shown in Figure 1.23, and the authors postulated that such motion may be observable as quasi-gravitational waves (i.e. gravitational waves that are not the direct result of a merger, Mewes et al., 2016). Further to this, Kiuchi et al. (2011) suggested that the PPI may form in tori made during γ -ray bursts. Finally, although the PPI has also been invoked to understand QPOs (§ 1.4), there are a number of open questions about this interpretation. Dönmez (2014) suggested that the over-densities that formed in a torus around a black hole resulted in an oscillatory mass accretion rate, leading to an oscillatory light curve. However, observations suggest that some QPOs are strongly a function of the black hole spin while the PPI is not affected by this (from both the analytical work summarised above and from simulations like Mewes et al., 2016). The strong observed geometric dependence ruled out LFQPOs (Motta et al., 2015), but frequencies that are commensurate with HFQPOs have not yet been generated with this method (Dönmez, 2014). Finally, because the PPI only occurs once in a torus (i.e. once the angular momentum is redistributed and the torus is accreting, the torus is stable to the PPI), this interpretation offers no explanation for the consistent nature of QPOs or their ability to drift in frequency during observations.

The effect of magnetic fields on the development of the PPI in thick tori has recently been examined by Bugli et al. (2017). They consider two simulations, one including the MRI and the other purely hydrodynamical and where both are initialised with random perturbations (rather than seeding the dominant $m = 1$ mode). Their simulations suggest that the inclusion of the MRI completely inhibits the PPI from developing as accretion is driven immediately (for the ratio of gas to magnetic pressure of 100, i.e. a weak field).

1.5.6 Tidal disruption events and the Papaloizou-Pringle instability summary

The formation of a geometrically thick torus from a tidal disruption event is expected when the material is able to circularise faster than it is able to radiatively cool (e.g. Bonnerot et al., 2016). When this torus accretes before the rest of the stellar debris returns to the black hole, the light curve observed is expected to take the form of $t^{-5/3}$ (Phinney, 1989) consistent with observations (Gezari et al., 2012, e.g.). The assumption that the torus is

able to circularise and accrete quickly depends on competition between the viscous, cooling and circularisation timescales (Evans and Kochanek, 1989) and both a geometrically thin or thick structure is supported in the literature.

In the event that a torus does form and it is not able to cool efficiently, it may be susceptible to the Papaloizou-Pringle instability (Papaloizou and Pringle, 1984). In a torus with parameters expected after a tidal disruption event (i.e. a radially wide torus) the dominant $m = 1$ mode has a growth rate comparable to that of the MRI, but the PPI saturates after only a few orbits (e.g. Zurek and Benz, 1986). Additionally, as magnetic fields in stars are weak the initial strength of the magnetic field in the torus is expected to be small, such that the MRI will take a number of orbits to establish itself. Previous numerical simulations show that the development of the PPI drives angular momentum redistribution and suggest that it may be responsible for angular momentum transport (e.g. Hawley, 1987; Zurek and Benz, 1986). In the case that a tidal disruption remnant is susceptible to the PPI, it may thus be possible for angular momentum transport to be driven by the PPI before the MRI is established in the torus.

In the case of a highly inclined encounter, nodal precession may significantly delay the circularisation process and the above scenario may not occur (Guillochon and Ramirez-Ruiz, 2015). Additionally, strongly magnetised tori that are susceptible to the PPI may form during a γ -ray burst (although Wen and Dong, 2011, suggests that this does not prevent the PPI from occurring). We thus restrict the discussion of the PPI in tori from tidal disruption events to those which have low inclination encounters, are able to circularise quickly and have an initially small magnetic field, such as simulated by Bonnerot et al. (2016).

Chapter 2

Numerical methods

Both grid and particle based methods are commonly used for numerical simulations of astrophysical fluids. The difference in these methods stems from their interpretation of the governing equations of motion; particle methods are derived from the Lagrangian representation and grid based from the Eulerian. Practically this means that grid based simulations discretise the problem onto a grid of small volumes and consider the inflow and outflow of material between these volumes. In contrast, particle methods discretise the problem on to particles of fixed mass which carry the properties of the fluid as they flow. As highlighted in the literature review, each of these methods have particular advantages and disadvantages when applied to the problem of accretion discs.

Grid based methods are able to incorporate both magnetic fields to generate the magnetorotational instability (MRI, e.g. Fragile et al., 2007) and full general relativistic (GR, e.g. Morales Teixeira et al., 2014) effects. However, conservation of properties like mass, angular momentum and energy can be compromised when the flow of material is misaligned to the grid used to define the simulation (Robertson et al., 2010). This is especially important when considering the evolution of high inclination discs, as material that starts aligned to the grid will inevitably become misaligned. As a result grid based simulations of inclined accretion discs are generally restricted to small inclinations where this effect is small (e.g. Krolik and Hawley, 2015). The most recent grid based simulations with an inclined disc consider an inclination of 30° for ~ 110 orbits at the innermost edge of the disc (with $R_{\text{in}} = 12.5R_g$, Liska et al., 2017).

While grid based methods must make use of adaptive meshes or nested grids (called ‘turduckening’, Mewes et al., 2016), particle methods naturally have higher resolution in the inner regions of the disc where it is needed. Additionally, due to their formulation they conserve properties like the linear and total angular momentum exactly (as accurately as the time-stepping method allows). Importantly, conservation of these properties is independent of the inclination of the disc, making this method well suited to highly inclined discs. However, existing SPH simulations are not able to include full GR effects or magnetic fields (e.g. Nixon et al., 2012). Instead, they make use of a post-Newtonian approximation to include GR effects from the rotation of the black hole and use the ‘isotropic α ’ viscosity model.

We use smoothed particle hydrodynamics to conduct simulations of misaligned accretion discs, making use of its advantages for strongly misaligned flows. We use the publicly available smoothed particle magnetohydrodynamics code PHANTOM (Price et al., 2017).

In this chapter we introduce the fundamentals of SPH, the main features implemented in PHANTOM and the limitations we anticipate with this numerical technique. This chapter is based primarily on the extensive review Price (2012) and the companion paper to PHANTOM (Price et al., 2017). Extensions of PHANTOM specific to this work include the addition of a velocity correction for a rotating black hole (the equivalent of Equation 2.53, discussed in Chapter 3), developing analysis routines to measure azimuthally averaged properties and two setup routines (the first for warped discs and the second for tori). A 1D SPH code was written (from scratch) to demonstrate the importance of artificial viscosity (see the Sod shock tube example, Figures 2.2 and 2.3). The tests documented in Sections 2.2.4 and 2.2.5 have been published in Nealon, Price, and Nixon (2015) and additional code development is described in the chapter that it is relevant to.

2.1 Smoothed particle hydrodynamics

2.1.1 Discretising properties onto particles

Smoothed particle hydrodynamics (SPH) was introduced by Lucy (1977) and Gingold and Monaghan (1977). It is a Lagrangian particle method used to solve the equations of hydrodynamics. The fluid to be simulated is discretised onto a set of representative particles with each one assigned properties like mass, density, position, velocity and energy. The movement of the particles is dictated by the equations of motion and they carry their properties with them as they move.

For each particle in the fluid, the properties that are allowed to evolve (e.g. density, pressure and energy) depend on the neighbouring particles and in turn, their properties (by neighbouring particles, we mean the particles that are close to the particle of interest). To illustrate the fundamental SPH calculation, we consider how to calculate the density for the example particles shown in Figure 2.1 where a particle at a location of high density is indicated in red and one at low density location in green. As density is just mass divided by volume, the relative positions and masses of the particles can be used to determine the local density. For example, the more tightly packed particles in the red region represent a higher density than those in the green region. As the density at a particular particle is mostly influenced by its closest neighbours, the contribution of particles should be weighted so that the nearby neighbours have the greatest influence and the furthest have the least. In Figure 2.1, this weighting is indicated by gradient in the shaded region around the particles. These concepts are brought together in the SPH density summation, where the

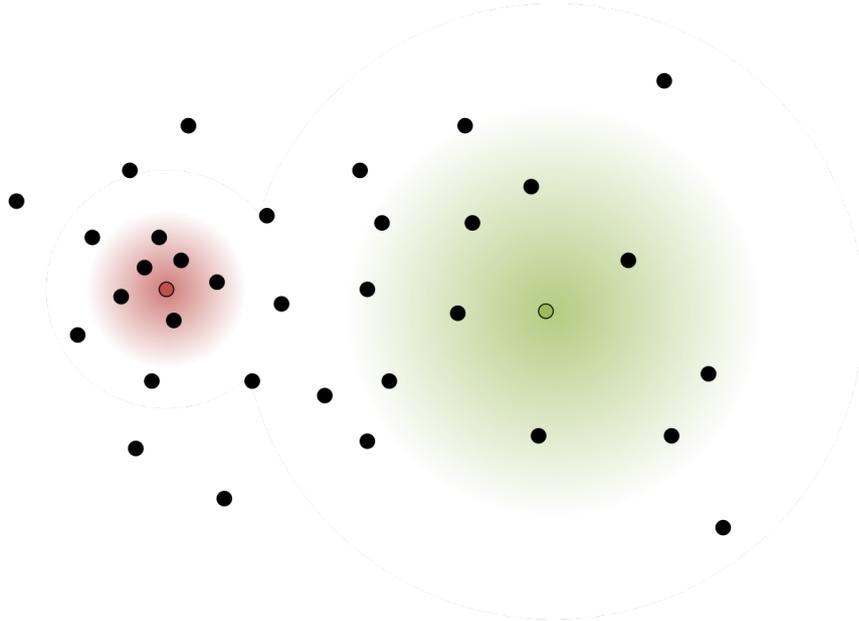


FIGURE 2.1: Particles representing a fluid. A high density region is indicated at the red particle and a low density region at the green particle, with the shaded regions representing the relative smoothing length. The gradient in the shaded regions indicates the weighting of the kernel function (based on Figure 1c of Price, 2012).

density ρ at the location of particle a is calculated by

$$\rho_a = \sum_b m_b W(\mathbf{r}_a - \mathbf{r}_b, h_a). \quad (2.1)$$

Here the subscript b refers to the neighbour particles that are included in the summation, each with mass m . The kernel W is a weighting function that specifies the contribution of each neighbouring particle based on the distance $\mathbf{r}_a - \mathbf{r}_b$ and is constrained by the smoothing length h , a characteristic length related to the particle spacing. The smoothing length determines the distance over which the contribution of the particles should be weighted. In the example Figure 2.1 this is indicated by the size of the shaded regions, with the less dense (green) region having a larger smoothing length and hence larger region than the more dense (red) region. The smoothing length is then dependent on the density, suggesting the following description for equal mass particles (i.e. $m_a = m_b = m$) in three-dimensions:

$$h_a = h_{\text{fact}} \left(\frac{m}{\rho_a} \right)^{1/3}. \quad (2.2)$$

Here h_{fact} is a factor in units of the mean local particle spacing. The final component to complete the density calculation is the kernel that decides the weighting of the neighbour particles. The kernel must satisfy a number of specific properties in order to be useful (from here we refer to $W(\mathbf{r}_a - \mathbf{r}_b, h_a)$ as $W_{ab}(h_a)$):

1. The kernel must be a positive function that decreases smoothly towards zero at larger distances.

2. The kernel must be symmetric ($W_{a,b} = W_{b,a}$) so that the influence of a neighbouring particle depends on its distance, not position.
3. At small distances between particles, the gradient of the kernel must be shallow so that if the nearest neighbours move slightly the density estimate does not change much.
4. The magnitude of the kernel function is normalised using

$$\int_V W_{ab}(h_a) dV' = 1. \quad (2.3)$$

Although a Gaussian profile satisfies all of these properties, it requires that all neighbours in the fluid contribute to the density for each particle — even though those that are far away do not contribute very much. For the green and red example particles in Figure 2.1, this would mean that all the particles in the domain are used in the density calculation, despite it being mostly determined by only the particles in the shaded regions. This leads to kernel functions that have a similar profile to a Gaussian but with ‘compact support’ that tend to zero within a few smoothing lengths. By comparison, these functions have an $\mathcal{O}(N_{neighbours}N)$ computational cost rather than an $\mathcal{O}(N^2)$ computational cost (computational cost is a measure of the time taken to complete a step in the computation and is strongly a function of the resolution in the simulation). Here $N_{neighbours}$ represents the neighbour particles that are included within the compact radius (within the shaded regions) and N is the total number of particles. The cubic spline kernel is one of the most commonly used (e.g. the M_4 B-spline, Monaghan and Lattanzio, 1985), where the kernel is set to

$$W_{ab}(h_a) = \sigma \frac{w(q)}{h_a^d}, \quad (2.4)$$

with the parameter q defined by

$$q \equiv \frac{|\mathbf{r}_a - \mathbf{r}_b|}{h_a}, \quad (2.5)$$

and d as the number of dimensions. The normalisation constant σ depends on d (e.g. $\sigma = 2/3, 10/(7\pi)$ or $1/\pi$ in 1, 2 or 3 dimensions). The function $w(q)$ is described by

$$w(q) = \begin{cases} \frac{1}{4}(2-q)^3 - (1-q)^3, & 0 \leq q \leq 1; \\ \frac{1}{4}(2-q)^3, & 1 \leq q \leq 2; \\ 0, & q \geq 2. \end{cases} \quad (2.6)$$

The M_4 kernel is used in our subsequent work but we note that kernels with a larger compact support radius (e.g. that go out to more than $2h$) have been shown to be more accurate (Dehnen and Aly, 2012) and other kernels have been explored (e.g. Wendland, 1995; Morris, 1996). With Equations 2.1 and 2.2 and an appropriate kernel we can thus calculate the density at the position of any particle in the fluid. This calculation is the

core of SPH, as the calculation of other properties like the pressure and velocities depend on the density.

2.1.2 General form and identities used in SPH

The SPH density summation introduced in Equation 2.1 is an example of using an interpolation function. If we consider a scalar property A , a more general interpolation function is given by

$$A(\mathbf{r}) = \int A(\mathbf{r}')\delta(\mathbf{r} - \mathbf{r}')d^3\mathbf{r}', \quad (2.7)$$

where \mathbf{r} is the position and δ is the Dirac delta function. The delta function may then be replaced with a kernel function that has a width determined by h , where $W \rightarrow \delta$ as $h \rightarrow 0$ such that

$$A(\mathbf{r}) = \int A(\mathbf{r}')W(\mathbf{r} - \mathbf{r}', h)d^3\mathbf{r}' + \mathcal{O}(h^2). \quad (2.8)$$

The last term here represents the error that has been introduced by using a finite kernel and is discussed below. To discretise this so that it may be applied to a set of particles, the integral is replaced by a summation and the mass element ρdV with the particle mass m . The scalar field on particle a may then be computed using

$$A(\mathbf{r})_a \approx \sum_b^N m_b \frac{A_b}{\rho_b} W_{ab}(h_a). \quad (2.9)$$

It is clear that if A is replaced by ρ , the above equation reduces to Equation 2.1. The importance of this general form is that it is differentiable and can be extended to vectors (\mathbf{A} rather than A), prompting the following useful identities (e.g. Price, 2012):

$$\nabla A(\mathbf{r}_a) \approx \sum_b^N m_b \frac{A_b}{\rho_b} \nabla W_{ab}(h_a), \quad (2.10)$$

$$\mathbf{A}(\mathbf{r}_a) \approx \sum_b^N m_b \frac{\mathbf{A}_b}{\rho_b} W_{ab}(h_a), \quad (2.11)$$

$$\nabla \cdot \mathbf{A}(\mathbf{r}_a) \approx \sum_b^N m_b \frac{\mathbf{A}_b}{\rho_b} \cdot \nabla W_{ab}(h_a). \quad (2.12)$$

These relations may be used to write the equations of motion in SPH form (further useful general vector identities are given in Equations 76 - 81 of Price, 2012). We also introduce some important relations for the kernel function (quoting in three-dimensions, e.g. Price, 2012):

$$\nabla_a W_{ab} = \hat{\mathbf{r}}_{ab} F_{ab}, \quad (2.13)$$

where

$$F_{ab} \equiv \frac{\sigma}{h_a^4} w'(q), \quad (2.14)$$

and the related term $G_{ab} \equiv -2F_{ab}/|r_{ab}|$. For completeness we also include the expression for the derivate of the kernel with respect to the smoothing length, used in subsequent calculations as

$$\frac{\partial W_{ab}(h_a)}{\partial h_a} = -\frac{\sigma}{h_a^4} [3w(q) + qw'(q)]. \quad (2.15)$$

2.1.3 Governing equations in SPH form

Along with the density, we must also be able to calculate the accelerations, pressures and energies of particles. Here we use the above identities to write the equations of motion and the energy equation in SPH summation form — although they may be derived directly from Equation 2.1 and the discretisation of the Lagrangian. The latter method guarantees conservation of total angular and linear momentum that we will confirm below. Additionally, in this work we do not make use of magnetohydrodynamics or cooling so these terms are not included. For now, we also consider the equations of motion without dissipation.

We begin by considering the density formulation in Equation 2.1. In constructing this, the total mass of the system was implicitly conserved because the mass on each particle is fixed. Because the continuity equation is another representation of the conservation of mass, we can show that the density summation is equivalent to a discretised version of the continuity equation. Starting with the continuity equation for a particle a (e.g. Monaghan, 2012; Price, 2012)

$$\begin{aligned} \frac{d\rho_a}{dt} &= -\rho_a(\nabla \cdot \mathbf{v})_a \\ &= \mathbf{v}_a \cdot \nabla \rho_a - \nabla \cdot (\rho_a \mathbf{v}_a). \end{aligned} \quad (2.16)$$

Using the general summation in Equation 2.9, this may be re-written as

$$\begin{aligned} \frac{d\rho_a}{dt} &= \mathbf{v}_a \cdot \sum_b m_b \nabla_a W_{ab} - \sum_b \frac{m_b}{\rho_b} (\rho_b \mathbf{v}_b) \cdot \nabla_a W_{ab}, \\ &= \sum_b m_b (\mathbf{v}_a - \mathbf{v}_b) \cdot \nabla_a W_{ab}. \end{aligned} \quad (2.17)$$

Recalling that $W_{ab}(h_a) = W(\mathbf{r}_a - \mathbf{r}_b, h_a)$, this is just the derivative of Equation 2.1 and confirms that the continuity equation is inherently contained in the original SPH density summation. The equations of motion for compressible hydrodynamics (ignoring dissipation terms for the moment) can be written as

$$\frac{d\mathbf{v}}{dt} = -\frac{\nabla P}{\rho} + \mathbf{a}_{\text{ext}}. \quad (2.18)$$

The \mathbf{a}_{ext} term represents the acceleration due to external forces per mass, although we are neglecting terms related to self gravity and sink particles in this work (compared to Equations 24 and 25 from Price et al., 2017). In the case of a compact central object with a gravitational potential Φ , the acceleration from this external body is

$$\mathbf{a}_{\text{ext}} = -\nabla\Phi. \quad (2.19)$$

In SPH summation form, the equations of motion are then expressed as

$$\frac{d\mathbf{v}_a}{dt} = -\sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla_a W_{ab}(h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla_a W_{ab}(h_a) \right] + \mathbf{a}_{\text{ext},a}. \quad (2.20)$$

For the adaptive smoothing length introduced above (where the smoothing length of each particle is a function of the density and changes during the course of a simulation), the convenient summation term Ω_a is defined (in three-dimensions) as

$$\Omega_a \equiv \left[1 + \frac{h_a}{3\rho_a} \sum_b m_b \frac{\partial W_{ab}(h_a)}{\partial h_a} \right]. \quad (2.21)$$

The full set of governing equations are completed with an equation to control changes in energy of each particle, u_a . Ignoring any cooling terms completely and viscous and shock heating for now, the energy is evolved through

$$\frac{du}{dt} = -\frac{P}{\rho} (\nabla \cdot \mathbf{v}). \quad (2.22)$$

The above equation is written in SPH form as

$$\frac{du_a}{dt} = \frac{P_a}{\Omega_a \rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab}(h_a), \quad (2.23)$$

where we have introduced $\mathbf{v}_{ab} = \mathbf{v}_a - \mathbf{v}_b$.

2.1.4 Conservation properties

As mentioned before, the governing equations in SPH form can be derived from just the density summation and the Lagrangian (e.g. see Price, 2012). This relationship dictates that the governing equations should exhibit translational and rotational symmetries: because the Lagrangian and density estimate are invariant to translations, the total linear momentum of the particles should be conserved. Similarly, they are invariant to rotations so the total angular momentum should also be conserved. Finally, because the Lagrangian is symmetric with respect to time and invariant to translations in time, the total energy of the system should also be conserved. We can show analytically that this is true by considering the time derivative of each of these in turn. Firstly, the total linear momentum

(Price, 2012),

$$\frac{d}{dt} \sum_a m_a \mathbf{v}_a = \sum_a m_a \frac{d\mathbf{v}_a}{dt} = - \sum_a \sum_b m_a m_b \left(\frac{P_a}{\rho_a^2} + \frac{P_b}{\rho_b^2} \right) \nabla_a W_{ab} = 0. \quad (2.24)$$

The final part of this identity makes use of the symmetry of the kernel gradient (Equation 2.13 implies that $\nabla_a W_{ab} = -\nabla_a W_{ba}$). Using a similar method and recalling that $(\mathbf{r}_a \times \mathbf{r}_b) = -(\mathbf{r}_b \times \mathbf{r}_a)$, the total angular momentum is also conserved:

$$\frac{d}{dt} \sum_a \mathbf{r}_a \times m_a \mathbf{v}_a = \sum_a m_a \left(\mathbf{r}_a \times \frac{d\mathbf{v}_a}{dt} \right) = - \sum_a \sum_b m_a m_b \left(\frac{P_a}{\rho_a^2} + \frac{P_b}{\rho_b^2} \right) \mathbf{r}_a \times \nabla_a W_{ab} = 0. \quad (2.25)$$

Finally, the (dissipationless) specific energy $e = \frac{1}{2}v^2 + u$, is used to confirm conservation of the total energy E as

$$\frac{dE}{dt} = \sum_a m_a \frac{de_a}{dt} = - \sum_a \sum_b m_a m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \mathbf{v}_b \cdot \nabla_a W_{ab}(h_a) + \frac{P_b}{\Omega_b \rho_b^2} \mathbf{v}_a \cdot \nabla_a W_{ab}(h_b) \right] = 0. \quad (2.26)$$

Although we can show algebraically that these properties are conserved, in practice they are only conserved to the accuracy of the time-stepping scheme that is employed. It is these conservation properties that makes SPH particularly useful for misaligned flows; conservation of angular and linear momentum is independent of the inclination of the gas.

2.1.5 Errors and derivatives

The general form introduced in Equation 2.7 is exact when the kernel function used is a δ function. However, when we chose to use a different kernel with a finite radius an error of order h^2 is introduced. This is confirmed using a Taylor series expansion (Benz, 1990; Monaghan, 1992), relying on the symmetry of the kernel to eliminate the error terms with odd powers. The error associated with the density approximation can thus be reduced by increasing the number of particles (the resolution in the simulation) which decreases the smoothing length, h .

In order to truly be second order accurate, Price (2012) state that the two following conditions should be satisfied:

$$\sum_b \frac{m_b}{\rho_b} W_{ab} \approx 1; \quad \text{and} \quad \sum_b \frac{m_b}{\rho_b} (\mathbf{r}_b - \mathbf{r}_a) W_{ab} \approx 0. \quad (2.27)$$

However, the amount to which these conditions are satisfied is determined by the distribution of particles within the kernel as well as the ratio of the smoothing length to particle spacing. In general, a regular distribution of particles that are modelled with the commonly used kernels (e.g. the M_4 kernel introduced in Equation 2.6) lead to these being well approximated. However, the error introduced by implementing a discrete summation

with a kernel also results in errors in the calculation of the gradient of $\mathcal{O}(h^2)$, and these are generally bound by the same criteria (quantified using a Taylor series expansion, see Price, 2012).

A full Taylor series expansion of the discrete summation version of the gradient function identifies the particular form of these errors (e.g. see section 4.4 of Price, 2012). It is thus possible to remove the $\mathcal{O}(h^2)$ error in the calculation of the gradient by subtracting the well defined error terms. While it is tempting to implement this in order to calculate a more accurate gradient, the error is only reduced in the linear regime and it has been shown that retaining the error instead leads to a significant advantage: the $\mathcal{O}(h^2)$ error is positive, meaning that particles feel a small repulsive force. Poorly arranged particles simulated with this small repulsive force from the error will try to re-arrange themselves to a more regular arrangement, as this represents a lower energy state. A better arrangement of particles means that the above criteria (Equation 2.27) are better satisfied, which leads to a better estimate in the gradient terms. Examples showing the effect of keeping versus removing the $\mathcal{O}(h^2)$ error term for calculating derivatives can be found in Figure 4 of Price (2012) and Monaghan (2005).

Second derivatives calculated using compact kernels are not particularly accurate. Instead, second derivatives (for a scalar A) are calculated with (Brookshaw, 1985);

$$\nabla^2 A_a \approx 2 \sum_b \frac{m_b}{\rho_b} (A_a - A_b) \frac{F_{ab}}{|r_{ab}|}. \quad (2.28)$$

Conceptually, this is the same as the first derivative of the kernel function divided by the particle spacing. The advantage of this form is that it automatically meets the equivalent criteria listed above and so provides a good density estimate. Equivalent terms for vector second derivatives can be found in Español and Revenga (2003) and Monaghan (2005).

2.1.6 Shock capturing and dissipation terms

The equations that have been introduced so far do not include any dissipation, but this is essential in order to correctly capture shocks. The jump that is present across a shock in gas properties (for example, density) are effectively instantaneous. This physical scenario cannot be well resolved in any code; in grid simulations the shock should be modelled across one cell and in particle simulations there would be so few particles that it is similarly poorly resolved. In order to model shocks numerically then, the shock is smoothed over some length-scale by the implementation of shock capturing dissipation terms. In SPH, shock capturing dissipation terms were first introduced by Monaghan (1997) but have been extended by Price and Monaghan (2004) and Price and Monaghan (2005) to MHD and by Laibe and Price (2014) to dust-gas mixtures. Two terms are introduced into the energy and momentum equations; Π_{shock} is a dissipation term that is used to give a correct entropy increase at a shock front and Λ_{shock} is a viscous shock heating term used in the energy

equation. These terms are defined by an introduced artificial viscosity which acts to damp the particle motion and smooth shocks.

The q_a^{AV} artificial viscosity term is in turn defined by two artificial viscosity parameters, α^{AV} and β^{AV} . The q_a^{AV} term is described as (Price et al., 2017)

$$q_a^{\text{AV}} = \begin{cases} -\frac{1}{2}\rho_a v_{\text{sig},a} \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}, & \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab} < 0, \\ 0, & \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab} \geq 0. \end{cases} \quad (2.29)$$

Here the signal velocity, $v_{\text{sig},a}$ is defined as

$$v_{\text{sig},a} = \alpha_a^{\text{AV}} c_{s,a} + \beta^{\text{AV}} |\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}|. \quad (2.30)$$

The artificial viscosity α^{AV} is distinct from the disc viscosity α , as the artificial viscosity coefficient is strictly introduced to accommodate shocks. This artificial viscosity term takes values between 0 and 1, but typically a minimum value of $\alpha_{\text{min}}^{\text{AV}} = 0.1$ is used to ensure that the shock is being modelled correctly. To ensure that these dissipation terms only act at the location of shocks and there is little dissipation in the absence of shocks, Morris and Monaghan (1997) introduced a switch that evolves α^{AV} according to

$$\frac{d\alpha_a^{\text{AV}}}{dt} = \max(-\nabla \cdot \mathbf{v}_a, 0) - \frac{\alpha_a^{\text{AV}} - \alpha_{\text{min}}^{\text{AV}}}{\tau_a}, \quad (2.31)$$

where $\tau = \sigma_{\text{decay}} h / v_{\text{sig}}$ and $\sigma_{\text{decay}} = 0.1$. The first term increases the viscosity at the location of shocks and the second term decreases it exponentially over the timescale τ . Other switches have been proposed (e.g. Cullen and Dehnen, 2010). The β^{AV} term is used to prevent approaching particles from penetrating each other and is independent of the choice of α^{AV} . As previous studies have shown (see Meru and Bate, 2011; Meru and Bate, 2012) and as we shall confirm in Chapter 3, using $\beta^{\text{AV}} \geq 2.0$ is particularly important for accurately recovering shocks when waves are propagating in an accretion disc.

In SPH form, the Navier-Stokes equations read (and including dissipation for shocks)

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla_a W_{ab}(h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla_a W_{ab}(h_b) \right] + \Pi_{\text{shock},a} + \mathbf{a}_{\text{ext},a}. \quad (2.32)$$

The dissipation from shocks is encapsulated in the Π_{shock} term which is in turn described using the q_a^{AV} artificial viscosity terms with

$$\Pi_{\text{shock}}^a \equiv - \sum_b m_b \left[\frac{q_a^{\text{AV}}}{\Omega_a \rho_a^2} \nabla_a W_{ab}(h_a) + \frac{q_b^{\text{AV}}}{\Omega_b \rho_b^2} \nabla_a W_{ab}(h_b) \right]. \quad (2.33)$$

A similar term is required in the energy equation in order to dissipate the jump in energy. The term Λ_{shock} is added to the energy equation such that

$$\frac{du_a}{dt} = \frac{P_a}{\Omega_a \rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab}(h_a) + \Lambda_{\text{shock},a}. \quad (2.34)$$

This term is expressed as

$$\begin{aligned} \Lambda_{\text{shock},a} = & -\frac{1}{\Omega_a} \sum_b m_b v_{\text{sig},a} \frac{1}{2} (\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab})^2 F_{ab}(h_a) \\ & + \sum_b m_b \alpha_u v_{\text{sig}}^u (u_a - u_b) \left[\frac{F_{ab}(h_a)}{\Omega_a \rho_a} + \frac{F_{ab}(h_b)}{\Omega_b \rho_b} \right]. \end{aligned} \quad (2.35)$$

where the first term is the viscous shock heating and the second is the thermal conductivity (note here we ignore terms relating to artificial resistivity). For the thermal conductivity, the value $\alpha_u = 1.0$ is used and the signal velocity is distinct from Equation 2.30 (see Price, 2012). For problems that do not involve self-gravity, we use the following signal velocity for the conductivity term

$$v_{\text{sig}}^u = \sqrt{\frac{|P_a - P_b|}{\bar{\rho}_{ab}}}. \quad (2.36)$$

The importance of these shock terms is demonstrated by the canonical Sod shock tube test shown in Figures 2.2 and 2.3. The initial conditions for the 1D Sod shock specify an ideal gas with zero velocity with a pressure ratio of 10 and a density ratio of 8 on either side of the shock. When artificial viscosity is not included the particles adjacent to the shock become noisy and there is a characteristic blip in the energy and pressure due to the contact discontinuity. Inclusion of the artificial viscosity smooths the particles around the shock and including conductivity removes the temperature jump problem, providing a better comparison to the exact solution (shown in red). The low resolution results in Figures 2.2 and 2.3 are obtained using a 1D code but are analogous to the 3D examples in Figures 8 and 9 in Price (2012) and in Price et al. (2017).

2.1.7 Physical viscosity

Contributions from bulk and shear viscosity are included with a viscous stress tensor. The compressible Navier-Stokes equations can be written in tensor notation for each particle as

$$\frac{dv_a^i}{dt} = -\frac{1}{\rho} \frac{\partial S_a^{ij}}{\partial x_a^j} + \Pi_{\text{shock},a} + \mathbf{a}_{\text{ext},a} \quad (2.37)$$

with the viscous stress tensor, S , given by

$$S_a^{ij} = \left[P_a - \left(\zeta_a - \frac{2}{3} \eta_a \right) \frac{\partial v_a^k}{\partial x_a^k} \right] \delta^{ij} - \eta_a \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right). \quad (2.38)$$

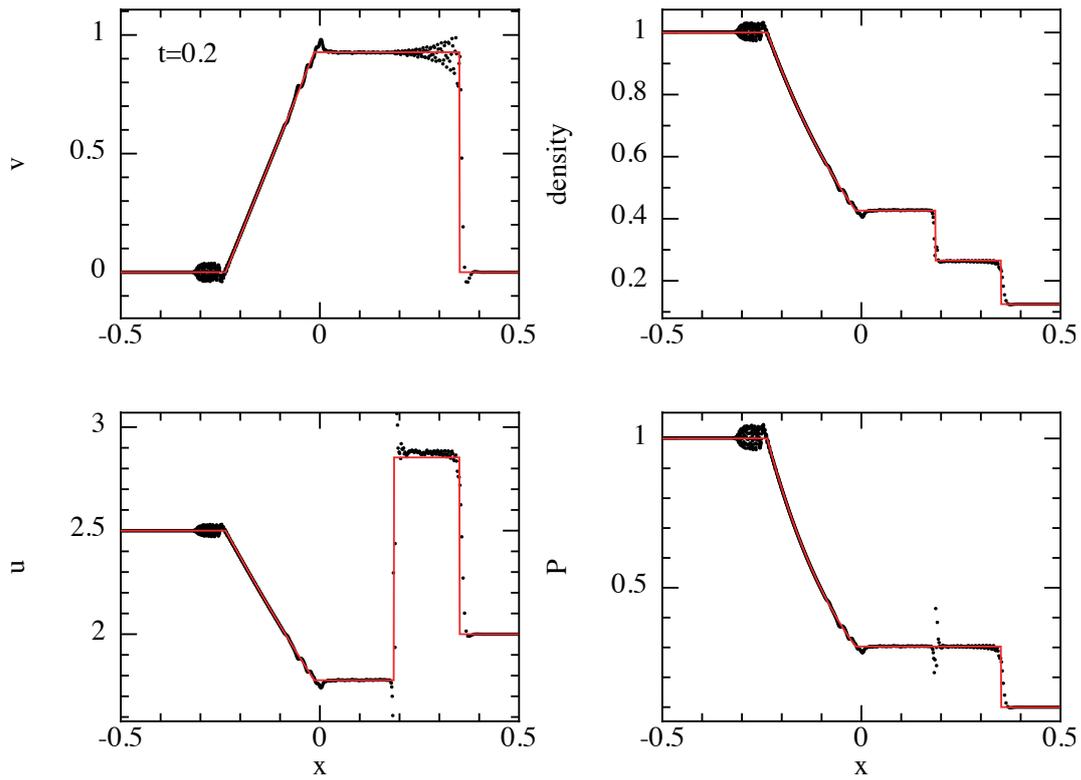


FIGURE 2.2: Sod shock tube example with negligible viscosity and conductivity terms. The red line shows the exact solution and the particles are in black.

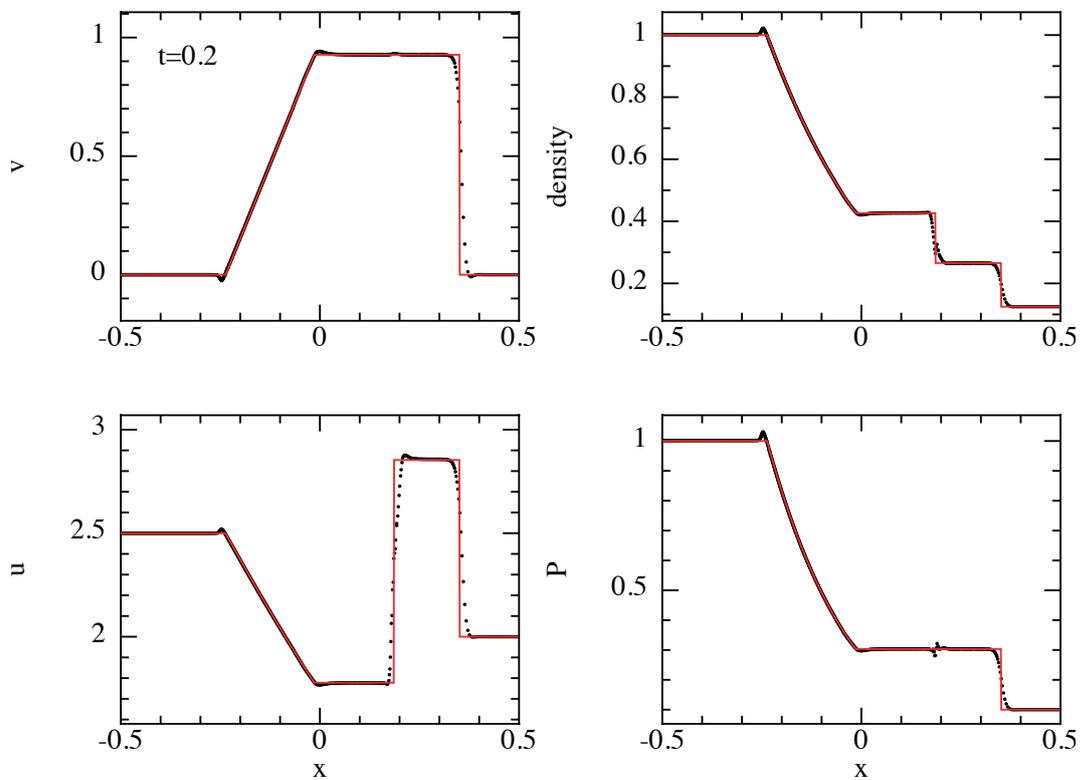


FIGURE 2.3: The same as Figure 2.2 but including the artificial viscosity terms and the switch described in Equation 2.31. The red line shows the exact solution and the particles are in black. Both of these tests were demonstrated with a 1D code but equivalent tests can be found in Price et al. (2017).

Here the bulk viscosity is ζ , shear viscosity η and the Kronecker-delta function δ . To calculate the influence of the viscous stress tensor the derivatives in S_a^{ij} must be evaluated and then the derivative of this is taken in the compressible Navier-Stokes equation (Equation 2.37) — this can be accomplished by either taking the first derivative twice or by taking a direct second derivative. The contribution to the energy equation from viscous heating, Λ_{visc} , alters the energy equation to read

$$\frac{du_a}{dt} = \frac{P_a}{\Omega_a \rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab}(h_a) + \Lambda_{\text{shock},a} + \Lambda_{\text{visc},a}. \quad (2.39)$$

The two different methods of taking the derivatives required in Equation 2.37 and 2.38 lead to two different expressions for the viscous heating term Λ_{visc} , described in depth below.

Physical viscosity with two first derivatives

In this method the first derivative is taken twice, introduced by Flebbe et al. (1994) and Watkins et al. (1996). The Navier-Stokes equation thus becomes

$$\frac{dv_a^i}{dt} = - \sum_b m_b \left[\frac{S_a^{ij}}{\Omega_a \rho_a^2} \nabla_a^j W_{ab}(h_a) + \frac{S_b^{ij}}{\Omega_b \rho_b^2} \nabla_a^j W_{ab}(h_a) \right] + \Pi_{\text{shock},a}^i + a_{\text{ext},a}^i. \quad (2.40)$$

In order to guarantee energy conservation and a positive definite entropy increase from dissipative terms, the equation below is used to calculate the velocity gradients:

$$\frac{\partial v_a^i}{\partial x_a^j} = \frac{1}{\Omega_a \rho_a} \sum_b m_b v_{ab}^i \nabla_a^j W_{ab}(h_a). \quad (2.41)$$

With this description, S^{ij} can be computed from Equation 2.38 and the viscous heating term can be written as (Price et al., 2017)

$$\Lambda_{\text{visc},a} = \left(\zeta_{v,a} - \frac{2}{3} \nu_a \right) (\nabla \cdot \mathbf{v})_a^2 + \frac{\nu_a}{2} \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right)^2, \quad (2.42)$$

where $\zeta_\nu = \zeta/\rho$ and $\nu = \eta/\rho$. Although this method is straightforward, it requires an extra six values to be stored in the stress tensor for each particle (taking into account symmetries).

Physical viscosity with direct second derivatives

Here the derivatives in the stress tensor and the Navier-Stokes equation are computed in one step with a second derivative. Following Español and Revenga (2003), this can be achieved by introducing two new variables τ and κ . These are defined using the bulk and

shear viscosity terms such that

$$\tau = \frac{5}{4} \left(\zeta + \frac{\eta}{3} \right) \quad \text{and} \quad \kappa = \frac{1}{12} (5\eta - 3\zeta). \quad (2.43)$$

The Navier-Stokes equation can then be re-written with these terms as

$$\begin{aligned} \left(\frac{dv_a^i}{dt} \right)_{\text{visc}} &= \sum_b m_b \left[\frac{\tau_a G_{ab}(h_a)}{\rho_a} + \frac{\tau_b G_{ab}(h_b)}{\rho_b} \right] (\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}) \hat{r}_{ab}^i \\ &+ \sum_b m_b \left[\frac{\kappa_a G_{ab}(h_a)}{\rho_a} + \frac{\kappa_b G_{ab}(h_b)}{\rho_b} \right] v_{ab}^i, \end{aligned} \quad (2.44)$$

and these are then used to define the viscous heating term directly as

$$\Lambda_{\text{visc},a} = \frac{\tau_a}{\rho_a} \sum_b m_b (\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab})^2 G_{ab}(h_a) + \frac{\kappa_a}{\rho_a} \sum_b m_b (\mathbf{v}_{ab})^2 G_{ab}(h_a). \quad (2.45)$$

The advantage of this method is that it does not require any additional storage (whereas the previous method does). Español and Revenga (2003) used this method and Price et al. (2017) reported that little difference has been found between the two methods. As a result we use the latter method in our calculations where physical viscosity is applied (e.g. physical viscosity is used in our simulations exclusively in Chapter 5).

2.1.8 Summary of governing equations

The main equations used to move the particles and that dictate the evolution of their properties are listed here, starting with the density summation (continuity equation) and the smoothing length;

$$\rho_a = \sum_b^N m_b W_{ab}(h_a) \quad \text{and} \quad h_a = h_{\text{fact}} \left(\frac{m_a}{\rho_a} \right)^{1/3}. \quad (2.46)$$

The equations of motion including shock and viscosity terms reads as

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla_a W_{ab}(h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla_a W_{ab}(h_b) \right] + \Pi_{\text{shock},a} + \mathbf{a}_{\text{ext},a}. \quad (2.47)$$

Here the shock capturing dissipative term $\Pi_{\text{shock},a}$ is defined in Equation 2.33. Including contributions from viscous and shock heating, the energy equation becomes

$$\frac{du_a}{dt} = \frac{P_a}{\Omega_a \rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab}(h_a) + \Lambda_{\text{shock}} + \Lambda_{\text{visc},a}, \quad (2.48)$$

with Λ_{shock} defined in Equation 2.35 and $\Lambda_{\text{visc},a}$ with either Equation 2.42 or 2.45.

2.1.9 Equation of state

This set of equations is completed with an appropriate equation of state. For an ideal gas, this is given by

$$P_a = (\gamma - 1)\rho_a u_a, \quad (2.49)$$

where $c_s = \sqrt{\gamma P/\rho}$ and γ is the ratio of specific heats. For a locally isothermal fluid with $P = c_s^2 \rho$ (most commonly used throughout this work) this reduces to

$$P_a = c_{s,a}^2 (R_a) \rho_a, \quad (2.50)$$

where R is the radius from the central object. By assuming a mean molecular weight of the gas we could calculate the temperature, however this is not considered here. We note that in locally isothermal case, the energy of each particle u_a does not evolve through Equation 2.48, but is instead calculated by Equation 2.49. The above equations are solved for a set of particles using the SPH code PHANTOM, described below.

2.2 Phantom

We use PHANTOM, an efficient low-memory smoothed particle magnetohydrodynamics (SPMHD) code. The current full details of this code can be found in Price et al. (2017). PHANTOM has been used for a number of astrophysical problems over most of the last decade, including (but not limited to) studies of turbulence (Price and Federrath, 2010; Liptai et al., 2017), simulations that include a mixture of dust and gas (Dipierro et al., 2015; Dipierro et al., 2016) and non-ideal MHD problems (Wurster, Price, and Ayliffe, 2014; Wurster, Price, and Bate, 2017). From these studies PHANTOM includes the ability to model self-gravity, dust-gas mixtures, MHD (ideal and non-ideal), H₂ chemistry and situations where mass is injected into the simulation.

In the case of accretion disc problems, PHANTOM has been used to investigate the evolution of warped discs (Lodato and Price, 2010), strongly inclined discs around a black hole (Nixon et al., 2012; Nixon, King, and Price, 2012), inclined discs around a binary (Nixon, King, and Price, 2013; Facchini, Lodato, and Price, 2013), binaries with inclined planets (Martin et al., 2014a; Martin et al., 2014b; Martin et al., 2016) and the formation of discs through tidal disruption events (Bonnerot et al., 2016; Coughlin et al., 2016; Coughlin et al., 2017).

We use a few standard features and parameters across most of our simulations. Unless otherwise stated, we make use of the M₄ cubic spline with $h_{\text{fact}} = 1.2$. We use individual time-stepping (where particles are evolved on the longest time-scale possible) to speed up the code, noting that this means that total energy is expected to be less well conserved (Price et al., 2017). In the chapters relating to accretion discs we use the Morris-Monaghan

α^{AV} viscosity switch but in later chapters we use the default Cullen and Denhen switch (the Morris and Monaghan, 1997, switch was the default at the beginning of this work and we wanted to be consistent across accretion disc simulations). The simulations in Chapters 3 and 4 use a locally isothermal equation of state, Chapter 5 globally isothermal and Chapter 6 an adiabatic equation of state. When physical viscosity is explicitly modelled it is calculated with direct second derivatives. Modelling of the α parameter is accomplished from the numerical viscosity (and thus depends on α^{AV} , discussed below). Time integration in PHANTOM uses the velocity Verlet method (a leapfrog integrator, Verlet, 1967; Price et al., 2017).

The two features of PHANTOM that are most important to the evolution of accretion discs around rotating black holes are how the Lense-Thirring precession and disc viscosity are modelled. Tests for these can be found Nealon, Price, and Nixon (2015) and Price et al. (2017) but are included here for completeness. Here we describe the general setup and analysis that is common to the following chapters and how these two features are represented in PHANTOM.

2.2.1 Arranging the particles

The initial conditions require the positions and velocities of the particles to be specified. These are then used to calculate properties like the density on each particle before the simulation is started. For an accretion disc, the particles are set with a purely azimuthal velocity and a specified surface density profile. In the case of the torus, the particles are initialised on a grid with zero velocity (and no external potential) and relaxed to allow the particles to settle to a lower energy state.

Constructing a disc

For an accretion disc the initial particle positions are generated using a Monte Carlo placement method. This is constrained by the surface density profile for an accretion disc aligned with the black hole spin (Lynden-Bell and Pringle, 1974),

$$\Sigma(R) = \Sigma_{\text{in}} \left(\frac{R}{R_{\text{in}}} \right)^{-p} \left(1 - \sqrt{\frac{R_{\text{in}}}{R}} \right), \quad (2.51)$$

where the scaling parameter Σ_{in} is the surface density at the inner edge and the index p constrains the steepness of the surface density profile. In turn, each particle is given a random radius and angle in the x - y plane. From this position the sound speed is calculated. In the case that the disc is isothermal,

$$c_s(R) = c_{s,\text{in}} \left(\frac{R}{R_{\text{in}}} \right)^{-q}. \quad (2.52)$$

where the index q relates to the disc thickness profile and the sound speed at the inner edge $c_{s,\text{in}}$ is related to Σ_{in} . Assuming the disc is Keplerian, the scale height at each radius ($H = c_s/\Omega_k$) can then be estimated. The z position of each particle is then chosen randomly but restricted to be $\pm 3\sqrt{2}H$ from the $z = 0$ plane such that the density perpendicular to the disc has a truncated Gaussian profile.

Once the positions are set, the value of Σ_{in} is found by assuming the disc to be non self-gravitating, corresponding to a mass ratio between the disc and black hole of $\sim 10^{-4}$. This assumed Σ_{in} value is used to construct and calculate the mass of the disc by discretising it into concentric rings. The total mass of the disc divided by the number of particles then determines the mass of each particle.

The velocities of the particles are set to give centrifugal balance, assuming that the vertical and radial velocities are zero. In the case of a non-rotating black hole, this reduces to the almost Keplerian expression

$$v_\phi = \sqrt{\frac{GM}{r} + \frac{1}{\rho}\nabla P}. \quad (2.53)$$

The equivalent expression for a rotating black hole is introduced in Chapter 3. Once the particles have been set up with the correct surface density profile, sound speed and radial velocity they are rotated through the inclination β . The inclined positions (x', y', z') and velocities (v'_x, v'_y, v'_z) are related to the aligned (x, y, z) and (v_x, v_y, v_z) through

$$\begin{aligned} x' &= x \cos \beta + z \sin \beta, & v'_x &= v_x \cos \beta + v_z \sin \beta, \\ y' &= y, & v'_y &= v_y, \\ z' &= -x \sin \beta + z \cos \beta, & v'_z &= -v_x \sin \beta + v_z \cos \beta. \end{aligned} \quad (2.54)$$

After this rotation the velocity of the particles must be corrected to take into account the rotation of the black hole, discussed in depth in Chapter 3. The setup routine also calculates the α^{AV} required for the calculated $\langle h \rangle/H$ value to achieve the specified α_{SS} in the disc.

Constructing a torus

The properties of the particles in a torus are constrained by the radius of the maximum density, R_0 , the shape factor d , and the maximum density. The sound speed of particles in the torus is constrained by the maximum density, ρ_{max} , using (Papaloizou and Pringle, 1984)

$$\rho_{\text{max}} = \left[\frac{GM}{(n+1)AR_0} \left(\frac{d-1}{2d} \right) \right]^n, \quad (2.55)$$

where $n = (\gamma - 1)^{-1}$ is the polytropic index, R_0 is the cusp radius (radius of maximum density), A is the polytropic constant and d is a shape factor that determines the cross section of the torus. Before constructing the torus, the mass of each particle is estimated

by summing the mass in spherical shells using the density prescription (Papaloizou and Pringle, 1984)

$$\frac{P}{\rho} = \frac{GM}{(n+1)R_0} \left[\frac{R_0}{r} - \frac{1}{2} \left(\frac{R_0}{r \sin \theta} \right)^2 - \frac{1}{2d} \right]. \quad (2.56)$$

Here θ is the angle above the $z = 0$ plane. The total mass found is then divided by the number of particles to give the particle mass.

Rather than being arranged randomly, here the particles are set up on concentric shells. Each shell has a grid in z and ϕ superimposed on it, with a particle located at the intersections of these grids. The separation of the shells in radius, δr , as well as the grid density in δz and $\delta \phi$ are all tuneable parameters. The number of shells is determined by constructing the first shell at R_0 , then constructing each shell at $r + \delta r$ outwards with decreasing density, stopping when the density is less than zero. This process is repeated starting at R_0 and moving inwards to build the inner half of the torus.

Velocities are added to the particles after the torus is built. The velocities are set to balance pressure gradients using

$$v_\phi^2 = \frac{GM R_0}{r^2}, \quad (2.57)$$

where $r^2 = x^2 + y^2$ is the cylindrical radius (tori do not have any z velocity dependence). In general, because the torus is set up on a regular grid it is necessary to evolve it in a ‘relaxing potential’ first — particles naturally do not like being arranged on a grid, and this process allows them to relax to their preferred, lower energy arrangement. In this case, the initial velocities are all set to zero and the torus is evolved using the effective potential

$$\Phi_{\text{relaxing}} = -\frac{1}{R^3} + \frac{R_0}{2r^2}, \quad (2.58)$$

where the centrifugal force balances both the pressure gradient and the gravity of the central point mass. This evolution continues until the potential of energy of the torus ceases fluctuating, which normally takes less than five orbits at R_0 . After relaxing, velocities are added according to Equation 2.57 and the potential updated to that of a point mass.

2.2.2 Comparing 3D simulations to 1D solutions

To measure the evolution of the surface density, tilt and twist profiles of an accretion disc we follow an averaging process described in Lodato and Price (2010). The disc is discretised into N concentric spherical shells, with constant radial thickness given by $\Delta_R = (R_{\text{out}} - R_{\text{in}})/N$. The averaged value of each property in the i^{th} shell is then found by summing over the k particles in that shell and then dividing by k . For example, the surface density is described by (Lodato and Price, 2010)

$$\Sigma_i = \frac{\sum_k m_k}{\pi [(R_i + \Delta/2)^2 - (R_i - \Delta/2)^2]}, \quad (2.59)$$

where k represents the number of particles in the i^{th} shell. Similarly, the average angular momentum components (and thus the tilt and twist of the disc) is calculated by

$$\mathbf{l}_i = \frac{\sum_k m_k \mathbf{r}_k \times \mathbf{v}_k}{|\sum_k m_k \mathbf{r}_k \times \mathbf{v}_k|}. \quad (2.60)$$

Averaging the disc properties azimuthally like this reduces them to a radial dependence, allowing them to be compared to 1D codes (such as that used in Chapter 3). Direct comparison is achieved by ensuring that the 1D code and simulations are output at the same time in the disc evolution.

This method is then extended to measure properties as a function of azimuthal angle. Here the disc is split into M azimuthal wedges, with an angular width $\Delta_\phi = 2\pi/M$. The averaged properties inside the j^{th} wedge is again determined by averaging over the particles contained within it. For example, the radially averaged density is given by

$$\rho_j = \frac{\sum_k \rho_k}{k}, \quad (2.61)$$

where k is the number of particles in the wedge. Radially averaging the particles allows properties such as the density perturbation from the mean density to be plotted as a function of the azimuthal angle.

2.2.3 Disc viscosity

We recall that the disc viscosity (parameterised by α) controls the evolution of the accretion disc completely and thus modelling it correctly is critical. As described by Lodato and Price (2010), there two ways to accomplish this. In either case, we seek to have a parameter α_{SS} that can be set at run-time so that the viscosity in the disc can be specified. We note that despite following convention, the notation here becomes confusing: α is the disc viscosity, α_{AV} is the artificial viscosity included in the numerical method to accommodate shocks and α_{SS} is the viscosity parameter set at run time in simulations. In this section we are describing two methods to ensure that α and α_{SS} are the same. Testing the validity of these two approaches to modelling the disc viscosity with PHANTOM is the topic of Lodato and Price (2010), so we do not conduct a test here — in particular, see their Figure 4.

Using the Navier-Stokes viscosity to represent disc viscosity

The more direct method is to calculate the viscous terms directly from the Navier-Stokes equations. In this case, the parameter α_{SS} is specified and defines the kinematic viscosity for each particle as

$$\nu_a = \alpha_{\text{SS}} c_{s,a} H_a, \quad (2.62)$$

where H_a is the scale-height of the disc. The scale-height is calculated by assuming Keplerian rotation such that

$$H_a = \frac{c_{s,a}}{\Omega(R_a)}. \quad (2.63)$$

The advantage of this method is that the shear viscosity is directly defined and is independent of the smoothing length. However, because of the approximation used, this method must be restricted to situations where the orbits are near Keplerian and the radius R_a is physically meaningful (e.g. this method cannot be used around binary stars, Price et al., 2017). Additionally, Lodato and Price (2010) showed that some bulk viscosity needs to be included in order to capture shocks correctly and to prevent particle interpenetration.

Using numerical viscosity to represent disc viscosity

The second method used to represent disc viscosity relates the numerical viscosity imposed in the Navier-Stokes equations to the bulk and shear viscosity of the fluid (Artymowicz and Lubow, 1994; Murray, 1996). Espaol and Revenga (2003) showed that

$$\sum_b \equiv \frac{m_b}{\bar{\rho}_{ab}} (\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}) \frac{\nabla_a W_{ab}}{|r_{ab}|} \approx \frac{1}{5} \nabla(\nabla \cdot \mathbf{v}) + \frac{1}{10} \nabla^2 \mathbf{v}. \quad (2.64)$$

This form allows Equation 2.33 to be re-written in terms of Navier-Stokes terms. In other words, the Π_{shock} terms are the same as the Navier-Stokes viscosity with a constant ratio between the bulk and shear viscosity (Murray, 1996; Espaol and Revenga, 2003; Monaghan, 2005; Jubelgas, Springel, and Dolag, 2004; Lodato and Price, 2010; Meru and Bate, 2012). In this form, q_a^{AV} is redefined as

$$q_a^{\text{AV}} = \begin{cases} -\frac{1}{2} \rho_a \frac{h_a}{|r_{ab}|} (\alpha^{\text{AV}} c_{s,a} + \beta_{\text{AV}}) \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}, & \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab} < 0, \\ -\frac{1}{2} \rho_a \frac{h_a}{|r_{ab}|} (\alpha^{\text{AV}} c_{s,a} \mathbf{v}_{ab}) \cdot \hat{\mathbf{r}}_{ab} & \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab} \geq 0. \end{cases} \quad (2.65)$$

Making use of the above Espaol and Revenga (2003) identity, this means the bulk and shear viscosity can be written as (e.g. Lodato and Price, 2010)

$$\nu^{\text{AV}} \approx \frac{1}{10} \alpha^{\text{AV}} v_{\text{sig}} h, \quad (2.66)$$

$$\zeta_{\nu}^{\text{AV}} \approx \frac{1}{6} \alpha^{\text{AV}} v_{\text{sig}} h. \quad (2.67)$$

The above form may be used to represent the Shakura and Sunyaev (1973) viscosity by using the artificial viscosity (that was introduced to accommodate shocks) α^{AV} directly. However, to do so the following needs to be assumed (Lodato and Price, 2010):

1. Viscosity should be applied to both approaching and receding particles (rather than just approaching).

2. The signal velocity is altered to $v_{\text{sig}} = c_s$.
3. Constant α^{AV} is set and shock detection switches are turned off (so that the artificial viscosity is constant for all particles).
4. The viscosity term is scaled by the factor $h/|r_{ab}|$.

The resultant prescription for the α_{SS} viscosity is thus

$$\alpha_{\text{SS}} \approx \frac{1}{10} \alpha^{\text{AV}} \frac{\langle h \rangle}{H}, \quad (2.68)$$

where $\langle h \rangle$ is the shell averaged smoothing length (i.e. the average smoothing length as a function of radius). For accretion discs, it is common to desire that α be constant throughout the disc (Lodato and Pringle, 2007). Since α^{AV} is a constant, this can be achieved by forcing the radial dependence of $\langle h \rangle$ and H to be the same. The radial dependence of H is determined from its relation with the sound speed and angular velocity profile such that

$$H = \frac{c_s}{\Omega} \propto R^{3/2-q}, \quad (2.69)$$

where q determines the profile of the sound speed. The shell averaged smoothing length can be related to the density such that its radial dependence is

$$\langle h \rangle \propto \rho^{-1/3} \propto \left(\frac{\Sigma}{H} \right)^{-1/3} \propto R^{(p-q)/3+1/2}. \quad (2.70)$$

Equating these dependencies, we thus require that $3/2 - q = (p - q)/3 + 1/2$ for α to be constant with radius. We adopt the convenient choice of $p = 1.5$ and $q = 0.75$ (i.e. $\langle h \rangle \propto H \propto R^{3/4}$), unless otherwise stated.

2.2.4 Simulating wave-like accretion discs

To confirm that we can correctly describe the propagation of warps in the wavelike regime, we use the test described by Fragner and Nelson (2010). They simulated a wavelike accretion disc with a point mass potential and compared to a 1D calculation, finding agreement at the $\sim 10\%$ level. We choose to compare to this 1D solution instead of the solution from Equations 1.27 and 1.28 because the linear solution from Fragner and Nelson (2010) allows the surface density to evolve, as occurs in our 3D simulations.

We conduct a simulation using the same parameters cited in figure 1 of Fragner and Nelson (2010). Here $H/R = 0.03$, $\alpha = 0.001$ and an initial disturbance of 5° is used. In this case we do not drive the evolution with Lense-Thirring precession, so that we can isolate the behaviour due to warp propagation only. Our results are shown in Figure 2.4 using 10^6 and 10^7 particles. As the disc evolves, the disturbance splits into two waves travelling at half the sound speed (as predicted by Papaloizou and Lin 1995); one inward and the other

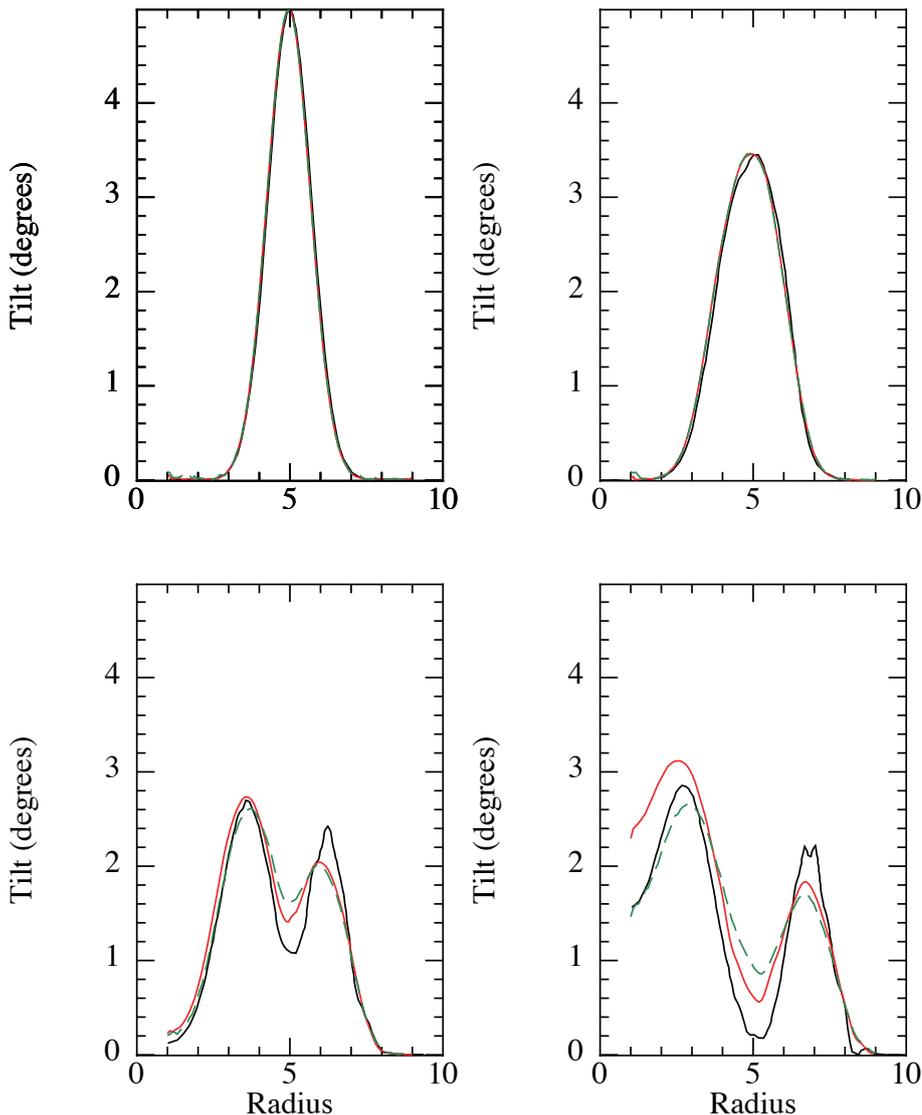


FIGURE 2.4: Evolution of bending waves in a disc, not subject to Lense-Thirring precession. The green (10^6 particles) and red (10^7 particles) lines show the results from our 3D simulation. The black line shows the results from the 1D code of Fragner and Nelson (Figure 1, 2010), using the same initial parameters. The agreement between these two solutions confirms that SPH can be used to describe the evolution of warp propagation in the wavelike regime.

outward. By the end of the simulation these have fully separated and are beginning to interact with the boundaries.

The SPH solution shows the same behaviour as the 1D solution, and increasing the resolution reduces the discrepancy. However, at late times near the inner edge of the disc there is increasing disagreement, most likely due to differences in the inner boundary condition. For example, in the 1D code assumptions must be made about the surface density profile at the inner boundary that are generated self-consistently in the 3D simulation. Although we are not able to resolve these differences, we note that the lower resolution test conducted in 3D by Fragner and Nelson (2010) using 20,000 particles experiences a similar discrepancy at the inner edge and that the propagation of the wave (represented by the radial location

of the two peaks) is consistent with the 1D solution. This test confirms that PHANTOM can be used to describe the propagation of warps in the wavelike regime.

2.2.5 Modelling Lense-Thirring precession

In PHANTOM, the precession effects due to the rotation of the black hole are modelled using a post-Newtonian approximation. In other words, we use the Newtonian potential due to a point mass and subsequently add a first order correction (in v/c) to mimic the effect of the black hole rotation. In our notation, this is represented by the external forces expressed as

$$\mathbf{a}_{\text{ext}} = -\nabla\Phi + \mathbf{v} \times \mathbf{h}, \quad (2.71)$$

where $\Phi = -GM/r$. The second term is the gravo-magnetic acceleration and is proportional to the spin of the black hole. This term is estimated using a dipole approximation with

$$\mathbf{h} \equiv \frac{2\mathbf{S}}{R^3} - \frac{6(\mathbf{S} \cdot \mathbf{r})\mathbf{r}}{R^5}, \quad (2.72)$$

where $\mathbf{S} = a(GM)^2\mathbf{k}/c^3$ and \mathbf{k} is a unit vector parallel to the black hole spin. A complication to the implementation in the code is that we use a leapfrog integrator in the ‘Velocity-Verlet’ form, where the positions and velocities of the particles are updated from time t^n to t^{n+1} according to

$$\mathbf{v}^{n+\frac{1}{2}} = \mathbf{v}^n + \frac{1}{2}\Delta t\mathbf{a}^n, \quad (2.73)$$

$$\mathbf{x}^{n+1} = \mathbf{x}^n + \Delta t\mathbf{v}^{n+\frac{1}{2}}, \quad (2.74)$$

$$\mathbf{v}^{n+1} = \mathbf{v}^{n+\frac{1}{2}} + \frac{1}{2}\Delta t\mathbf{a}^{n+1}. \quad (2.75)$$

However, the acceleration caused by Lense-Thirring precession depends on velocity. Thus (2.75) becomes implicit. This can be easily solved by writing the corrector step in the form

$$\mathbf{v}^{n+1} = \mathbf{v}^{n+\frac{1}{2}} + \frac{1}{2}\Delta t\mathbf{a}_{\text{pos}}^{n+1} + \frac{1}{2}\Delta t(\mathbf{v}^{n+1} \times \mathbf{h}^{n+1}), \quad (2.76)$$

where $\mathbf{a}_{\text{pos}}^{n+1}$ contains the position-dependent terms. This forms a set of three linear equations for each component of \mathbf{v}^{n+1} , that we solve analytically by inverting the resulting 3 x 3 matrix.

We also perform a simple test of the Lense-Thirring precession. We simulate a disc consisting of test particles with no viscosity and zero sound speed (i.e. $\alpha = c_s = 0$) subject to Lense-Thirring precession. The initial velocities are set including a correction due to the Lense-Thirring precession, discussed in depth in Chapter 3, and the disc is inclined at 30° . We then calculate the precession in the disc as a function of the radius using the procedure outlined in Appendix A.

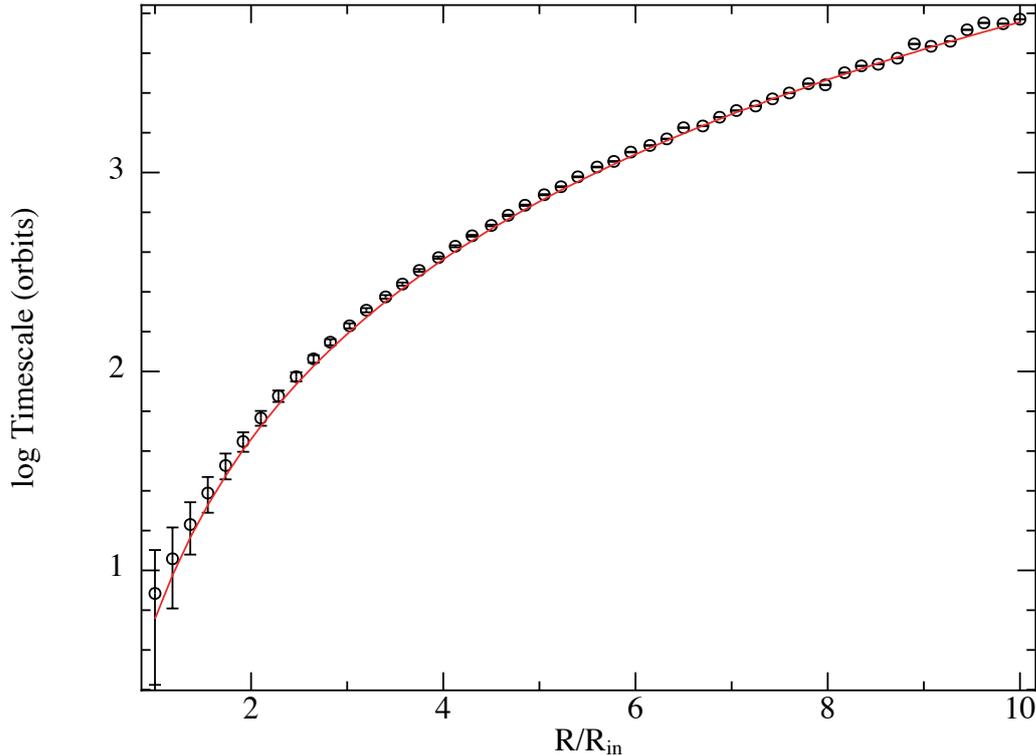


FIGURE 2.5: Precession timescale measured from an inviscid and pressureless 3D disc as a function of radius (black circles), compared to the expected Lense-Thirring precession (red line). $R_{\text{in}} = 4R_{\text{g}}$.

Figure 2.5 shows the comparison between the precession measured from our disc and the predicted precession in the disc, given by $t_{\text{precession}} = R^3/(2a)$. We find agreement within measurement uncertainties throughout the disc.

2.3 Limitations

The accuracy of all numerical techniques is dictated by the resolution and time-stepping scheme used. As we showed in this chapter, calculation of variables like the density is accurate to $\mathcal{O}(h^2)$, so we can repeat the simulation at multiple resolutions to quantify the effect of this error. While the time-stepping scheme does introduce an error in the conserved variables, PHANTOM conserves linear and angular momentum to machine precision when all particles are evolved on the same time-step (and a relative error of $\sim 10^{-3}$ in total energy when individual time-steps are used Price et al., 2017).

Particle methods offer a clear advantage for investigating highly inclined discs, as they do not suffer from orientation issues found in mesh-based approaches. However, magnetohydrodynamic simulations of tearing regions using particle methods is currently untenable, as there is low resolution at the location of the disc break (by definition) but accurate simulations including magnetic effects require extremely high resolution. These restrictions present a major problem: particle methods are able to simulate at inclinations that result in tearing but cannot include magnetic effects, while grid based methods are able

to include magnetic effects but cannot accurately simulate at such high inclinations. This limits us to using purely hydrodynamic simulations currently, and although this is justified in low inclination discs (as we will show in Chapter 4) this has not yet been confirmed in tearing discs.

PHANTOM is specifically limited by its ability to represent relativistic effects, as it currently implements a post-Newtonian approximation. Whilst this approximation is valid at moderate distances from the black hole, the rate of precession $\lesssim 10R_g$ can be incorrect by up to 50%. A full general relativistic implementation is outside the scope of this work.

The physical extent of the discs we are able to simulate is also problematic, as a larger disc is more computationally expensive for the same resolution. In Chapter 5 we consider applying our findings to low mass X-ray binaries, which have discs that are of the order of the separation of the binary (1000s of R_g). Although this is far too large to simulate, in Chapter 3 we discuss that on the condition that the outer radius is large enough to demonstrate tearing, the disc behaviour we identify is independent of the location of the outer radius. On this basis we expect our results to be consistent with discs that have a larger radial extent than we can currently simulate.

The final major limitations faced in this work were computational. As discussed in later chapters, the accretion disc must have an outer radius well outside the tearing radius and be well resolved ($\langle h \rangle / H \lesssim 1$) in order to demonstrate tearing — but these are competing requirements as increasing the outer radius for a given number of particles decreases the resolution in the disc. Compounding this problem, hundreds of orbits at the inner edge are required to observe the evolution of the disc, making the simulations computationally expensive. In Chapter 3 achieving a compromise between these requirements was manageable, but in Chapter 5 this caused major difficulties.

Declaration for Chapter 3

Declaration by Candidate

The work presented in § 3.1 to § 3.5 and § 2.2.4, 2.2.5 has been published in Nealon, Price, and Nixon (2015).

I declare that my contribution of the work presented in this chapter involved the following:

The writing of the publication, running the 1D code, testing wave propagation in the wave-like regime, testing the Lense-Thirring precession rate, and conducting, analysing and presenting the 3D simulations (more than 50 were conducted).

Extent of contribution by candidate: 80%

Co-author names: Chris Nixon and Daniel Price

Nature of co-author's contribution: Providing codes used for simulations, guidance in comparing the 1D and 3D results, guidance in choosing the simulation parameters, comments and feedback during the writing of the paper.

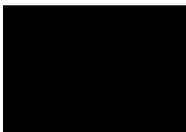
I have renumbered sections, equations and figures of the published paper in order to generate a consistent presentation within the thesis.

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the student's and co-authors' contributions to this work.

Student signature: _____

Date: November 20, 2017

Main Supervisor signature:



Date: November 20, 2017

Chapter 3

On the Bardeen-Petterson Effect in black hole accretion discs

Does the Bardeen-Petterson Effect occur in numerical simulations of inclined discs? There is currently no consensus. Here we re-examine the evolution of inclined discs and hence the behaviour of the Bardeen-Petterson Effect using the SPH code PHANTOM. Complementary to previous numerical studies, we focus on the disc regime where communication of disturbances in the disc is governed by pressure waves. As previous analytical studies in this regime have been restricted to the linear inclinations, this chapter begins by considering small inclinations before moving on to larger, non-linear inclinations.

Two mechanisms that may prevent the Bardeen-Petterson Effect from occurring are considered. First, the possible misalignment of the inner disc due to the development of a steady state oscillatory tilt near the black hole as predicted analytically by Lubow, Ogilvie, and Pringle (2002). Second, we investigate the ability of the disc to accrete material misaligned from large radii, comparing to previous simulations by Fragile et al. (2007).

The possibility of disc tearing is investigated in the non-linear regime. Focussing on a similar study by Nelson and Papaloizou (2000), we quantify the resolution criteria for disc breaking. Complementary to a study by Nixon et al. (2012) in the diffusive regime, we consider the mass accretion rate of broken and tearing discs. Finally, we investigate the location of the tear and the width of the rings torn.

A caveat is that we do not consider magnetic fields, despite it being widely accepted that magnetorotational instability (MRI) is the controlling mechanism for viscosity in the disc (Balbus and Hawley, 1991). However, magnetic fields have been shown to have little effect on the geometrical evolution (Sorathia, Krolik, and Hawley, 2013) meaning this can be accurately approximated by a viscous evolution (or by the ‘ α model’). Following Nelson and Papaloizou (2000), we adopt a post-Newtonian approach instead of general relativity (GR). One of the findings of this chapter is that this approximation must be considered carefully in order to capture the combination of relativistic effects that lead to tilt oscillations.

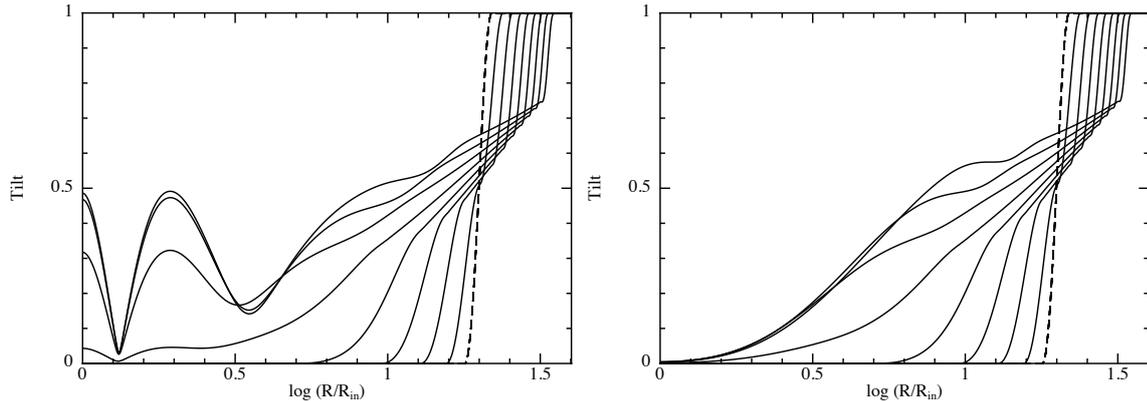


FIGURE 3.1: The steady state tilt profile found by LOP02, as in their Figure 3 and shown to the same time (at eight equally spaced times) but with a logarithmic scale for clarity in the inner region. The original apsidal and nodal frequencies are used in the left panel (Equations 3.5 and 3.6) with the same sign. In the right panel we have reversed the sign of the nodal frequency so the precession frequencies have different signs, and there are no oscillations present in this steady state (similar to LOP02, Figure 6). Here the tilt is shown as a fraction of the maximum tilt, the initial condition is shown with the dashed line and $R_{\text{in}} = 4R_g$.

3.1 Does the Bardeen-Petterson effect hold in the wavelike regime?

We consider three possible ways that the Bardeen-Petterson effect may be violated in wavelike discs. Firstly, radial oscillations in the tilt of the disc may prevent the disc from aligning at the inner edge. Secondly, the smooth transition between aligned and misaligned material may be broken if the disc tears, as has been observed in the diffusive regime (Nixon and King, 2012; Nixon et al., 2012). Finally, it may not be possible for the disc to find a steady state if the disc is relatively thick and the viscous time is short.

3.1.1 Is the inner disc aligned?

Lubow, Ogilvie, and Pringle (2002) (LOP02) considered warps in geometrically thin, almost Keplerian discs described by a surface density $\Sigma(R)$ and angular velocity $\Omega(R)$. The scale height of the disc is given by $H(R) \equiv c_s/\Omega$, where $c_s(R)$ is the sound speed in the disc. Their description is one dimensional in the sense that the total angular momentum in the disc \mathbf{L} is a function only of the cylindrical radial coordinate, R . The disc is then discretised into a series of rings, each described by the orientation of its tilt and twist angle. The tilt angle β is measured from the z -axis, and if this angle varies with radius the disc is considered to be warped. The twist angle γ is measured from an axis that is perpendicular to the z -axis, and similarly, if the twist angle varies with radius the disc is twisted. These two angles can be related to the unit angular momentum vector by $\mathbf{l} = \mathbf{L}/L = (\cos \gamma \sin \beta, \sin \gamma \sin \beta, \cos \beta)$ (Pringle, 1996).

We assume an α disc viscosity where $\nu = \alpha c_s H$ (Shakura and Sunyaev, 1973). For accretion discs with $\alpha \lesssim H/R$, the warp propagates as a dispersive wave (Papaloizou and Pringle,

1983; Papaloizou and Lin, 1995). Assuming that the disc is nearly Keplerian and not self-gravitating, the equations of motion describing the wave propagation are (Lubow and Ogilvie, 2000; Lubow, Ogilvie, and Pringle, 2002)

$$\Sigma R^2 \Omega \frac{\partial \mathbf{l}}{\partial t} = \frac{1}{R} \frac{\partial \mathbf{G}}{\partial R} + \mathbf{T}, \quad (3.1)$$

$$\frac{\partial \mathbf{G}}{\partial t} - \left(\frac{\Omega^2 - \kappa^2}{2\Omega} \right) \mathbf{l} \times \mathbf{G} + \alpha \Omega \mathbf{G} = \Sigma R^3 \Omega \frac{c_s^2}{4} \frac{\partial \mathbf{l}}{\partial R}. \quad (3.2)$$

Here κ is the radial epicyclic frequency, \mathbf{G} represents the internal horizontal torque in the disc and \mathbf{T} is the external torque per unit area. LOP02 chose a complex representation where the warp is given by $W = l_x + i l_y$ and the internal torque as $G = G_x + i G_y$. This allows Equations 3.1 and 3.2 to be rewritten as

$$\Sigma R^2 \Omega \left[\frac{\partial W}{\partial t} - i \left(\frac{\Omega^2 - \Omega_z^2}{2\Omega} \right) W \right] = \frac{1}{R} \frac{\partial G}{\partial R}, \quad (3.3)$$

$$\frac{\partial G}{\partial t} - i \left(\frac{\Omega^2 - \kappa^2}{2\Omega} \right) G + \alpha \Omega G = \frac{PR^3 \Omega}{4} \frac{\partial W}{\partial R}. \quad (3.4)$$

These equations describe the propagation of a warp in the linear regime, and were solved numerically by LOP02 to find the steady state shape of the disc around a Kerr black hole. In this case the apsidal and nodal precession frequencies in the disc (scaled by Ω) can be approximated to first order from the Kerr metric as (Kato, 1990)

$$\eta_{\text{LOP}} = \frac{\kappa^2 - \Omega^2}{2\Omega^2} = -\frac{3}{2} \frac{R_s}{R}, \quad (3.5)$$

$$\zeta_{\text{LOP}} = \frac{\Omega_z^2 - \Omega^2}{2\Omega^2} = -\frac{a}{\sqrt{2}} \left(\frac{R_s}{R} \right)^{3/2}, \quad (3.6)$$

where $R_s = 2GM/c^2$ and a is the black hole spin. These frequencies are used in the solution by inserting them directly into Equations 3.3 and 3.4. An example of the solution by LOP02 is shown in the left panel of Figure 3.1, with the same parameters used in their work. The steady state solution is formed from the interaction of the ingoing and the outgoing bending waves, where the outgoing waves are created by the reflection of the ingoing waves at the inner boundary (LOP02). Here the oscillatory behaviour of the steady state near the inner edge is clear, as is the non-zero tilt at the inner edge.

It is known that the relative signs of the apsidal and nodal frequencies determines whether the solution is oscillatory or not (Ivanov and Illarionov, 1997). The frequencies used by LOP02 have the same sign, leading to radial oscillations in the steady state tilt profile. We confirm that the oscillatory profile is dependent only on the signs of the frequencies by changing the sign of ζ_{LOP} (equivalent to modelling a retrograde black hole, see LOP02,

Figure 6) in the right hand panel of Figure 3.1. The two solutions evolve in the same manner with the exception of the oscillations near the inner edge.

While one is free to set the precession frequencies directly when solving Equations 3.3 and 3.4, in 3D the nodal precession can be induced directly (e.g. using the post-Newtonian description of Lense-Thirring precession from a spinning black hole) and the apsidal precession (i.e. Einstein precession) arises indirectly from the central potential. Hence, it is possible for the choice of potential to preclude oscillations from the steady state solution in 3D simulations. It is then not surprising that simulations that do not take the apsidal precession into account as above also do not report tilt oscillations (Sorathia, Krolik, and Hawley, 2013). However, simulations by Nelson and Papaloizou (2000) did make use of a potential that resulted in apsidal and nodal precession frequencies with the same sign but did not find oscillations. Here we use high resolution simulations along with a potential that leads to precession frequencies of the same sign to investigate this discrepancy.

3.1.2 When does the disc break?

The derivation of Equations 3.1 and 3.2 assumes that the inclination of the disc is linear. From previous results in the non-linear regime we would anticipate that the disc may break when the external torque applied to the disc is stronger than the internal torque. Here the internal disc communication is governed by a combination of pressure and viscosity. The viscous torque that acts between successive, discrete rings in the disc is given by (Lynden-Bell and Pringle, 1974)

$$\mathbf{G} = 3\pi\nu\Sigma(GMR)^{1/2}. \quad (3.7)$$

Lense-Thirring precession causes the rings that make up the disc to precess. Per unit area on the disc, this torque is given by (e.g. Nixon et al., 2012)

$$\mathbf{T} = \frac{GM}{2a}\Sigma R^2\Omega|\sin\beta|\left(\frac{R_g}{R}\right)^3, \quad (3.8)$$

where $R_g = GM/c^2$, a is the black hole spin and β is the angle between the plane of the disc and the direction of the black hole spin. If the external torque applied to the disc is greater than the internal torque maintaining the disc, the rings will precess independently faster than the disc is able to communicate the precession (Nixon et al., 2012). This will result in the disc being separated and breaking, perhaps into differentially precessing rings. Assuming that the disc has no initial warp and that internal communication is dominated by viscosity, a comparison of the above torques predicts a maximum radius that it is possible for this to occur (Nixon, King, and Price, 2013)

$$R_{\text{break}} \lesssim \left(\frac{4a}{3\alpha}|\sin\beta|\left(\frac{H}{R}\right)^{-1}\right)^{2/3} R_g. \quad (3.9)$$

This approximate relationship places an upper bound on the breaking radius of the disc at a given angle. For the typical parameters used in this paper, we have $H/R = 0.05$, $R_{\text{out}} = 40R_{\text{in}}$, $\alpha = 0.01$ and $a = 0.9$. At the outer edge of the disc, the above relation then reduces to

$$R_{\text{break}} \lesssim 180 (\sin \beta)^{2/3} R_{\text{g}}. \quad (3.10)$$

This predicts that tearing may occur in the disc for inclinations of more than 6° . At this inclination or greater one would expect the discs to break rather than align. However, in the bending wave regime that we consider here, the internal communication is dominated by pressure. In this case we can estimate the radius at which the disc will break by comparing the sound crossing and the precession timescales in the disc. Following Nixon, King, and Price (2013) and assuming that the disc is inviscid (and hence not taking into account any wave damping) we find that

$$R_{\text{break,t}} \lesssim \left(4a |\sin \beta| \frac{R}{H} \right)^{2/3} R_{\text{g}}. \quad (3.11)$$

As Equation 3.10 predicts a smaller breaking radius than Equation 3.9, we expect the disc to break closer to the black hole than Equation 3.9.

3.1.3 Can the disc accrete misaligned?

A further assumption made in developing Equations 3.1 and 3.2 was that the viscous timescale in the disc is much larger than any other timescale, equivalent to assuming that the disc is replenished from radii outside the computational domain, or $\alpha \ll H/R$ (Lubow, Ogilvie, and Pringle, 2002). This implies that the surface density profile does not change during the evolution of these equations, which is valid until the warp reaches the outer boundary. We can quantify this approximation during the evolution of the equations by considering the ratio of the wave and viscous timescales (as in Lodato and Pringle, 2006; Facchini, Lodato, and Price, 2013)

$$\frac{t_{\text{wave}}}{t_{\nu}} = \frac{2R\nu}{c_s R^2} = 2\alpha \frac{H}{R}. \quad (3.12)$$

For $\alpha < H/R \ll 1$, the above relation implies that the viscous time is much greater than the sound crossing time that the warp communicates, so we can neglect mass accretion. Indeed, LOP02 neglected the evolution of the surface density profile completely in their solution, equivalent to assuming no mass is accreted at all. However for their disc $H/R = 0.1030$, and so it is not clear whether this assumption holds. Additionally, the viscous time can be written as

$$t_{\nu} = \frac{1}{\alpha\Omega} \left(\frac{H}{R} \right)^{-2}. \quad (3.13)$$

In this form, it is clear that increasing the aspect ratio of a disc results in a significant decrease in the viscous time. At a given radius R , when the viscous timescale is comparable

to (or smaller than) the precession timescale at that same radius, accretion dominates. In thick discs (or tori) it may then be possible for the material in the outer disc to be accreted before it has a chance to align. The tilt profile in this case will not reach a steady state but instead be determined by the inward flux of angular momentum. Thus in relatively thick discs no tilt oscillations would be expected.

3.2 Numerical Method

3.2.1 Potentials

We make use of two gravitational potentials in this work. The first was previously introduced by Nelson and Papaloizou (2000), referred to as the Einstein potential (see their Equation 8). In our notation, with r as the spherical radius, it is given as

$$\Phi_{\text{E}}(r) = -\frac{GM}{r} \left(1 + \frac{3R_{\text{g}}}{r} \right). \quad (3.14)$$

This potential was introduced because it prevents the gravitational force tending to infinity as the radius decreases. However, it also results in the correct apsidal precession frequency at large distances from the black hole and has the same sign as the nodal frequency (Nelson and Papaloizou, 2000). This is in contrast to the standard Keplerian potential

$$\Phi(r) = -\frac{GM}{r}. \quad (3.15)$$

The standard potential (3.15) was used in all of the non-linear simulations, except for Figures 3.2–3.5 where (3.14) was used.

While a full GR treatment would be preferred, it is outside the scope of this work. However, the precession frequencies modelled using the Einstein potential are within $\sim 10\%$ outside of $\sim 10R_{\text{g}}$ (we demonstrate this later, see Figure 5.1). Inside of this radius, care needs to be taken as both the apsidal and nodal precession frequencies underestimate the values predicted by the Kerr metric (the worst is the apsidal precession at the ISCO, which is roughly half of the expected value). By contrast, we will demonstrate that the Keplerian potential (i.e. Equation 3.15) is less accurate in modelling the precession frequencies.

3.2.2 Precession Frequencies

We calculate the apsidal and nodal precession frequencies in our disc using the standard (Newtonian) definitions for the epicyclic and vertical frequencies,

$$\kappa^2 = 4\Omega^2 + r \frac{d}{dr}(\Omega^2), \quad (3.16)$$

$$\Omega_z^2 \equiv \left. \frac{\partial^2 \Phi(r)}{\partial z^2} \right|_{z=0}. \quad (3.17)$$

Following Nelson and Papaloizou (2000), we compute these using an effective potential that takes in to account the right hand side of Equation 2.71,

$$\Phi_{\text{eff}}(r) = \Phi(r) + \frac{4S\sqrt{GM}}{5r^{5/2}} - \frac{6Sz^2\sqrt{GM}}{r^{9/2}}. \quad (3.18)$$

This potential accounts for both the normal Keplerian potential and the correction due to the $\mathbf{v} \times \mathbf{h}$ term, where $\Phi(r)$ could be represented by either Equation 3.14 or 3.15. Firstly considering the Einstein potential, using Equations 3.16 and 3.17 the post-Newtonian apsidal and nodal precession frequencies (scaled by Ω) are given by

$$\eta_E = \frac{-1}{2\Omega^2} \left[\frac{6GMR_g}{r^4} - \frac{3S\sqrt{GM}}{r^{9/2}} \right], \quad (3.19)$$

$$\zeta_E = \frac{-2S\sqrt{GM}}{r^{9/2}\Omega^2}, \quad (3.20)$$

where Ω^2 is given by

$$\Omega_E^2 = \frac{GM}{r^3} \left(1 + \frac{6R_g}{r} \right) - \frac{2S\sqrt{GM}}{r^{9/2}}. \quad (3.21)$$

As the signs of the apsidal and nodal precession frequencies here are the same throughout the disc, this potential will allow an oscillatory profile to develop. Considering now a standard post-Newtonian potential, given by Equation 3.15, we find the apsidal and nodal precession frequencies to be (again scaled by Ω)

$$\eta_{\text{PN}} = \frac{3S}{2\sqrt{GM}r^{3/2} - 4S}, \quad (3.22)$$

$$\zeta_{\text{PN}} = \frac{-4S}{2\sqrt{GM}r^{3/2} - 4S}. \quad (3.23)$$

Here we note an important difference with respect to the solution used by LOP02. As the signs of the precession frequencies here are opposite, the steady state tilt profile will not have oscillations if the potential in Equation 3.15 is used.

3.2.3 Initial Conditions and Scope

Unless otherwise stated, the discs presented all made use of 10^7 particles, and simulations with 10^6 and 10^5 were also conducted to check convergence. We note that each time the resolution is changed between simulations with otherwise the same parameters, the artificial viscosity is altered according to the scaling described in Lodato and Price (2010) so that the discs have the same α_{SS} independent of which resolution is used. The locally isothermal sound speed in the disc was set to $c_s(R) = c_{s,\text{in}}(R/R_{\text{in}})^{-q}$ and the surface

Simulation	β ($^\circ$)	a	α	α_{AV}
PS1	30	0.1	0.01	0.395
PS2	30	0.1	0.03	1.186
PS3	30	0.3	0.01	0.395
PS4	30	0.3	0.03	1.186
PS5	30	0.5	0.01	0.395
PS6	30	0.5	0.03	1.186
PS7	30	0.7	0.01	0.395
PS8	30	0.7	0.03	1.186
PS9	30	0.9	0.01	0.395
PS10	30	0.9	0.03	0.186
A1	0	0.9	0.01	0.395
A2	15	0.9	0.01	0.395
A3	30	0.9	0.01	0.395
A4	45	0.9	0.01	0.395
A5	60	0.9	0.01	0.395
A6	90	0.9	0.01	0.395
A7	120	0.9	0.01	0.395
A8	150	0.9	0.01	0.395

TABLE 3.1: Simulation parameters, including the spin (a) and Shakura and Sunyaev (1973) viscosity (α) and artificial viscosity (α_{AV}). Unless otherwise noted, the accretion discs also had $H/R = 0.05$, an outer radius of $40R_{\text{in}}$ and made use of 10^7 particles.

density profile $\Sigma(R) = \Sigma_{\text{in}}(R/R_{\text{in}})^{-p}$, where $p = 3/2$ and $q = 3/4$ to give a constant α viscosity in the disc and uniform resolution (Lodato and Pringle, 2007). Each disc was initially set up aligned to the black hole spin, with the particles arranged using a Monte Carlo placement method. Each particle was then rotated (with respect to the centre of mass) by the inclination angle and assigned a velocity according to the following expression derived from Equation 2.71,

$$v_\phi = \frac{v_k^4}{c^3} \left[\sqrt{a^2 + \frac{R^3}{R_g^3}} - a \right] \cos(\beta), \quad (3.24)$$

where v_k is the Keplerian orbital velocity. The discs were therefore initially tilted to the black hole spin, but not warped. The results presented below have time shown in orbits at the inner edge and show the tilt as a function of radius only. This was found from the simulations using the method outlined in Section 3.2.6 of Lodato and Price (2010) where we used $N = 300$ spherical shells (see Chapter 2). For all of the simulations the inner radius was set as $R_{\text{in}} = 4R_g$, in order to compare to the LOP02 1D code. At small radii we note that the absence of GR limits the validity of our results, and indeed the need to carefully account for relativistic effects is one of our findings.

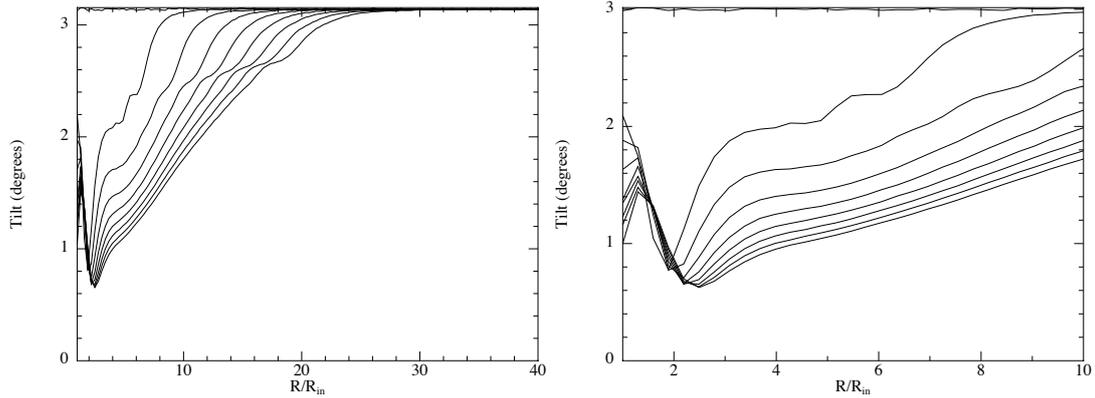


FIGURE 3.2: Time evolution of the angle between the disc plane and the black hole spin as a function of radius in a 3D disc subject to Lense-Thirring precession, using similar parameters to Lubow, Ogilvie, and Pringle (2002) and the same times as in Figure 3.1. The shape of this profile depends sensitively on the surface density at the inner edge and hence on resolution (see Figure 3.4). The right panel shows a zoom-in of the inner disc.

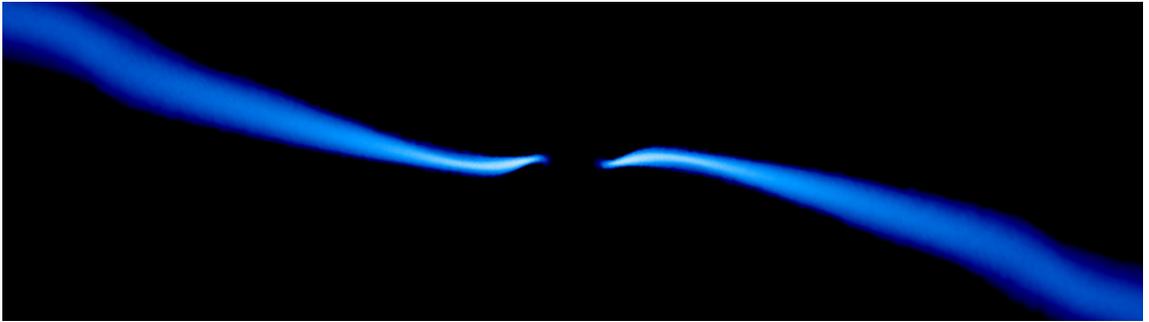


FIGURE 3.3: Cross-section view of the steady-state tilt oscillation formed in a disc initially inclined at 15° to the black hole spin (spin axis is vertical with respect to the page, i.e. along the z axis). The colour scale shows density with white being high density. Disc parameters are the same as in Figure 3.2, but with a larger initial inclination.

3.3 Results

3.3.1 Tilt Oscillations

We first investigated whether or not the tilt oscillations predicted by Lubow, Ogilvie, and Pringle (2002) are physical using 3D simulations. The disc is initiated with a constant misalignment of 3° , within the linear regime required by Equations 3.1 and 3.2. We chose parameters for our simulation similar to that of Lubow, Ogilvie, and Pringle (2002), with the exception of the surface density profile, the black hole spin and the disc thickness. Additionally, we made use of the Einstein potential outlined in Equation 3.14, in order to get precession frequencies of the same sign (Equations 3.19 and 3.20). We set $p = 1.5$ and $q = 0.75$ so that the disc is uniformly resolved, as discussed in Section 3.2.3. The disadvantage is that this results in lower amplitude oscillations in the 1D code. To combat this we encourage larger amplitude oscillations by increasing the spin to $a = 0.9$ and decreasing the disc thickness to $H/R = 0.05$. The evolution of the tilt as a function of radius is shown in Figure 3.2 from a simulation employing 10^7 particles.

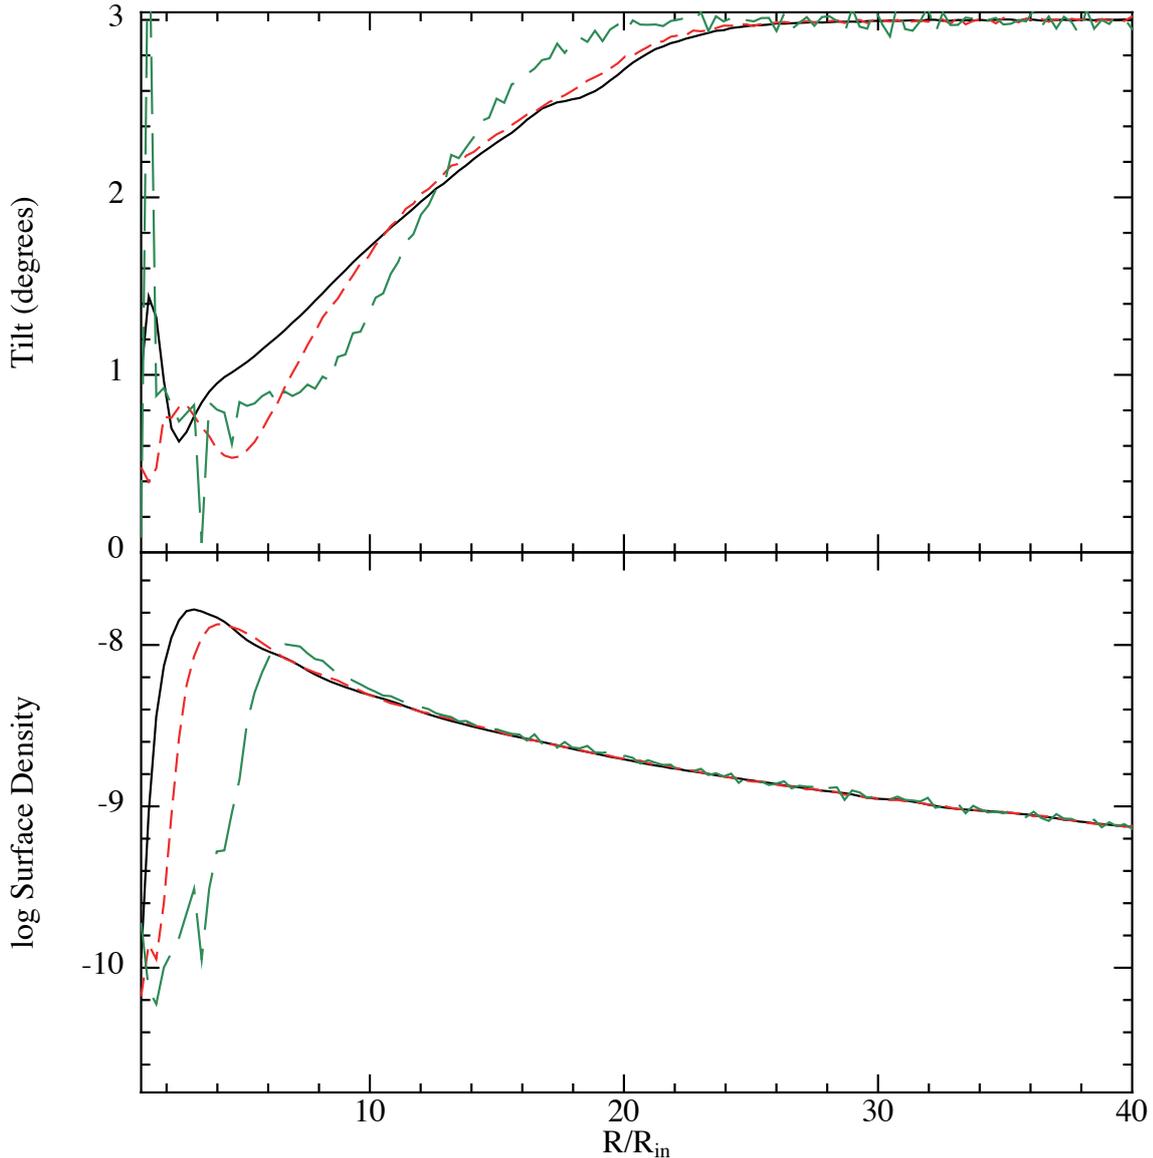


FIGURE 3.4: Resolution study showing the inclination (tilt angle) and surface density (in units of c^4/G^2M) as a function of radius at the final time of the disc shown in Figure 3.2, using 10^5 (long dashed green), 10^6 (short dashed red) and 10^7 (solid black) particles. Increasing the resolution better resolves the surface density profile at the inner edge, which strongly affects the final tilt profile found.

One of the main differences between this simulation and those conducted by Nelson and Papaloizou (2000) is the angle of inclination. While our 3° initial tilt was well in the linear regime required by the analytic description in Equations 3.1 and 3.2, the minimum inclination used by Nelson and Papaloizou (2000) was 10° . We explore the effect of non-linear inclinations in this potential by misaligning the same disc at 15° . Figure 3.3 shows a cross section of density in the inner disc from this calculation. The tilt profile after ~ 600 orbits is qualitatively similar to Figure 3.2, showing the same evolution. The quantitative tilt evolution is resolution-dependent, but nevertheless a non-zero tilt and oscillations were found at both medium (10^6 particles) and high (10^7 particles) resolution.

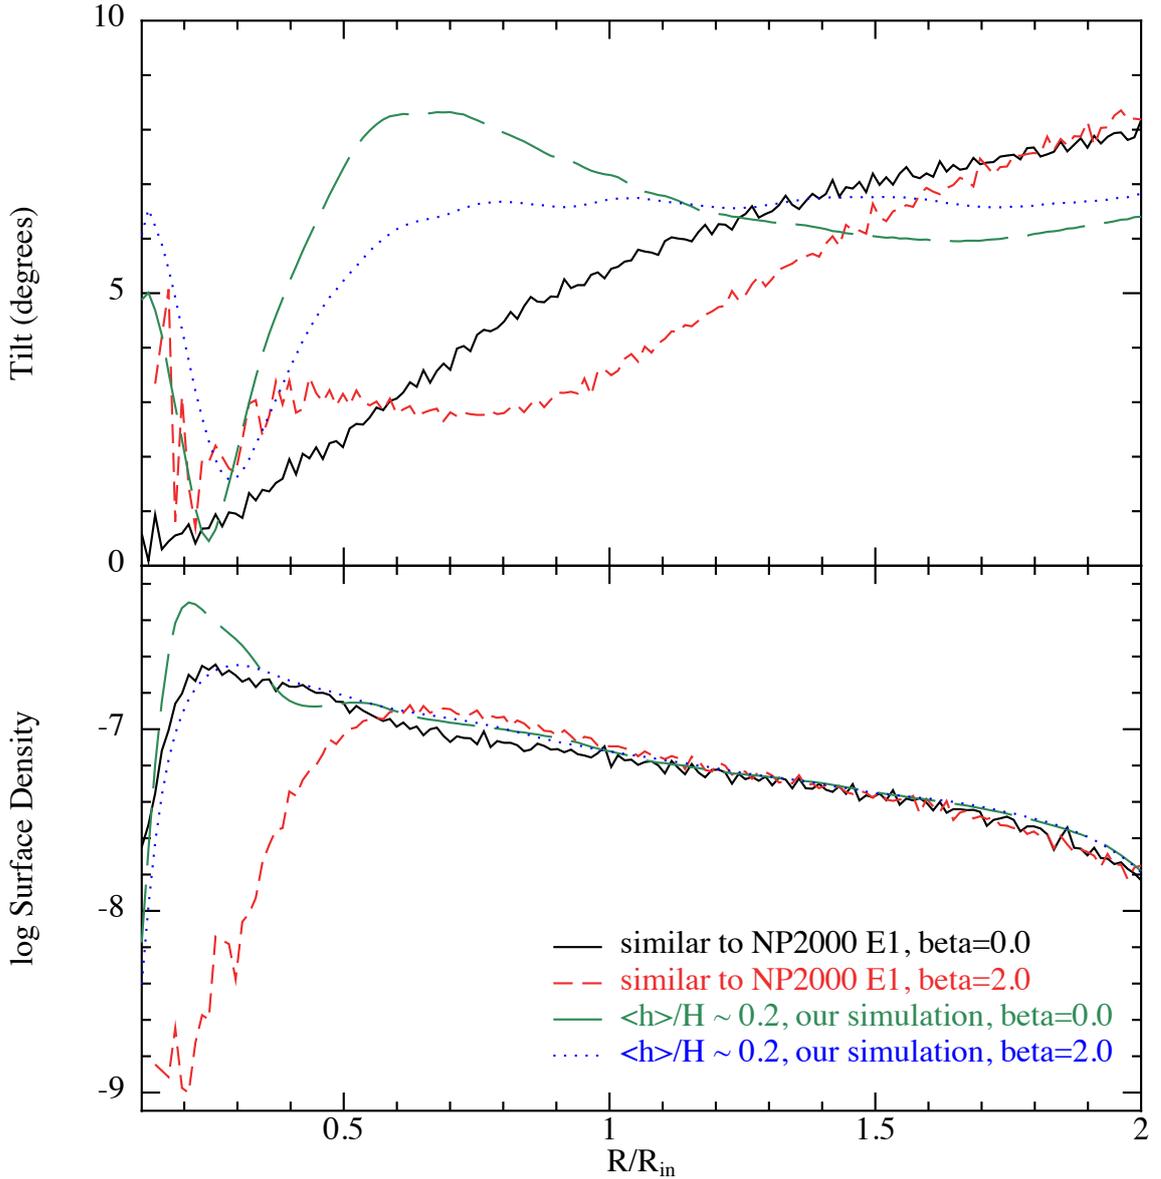


FIGURE 3.5: The effect of bulk viscosity on the final tilt of the accretion disc. The solid, black line uses the same parameters as simulation E1 of Nelson and Papaloizou (2000), with 5.2×10^4 particles and has no bulk viscosity. Bulk viscosity is used with both the red (short dashes, 5.2×10^5 particles) and green (long dashes, 5.2×10^6 particles) lines. At higher resolution and with bulk viscosity tilt oscillations are resolved, but the innermost parts of the disc remain unconverged.

Here $R_g = 0.04$ and the disc has been evolved until the warp has reached the outer edge.

Figure 3.4 shows a resolution study of the tilt and surface density profiles using 10^5 , 10^6 and 10^7 particles. The main artefact of low resolution is that accretion occurs faster and as a result there is less mass at the inner edge in the lower resolution calculations. Comparison of the surface density profiles indicates that $\Sigma(R)$ is not fully converged near the inner edge, which has a dramatic effect on the tilt profiles. However, in all discs there is a non-zero tilt at the inner edge and in the 10^6 and 10^7 discs radial oscillations are observed. The wavelength of these oscillations is consistent with the criteria given by Lubow, Ogilvie, and Pringle (2002). Even using the precession frequencies and surface density profiles in the 1D code that are appropriate to the 3D simulations (Equations 3.19 and 3.20) still does

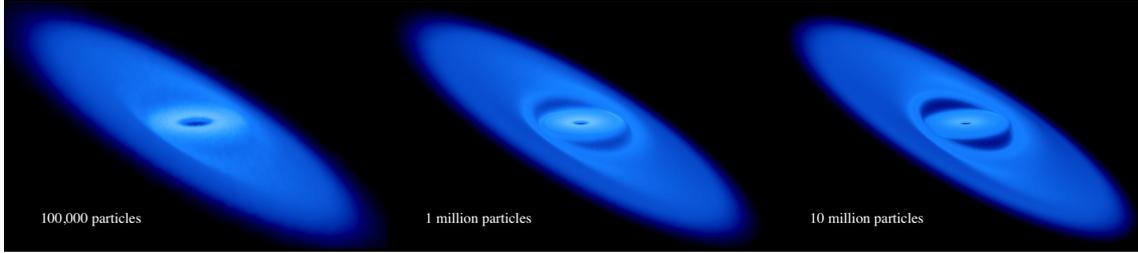


FIGURE 3.6: Structure of the disc with $a = 0.9$ and $\alpha = 0.03$ at increasing numerical resolution (left to right). At low resolution the Bardeen-Petterson Effect is observed, similar to the results of Nelson and Papaloizou (2000), but at high resolution the disc distinctly tears into two separate sections. The colour indicates density, with white being highest.

not provide a close match with the 3D results. It is not clear if this discrepancy is due to non-linear fluid effects, e.g. as discussed by Nelson and Papaloizou (2000), or simply requires higher resolution calculations to obtain numerically converged results. However it is clear that with the appropriate potential and system parameters, the disc can display radial tilt oscillations as predicted by Ivanov and Illarionov (1997) and LOP02.

Despite the resolution-dependence of our results, we were still able to observe tilt oscillations at resolutions used by Nelson and Papaloizou (2000) so long as Einstein precession was accounted for. We further investigated whether this might be due to the differences in the artificial viscosity parameters used, as we set the Von Neumann-Richtmyer viscosity coefficient $\beta^{\text{AV}} = 2.0$ (a term that prevents particle interpenetration and is important in describing shocks, see Price 2012) for all of our simulations whilst Nelson and Papaloizou (2000) used $\beta^{\text{AV}} = 0$. A nonzero β^{AV} viscosity is required to prevent particle penetration (Monaghan, 1989) and the absence of bulk viscosity is known to be problematic in disc simulations (Lodato and Price, 2010). Thus with $\beta^{\text{AV}} = 0$, the simulations of Nelson and Papaloizou (2000) might not have captured the wave interactions that create the tilt oscillations and the absence of bulk viscosity. To check this we conduct a low-resolution simulation equivalent to simulation E1 of Nelson and Papaloizou (2000) and $\beta^{\text{AV}} = 0$. Figure 3.5 shows the results (black solid line), compared to an equivalent simulation with $\beta^{\text{AV}} = 2$ (red dashed line) and also compared to a higher resolution simulations. At high resolution we find tilt oscillations regardless of the value of β^{AV} (green solid and blue dotted lines) but we find that using $\beta^{\text{AV}} = 0$ can indeed erase the tilt oscillations at low resolution. The lower panel of Figure 3.5 shows that this is not simply due to the effect on $\Sigma(R)$, since two of the calculations show very similar surface density profiles but rather different evolutions of the inner disc tilt.

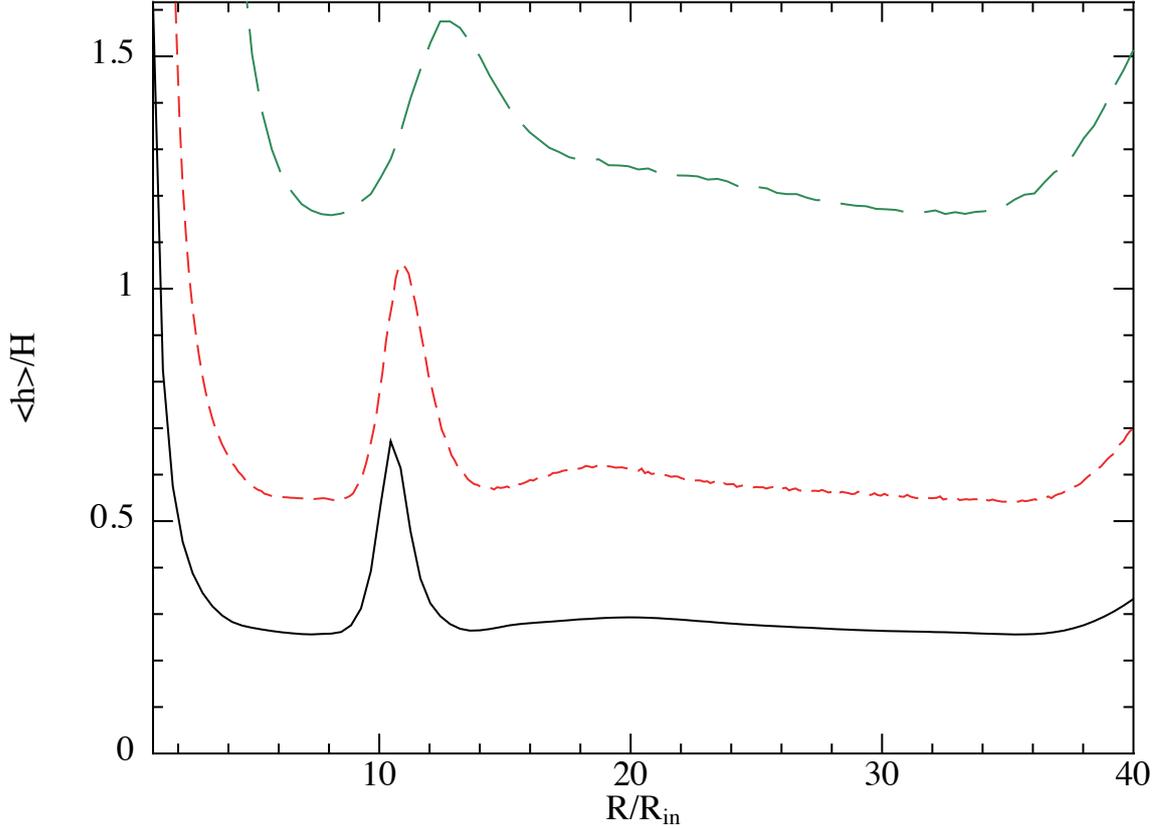


FIGURE 3.7: Resolution length as a fraction of the disc scale height ($\langle h \rangle / H$) for the three resolutions (10^5 in green long dashes, 10^6 in short red dashes and 10^7 in solid black) shown in Figure 3.6. As disc breaking occurs on length scales smaller than the scale height of the disc, the lowest resolution simulations here cannot resolve breaking behaviour (as the resolution length is greater than the scale height throughout the disc). The two higher resolution simulations are able to resolve this behaviour.

3.3.2 When does the disc break?

Bardeen-Petterson Alignment

A second possible violation of the Bardeen-Petterson picture may be that the disc breaks instead of maintaining a smooth transition between an aligned inner disc and a misaligned outer disc. In order to investigate this, we simulate a range of discs at 30° whilst varying α and a according to the list PS1-10 in Table 3.1. Here, for simplicity, we make use of a standard potential given by Equation 3.15, and hence do not expect any oscillatory behaviour.

Figure 3.6 shows a 3D rendering of density in one such simulation with $a = 0.9$ and $\alpha = 0.03$ at three different resolutions (PS10). Except for the potential used, this disc has similar parameters to simulation E3 of Nelson and Papaloizou (2000) which made use of 52,000 particles across a larger radial extent than our simulations, representing a lower resolution than any of those shown in Figure 3.6. At our lowest resolution (left panel of Figure 3.6), we also observe the inner disc aligning and smoothly transitioning to an outer, misaligned disc (see their Figure 12). However, at higher resolutions this behaviour is no

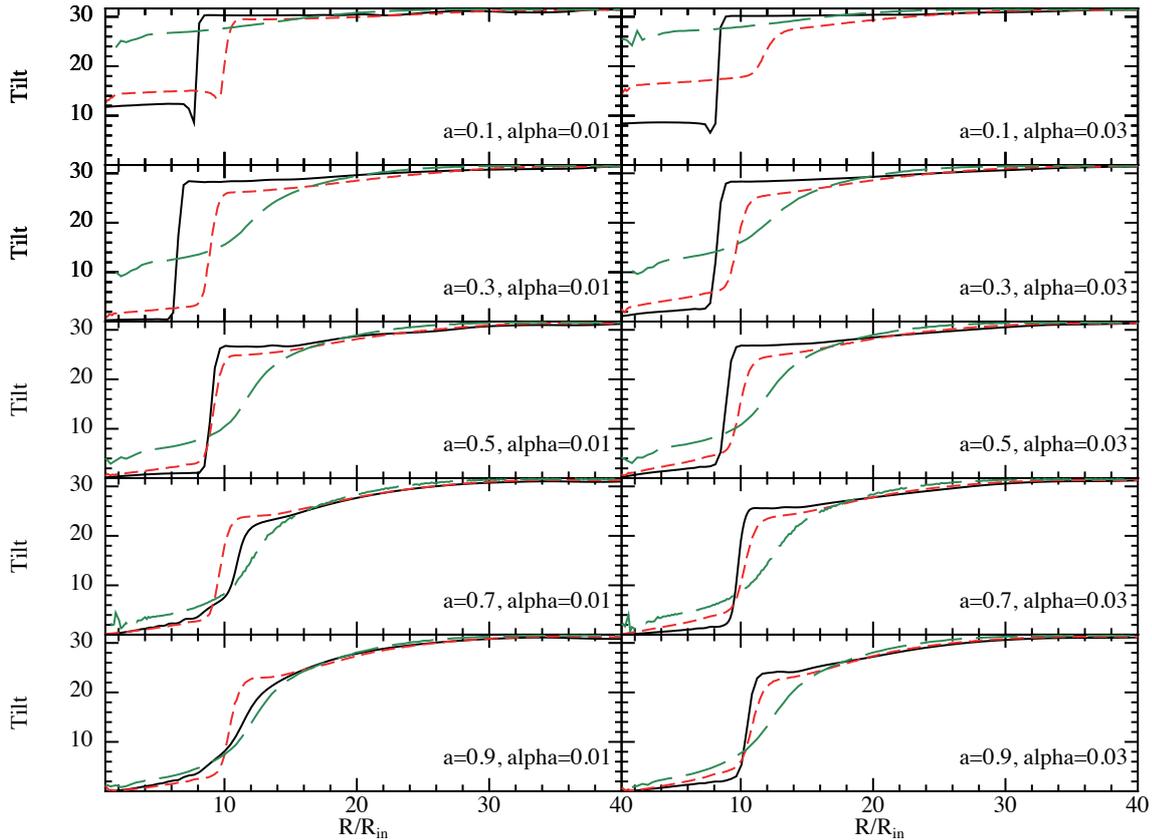


FIGURE 3.8: Disc tearing for different spin and viscosity combinations, with the same initial tilt of 30° and 10^5 particles shown in green (wide dashes), 10^6 particles shown in red (dashes) and 10^7 particles shown in black (solid). At the lowest resolution we observe the Bardeen-Petterson effect complete with a smooth transition for most discs. Increasing the resolution results in disc tearing (Bardeen-Petterson alignment) independent of our choice of viscosity or spin. As we do not include the affect of Einstein precession, we do not observe radial tilt oscillations in these discs.

longer observed — the disc instead breaks into two distinct sections, with the inner disc aligned with the black hole spin and the outer disc remaining misaligned.

Figure 3.7 shows that resolving disc breaking is mainly a question of resolving the disc scale height. For the lowest resolution simulations the resolution length is greater than the scale height of the disc and hence disc breaking (on a length scale smaller than H) cannot be resolved and a smooth transition is observed. By contrast, the two higher resolutions are able to resolve disc breaking. Figure 3.8 shows the same resolution study performed in Figure 3.6 for all of our discs at 30° , where the green line shows simulations that made use of 10^5 particles, red shows 10^6 and black shows 10^7 particles. The discs are shown after 1500 orbits at the inner edge, allowing the warp to propagate all the way to the outer radius. Increasing the spin of the black hole increases the rate at which the innermost part of the disc aligns.

Across all of the parameters chosen here, increasing the resolution changes the behaviour from the smooth tilt profile observed by Nelson and Papaloizou (2000) to a steepening of the tilt profile and ultimately a disc that is broken into distinct sections. The higher resolution results show an aligned inner edge, a misaligned outer edge and a sharp tilt

profile connecting these, representing a break in the disc. The discs simulated with lower spins appear to steepen and tear faster than those with higher spin as the break occurs further out (and hence a longer precession time). This is observed most clearly between the low viscosity, high spin cases. At $a = 0.5$, the tilt steepened and the disc tore before the end of the simulation. For the disc with $a = 0.7$, the tilt began steepening near the end of the simulation but was not able to separate, whilst at $a = 0.9$ the disc has not yet begun steepening. We have confirmed that this is the case by extending the high resolution simulations of the $a = 0.9$ case, and indeed observed steepening to occur at later times.

As with the previous simulations, Figure 3.8 demonstrates that the simulations are not fully converged, especially when considering the low spin cases ($a < 0.5$). For these discs, the discrepancy in the inner tilt is again due to the mass accreted at the inner edge of the disc. At low resolutions the inner part of the disc is accreted faster, resulting in less mass near the inner edge. The same Lense-Thirring torque then acts on less mass, and is thus not able to align the disc to the same extent. At increasing spins this effect is observed less, as the higher spin provides a larger torque and so even the lowest resolution discs are able to align. In the discs with $a < 0.5$, increasing the resolution leads to a more distinct tear in the disc suggesting that our results are consistent. Hence we can be confident that these discs do tear, and present an upper limit on the radius at which this occurs. As the tearing occurs outside of the radius where oscillations were found in Section 3.3.1, using similar parameters, this behaviour should not be affected by our choice of potential.

Disc Tearing

To investigate the dependence of disc tearing on the misalignment between the disc and the black hole spin in the wavelike regime we simulated a suite of discs at different inclinations. We again make use of the traditional post-Newtonian approximation given by Equation 3.18. We held $\alpha = 0.01$ and $a = 0.9$ constant and varied the inclination of the disc between 0° (aligned) and 150° , noted in Table 3.1 with A1-8. Figure 3.9 shows these simulations after more than 1500 orbits measured at the inner edge. Each disc was initially tilted but not warped. As the simulation progressed, a warp evolved in response to the Lense-Thirring torque and in the higher inclination cases resulted in the disc breaking.

At 15° (top right of Figure 3.9) the disc was observed to smoothly align to the spin of the black hole. At the end of the simulation, the tilt of the disc was consistent with the Bardeen-Petterson effect and is similar to results seen in previous simulations at 10° by Nelson and Papaloizou (2000). Extending the lower resolution version of this simulation (with 10^6 particles) for twice as long shows that the disc continues to align with the black hole spin, implying that the steady state for this disc is full alignment. Inclining the disc at 30° also did not yet result in disc tearing, however this is because for this particular choice of viscosity and spin this simulation has not been run long enough (see Section 3.3.2)

For discs at higher inclinations ($\gtrsim 45^\circ$; second, third and fourth rows of Figure 3.9), the inner section of the disc was found to align within 50 orbits and a smooth transition was

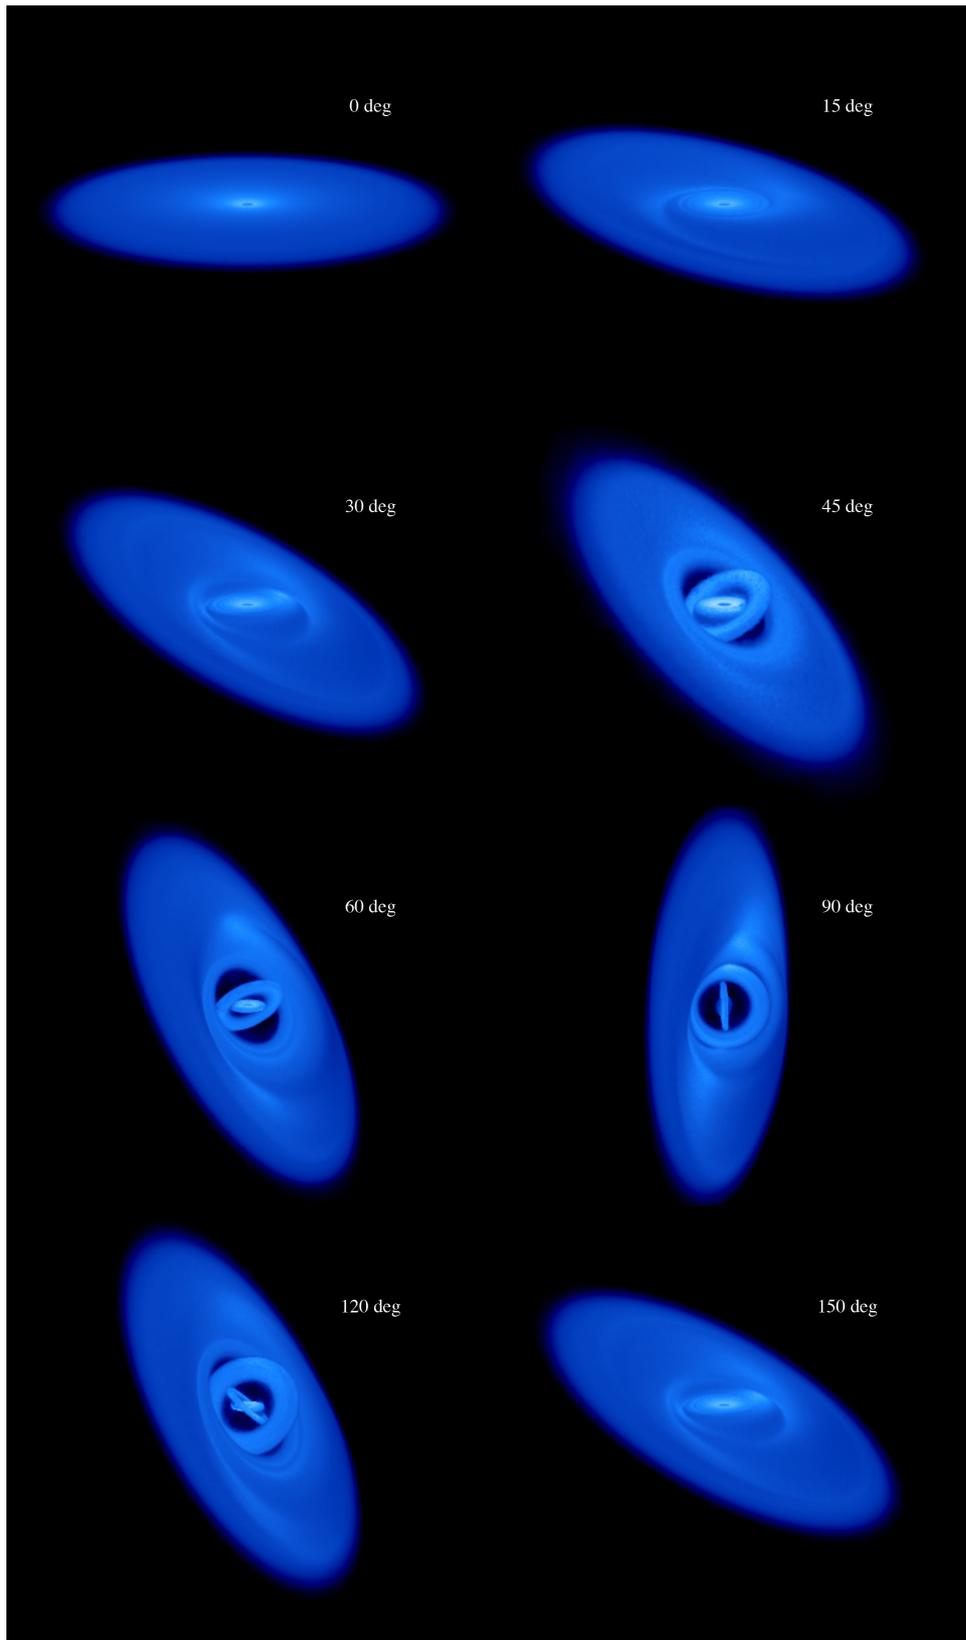


FIGURE 3.9: 3D renderings of discs that were initially misaligned with the black hole spin at various angles, with each simulation using 10^7 particles and shown after ~ 1500 orbits. The inability of the discs inclined by more than $\beta \gtrsim 45^\circ$ to communicate the Lense-Thirring precession causes the formation of discrete rings which ‘tear’ and precess effectively independently before undergoing direct cancellation of angular momentum and rapid accretion. The black hole spin in each of these images is vertical with respect to the page (i.e. along the z axis). The same density scale is used as in Figure 3.6.

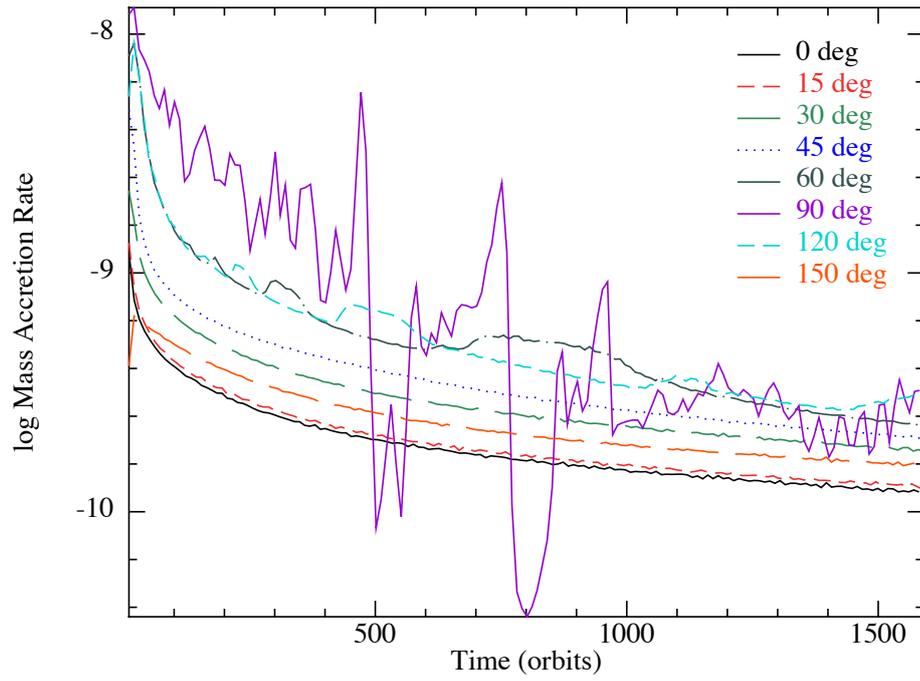


FIGURE 3.10: Instantaneous mass accretion (a direct output of PHANTOM) of the 0° , 15° , 30° , 45° , 60° , 90° , 120° and 150° discs run with 10^6 particles and time measured in orbits. The bin width is 10 times the orbit timescale and a logarithmic scale is used for convenience. In line with previous results, inclining the disc to the black hole results in mass accreting faster by almost an order of magnitude.

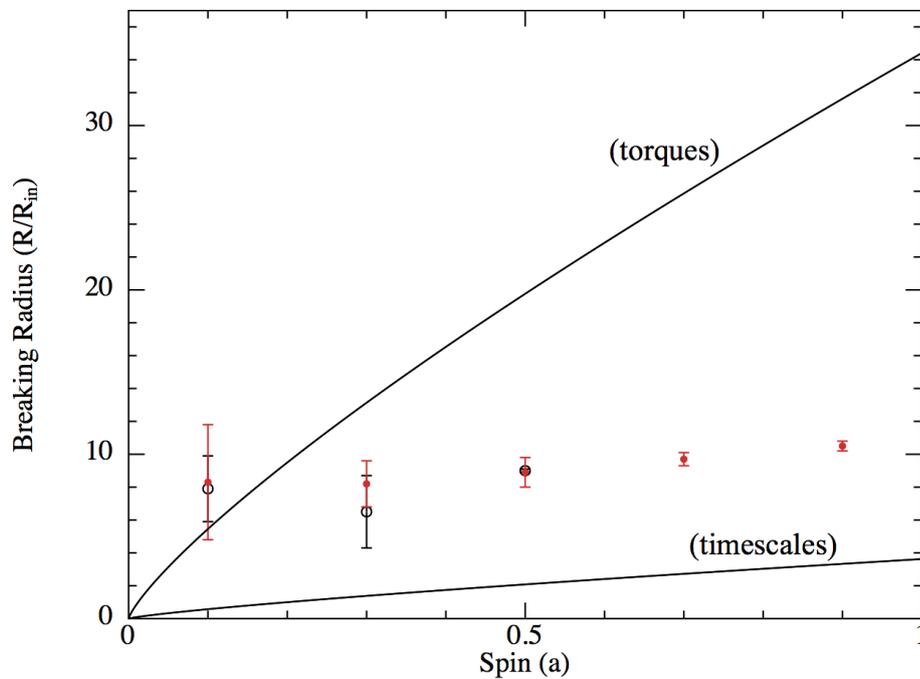


FIGURE 3.11: Comparison of the breaking radius measured from the discs inclined at 30° with our prediction of R_{break} (upper; assuming that $\alpha = 0.02$) found by considering the torques in the disc (Equation 3.9) and $R_{\text{break,t}}$ (lower) by comparing the sound crossing and precession timescales (Equation 3.11). Here the black points represent simulations using 10^7 particles and the red points simulations with 10^6 particles.

formed between this and the outer region of the disc. This transition then steepened until the disc broke into two sections that were connected by precessing rings of material. Multiple rings of material were torn off from the outer, misaligned disc and each was observed to precess effectively independently. Towards the end of the simulations, up to two rings were precessing at the same time (for example, 120° disc of Figure 3.9) and were present for up to ~ 400 orbits. Eventually each of these rings settled with and increased the inner, aligned region of the disc. The disc inclined at 90° (right hand panel in third row of Figure 3.9) also developed precessing rings of material that were accreted. However, for this inclination no inner aligned disc was observed and the rings of material were accreted directly onto the black hole.

Figure 3.10 shows the instantaneous mass accretion rate by the discs at different angles (one of the default outputs that PHANTOM generates, calculated by the particles that move inside the specified inner edge). It can be seen that inclining the disc to the spin of the black hole increases the rate of accretion by more than an order of magnitude when compared to an aligned disc, similar to previous findings (Nixon et al., 2012). The discs that form an inner aligned disc and precessing rings have even higher accretion rates, as the inner disc is continually fed by the rings as they align. In each case except the disc inclined at 90° , either the inner region aligns or a radially small, inner aligned disc forms (e.g. 45° , 60° and 120°). The aligned inner region provides constant mass accretion to the black hole, with any features from the interaction of the broken discs superimposed on this (e.g. the complementary behaviour of the 60° and 120° cases in Figure 3.10). The 90° case represents an exception to this; as no innermost aligned disc forms all the mass accreted is a result of the interaction between the rings, leading to a chaotic mass accretion rate. When taken in context with the results in the diffusive regime (Nixon et al., 2012), Figure 3.10 implies that regardless of whether the disc is thin or thick, mass accretion is faster when the disc is inclined.

Disc tearing has also been observed in the wavelike disc regime for circumbinary discs inclined at high angles (Facchini, Lodato, and Price, 2013). In a simulation of a circumbinary disc inclined at 60° , their disc separates into two sections and the inner one precessed effectively independently of the outer disc. As their disc is thicker than ours ($H/R = 0.1$) and has a higher viscosity ($\alpha = 0.05$), we would anticipate that a strong external torque would be required to tear the disc, and we observe their disc does tear at a smaller radius than any of ours.

Location of tearing radius

The disc is expected to tear when the Lense-Thirring torque is larger than the internal communication in the disc. If the internal communication in the disc is governed by viscosity, the torques given in Section 3.1.2 can be used to estimate the upper breaking radius given in Equation 3.9. However in our simulations the disc internal dynamics are dominated by pressure rather than viscosity, hence Equation 3.11 may be more appropriate.

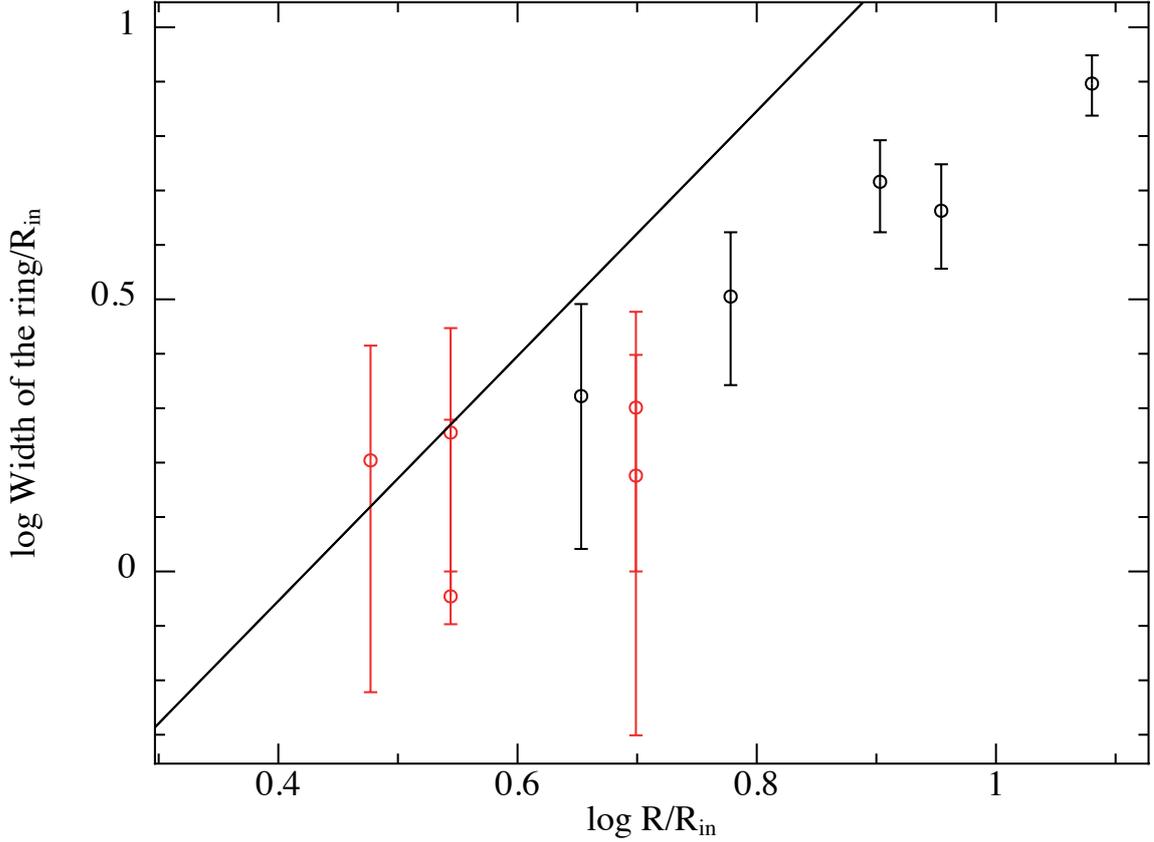


FIGURE 3.12: The solid line shows the expected width of each ring of gas torn off in our simulations, calculated by comparing the precession timescale to the distance that the wave can travel (Equation 3.26). The circles show the measured ring widths from the simulations, where black circles indicate short lived rings and red circles rings that are stable for more than ~ 20 orbits.

As the Lense-Thirring torque has a radial dependence, it is largest in the inner most parts of the disc and it is reasonable that these discs will break at a radius smaller than predicted by Equation 3.9. Figure 3.11 shows a comparison of the estimated break radius for the simulations inclined at 30° compared to the prediction from Equation 3.9 (upper line; assuming that $\alpha = 0.02$, the average for our simulations) and from Equation 3.11 (lower line). We find that the disc does break at radii lower than our prediction from the viscous torques alone, and that the breaking radius is intermediate between the predictions from Equations 3.9 and 3.11, indicating that the torques in our discs lie between these two extremes. The increasing uncertainties at low spin correspond to the decreasing convergence of our simulations due to mass accretion at the inner edge, seen in Figure 3.8.

The discrepancy between the predicted and the observed breaking radius appears to occur at all inclinations. Using Equation 3.9, the breaking radius for the 60° disc is found to be $R_{\text{break}} \sim 41R_{\text{in}}$ which is greater than R_{out} . However this disc is observed to break (at $R \lesssim 18R_{\text{in}}$), in line with the results of Figure 3.11. If we now consider the 15° disc, it is predicted to break at $R_{\text{break}} \sim 18R_{\text{in}}$ but from the simulation we do not observe tearing. This could occur if the actual tearing radius is less than R_{in} , consistent with the previous results.

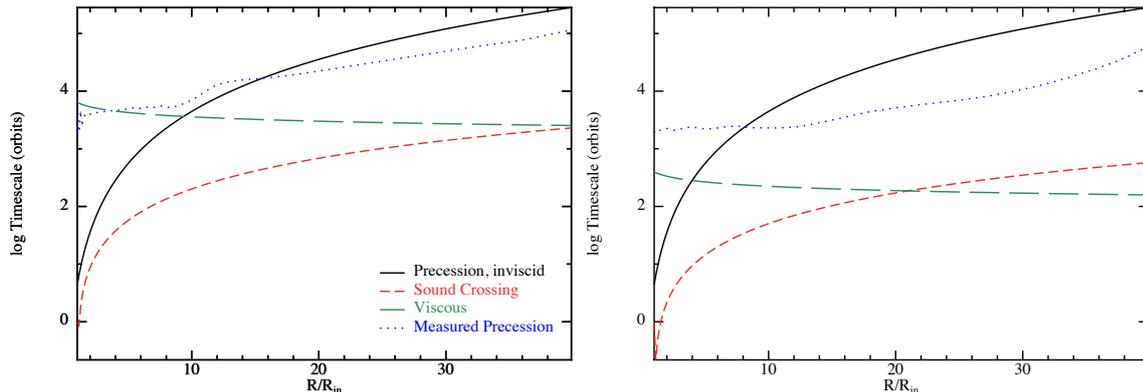


FIGURE 3.13: The timescales in our PS9 simulation (left) and a disc that is four times thicker (right). The precession timescale does not change much between the thin and thick discs, however the sound crossing and viscous timescales decrease as the disc becomes thicker. The decrease in viscous time means that the thicker disc is able to accrete misaligned material, preventing the development of a steady state.

Width of the rings

The rings that are torn off during the simulations appear to be much wider than those found in the diffusive regime (Nixon et al., 2012), some up to $\Delta R/H(R) \sim 25$ (where ΔR represents the ring width). It is possible for rings to form when the disc is able to break and differential precession is present, such as when the disc is subjected to Lense-Thirring precession. We therefore expect the width of the ring to be determined by a relative comparison between the sound crossing and precessional timescales in the disc. We can approximate this by letting ΔR be the distance that a wave can travel in a precession time such that

$$\int \frac{2}{c_s} dR \propto t_p, \quad (3.25)$$

across the ring. If we assume that the inner edge of the ring is at $R_{\text{in, ring}} = R - \Delta R/2$, the outer edge at $R_{\text{out, ring}} = R + \Delta R/2$ and use the expression for the sound speed, we get

$$R^3 \propto \frac{a}{(q+1)} \left[1 - \left(\frac{R_{\text{in, ring}}}{R_{\text{out, ring}}} \right)^{q+1} \right] \frac{R_{\text{out, ring}}}{c_s(R_{\text{out, ring}})}, \quad (3.26)$$

where R is the radius that a ring of thickness ΔR occurs at. Figure 3.12 compares the width of the rings measured from the simulations to this prediction. Although there are large uncertainties in the measurements the general trend of increasing ring width with R is reproduced.

3.3.3 Can the disc accrete misaligned?

Previous simulations of tilted accretion discs in the wavelike regime have not identified disc tearing when the disc is subjected to Lense-Thirring precession. The results of these thicker discs have found that the disc warps in the inner region, with a non-zero tilt at the inner edge, and then precesses as a solid body (Fragile and Anninos, 2005; Fragile et al.,

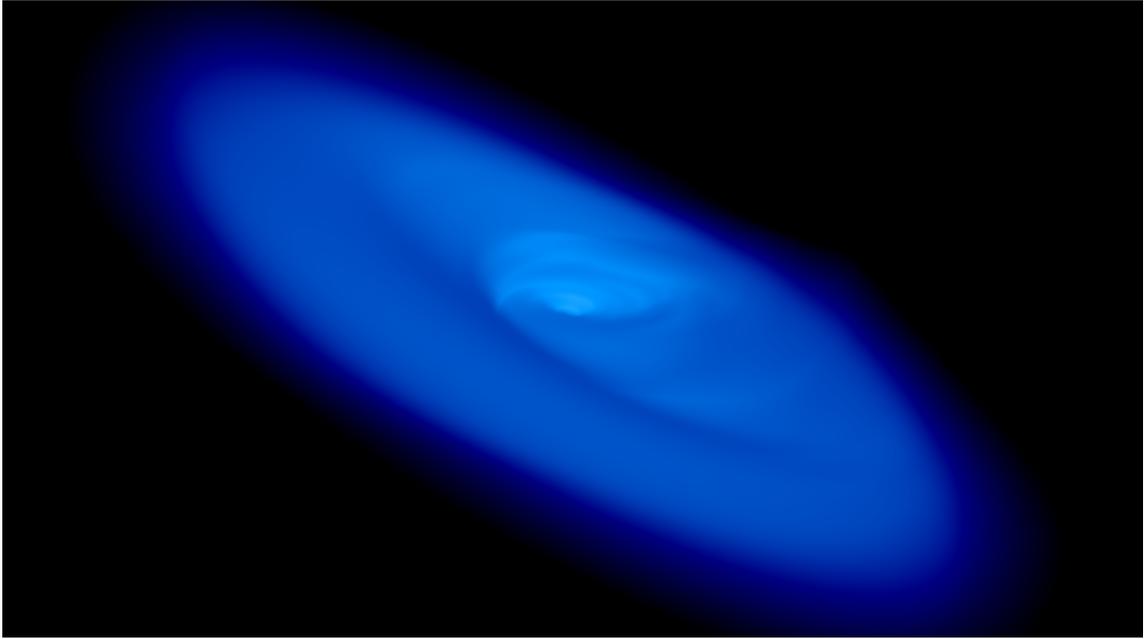


FIGURE 3.14: Our thick disc simulation, similar to that of Fragile et al. (2007) except that it is initialised at 30° and run for ten times longer. This disc is not observed to tear, as expected, but warping is observed in the inner regions and higher mass accretion than our thin disc. This figure is shown with the same density scale as Figures 3.6 and 3.9.

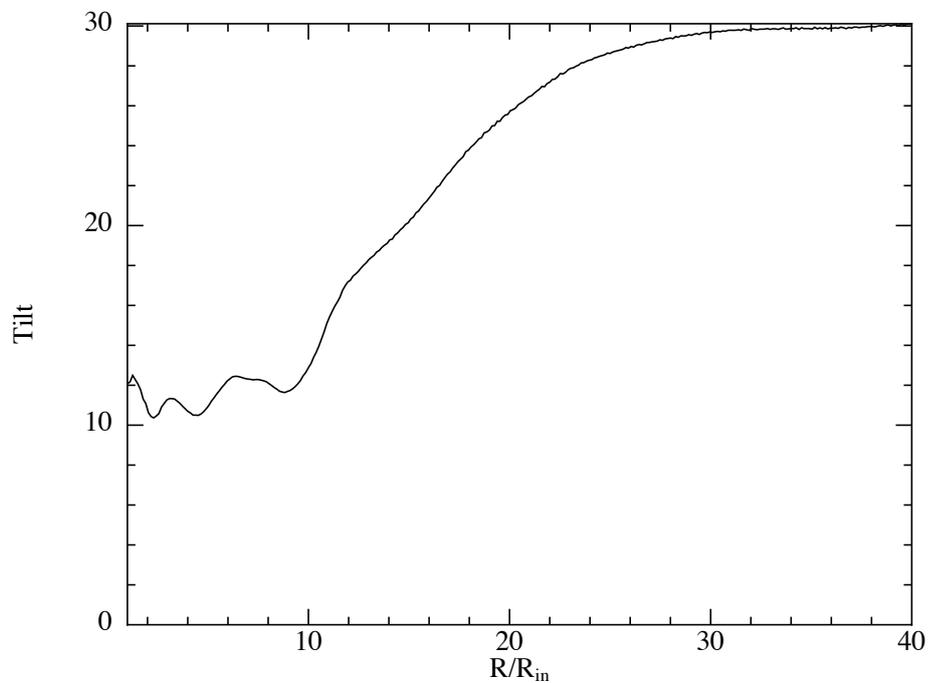


FIGURE 3.15: The final tilt profile of our thick disc simulation. Misaligned accretion occurs at the inner edge, causing a non-zero tilt and preventing a steady state from being formed. The inner edge features are not steady like the results of Figure 3.3.

2007). To examine this behaviour, we conduct a single simulation of a thick disc. We use the same parameters as PS9, but with an aspect ratio four times the initial value, such that $H/R = 0.2$ at the inner edge. This disc is similar to the simulation of Fragile et al. (2007), except that it has twice the initial tilt (and does not include magnetic fields).

The timescales for this disc are shown in the right of Figure 3.13. Comparison with the timescales from PS9 (left) shows that although the precession timescale has not changed appreciably, the viscous and sound crossing timescales have decreased substantially. In the outer half of the disc we note that the viscous timescale is the shortest, allowing the material located there to be accreted to the inner regions faster than it can align. This leads to material being accreted before it can align with the spin of the black hole, causing a non-zero tilt at the inner edge of the disc. A comparison between the mass accretion of this disc and our thinner PS9 simulation shows that there is more mass accreted by the thicker disc.

Simulating to approximately the same time as quoted by Fragile et al. (2007), we observe the thick disc to warp in the inner regions but not to tear. At this time in our thin disc simulations we also do not observe tearing, so we continue the simulation until approximately 200 orbits according to the time units specified by Fragile et al. (2007) (10 times longer than their lower resolution simulation). The results at this time are shown in Figure 3.14 and 3.15. We do not observe the large increase in the disc tilt at the inner edge that was found by Fragile et al. (2007) (see their Figure 12), however in their paper this is attributed to plunging streams which we also do not observe. Presumably this is due to our use of the post-Newtonian approximation in Equation 3.18.

As the disc is four times thicker than our simulation PS9, ν increases by a factor of 16 (even though α does not change). This increases the internal torque in the disc by the same factor (see Equation 3.7), but the external torque applied is the same as for our disc. This should make it much harder to tear the disc, and when we calculate the breaking radius using Equation 3.9 we find that it would be $R/R_{\text{in}} \sim 3$, inside the region where misaligned accretion is occurring. Indeed, from our results in Section 3.3.2, we would not expect this disc to tear at all.

3.4 Discussion

Despite using up to 10^7 particles, the simulations that have been presented are not yet converged. As shown in Figure 3.6 and comparison of our results with Nelson and Papaloizou (2000), increasing the resolution strongly affects the behaviour the disc displays. However, features like disc tearing and radial oscillations are present in both the medium and high resolution simulations and so we can draw conclusions about the qualitative behaviour. Additionally, whilst increasing the resolution decreases the breaking radius, it does so by a smaller amount each time, so we are confident that the measured tearing radii for our non-linear simulations is an upper limit and that our results are close to being converged.

The main point of discussion is why our results differ to those found by Nelson and Papaloizou (2000). The two main factors are the numerical resolution and the viscosity parameter β_{AV} . Simulations we performed at comparable resolution to Nelson and Papaloizou (2000) showed similar behaviour — namely a smooth transition between the aligned and misaligned regions. However, when we increased the resolution we found that the behaviour changes and these discs tear into two disconnected sections (the main criterion being to adequately resolve the disc scale height). This implies that low resolution is the primary reason that prevented Nelson and Papaloizou (2000) from observing disc tearing. However, we also showed that the inclusion of a β viscosity, even at low resolution, recovers steady-state oscillations in the tilt of the disc midplane with respect to the black hole spin axis similar to those predicted by the linear theory of Ivanov and Illarionov (1997) and LOP02. This is in contrast to the findings in Nelson and Papaloizou (2000), where it was suggested that the tilt oscillations were short wavelength features which could be damped out by non-linear effects. As shown by our 15° simulation and in agreement with LOP02, we found that the wavelength of the radial oscillation is of the order of the radius (R) and is not damped out by such effects.

The tilt oscillations that were found at linear inclinations do not match the description of the results obtained with the 1D code by LOP02, and increasing the resolution does not reduce the discrepancy. The difference is likely due to the 1D code assuming that the viscous timescale is negligibly large. In Section 3.3.3 it is found that the mass accretion is not necessarily negligible, as for discs with a larger aspect ratio we found it is possible for the material to accrete to the inner regions of the disc faster than it is able to align. This causes the disc to accrete misaligned material, which prevents a steady state from being formed and confirms that it is not possible to produce a tilt profile such as that described by the Bardeen-Petterson effect if the viscous time is too short (as predicted by Lodato and Pringle 2006). Recently the thinnest discs in relativistic simulations have been completed by Morales Teixeira et al. (2014), with $H/R = 0.08$. Their retrograde simulation showed partial alignment at the inner edge, but their prograde simulations displayed an inner edge tilt that was greater than the initial condition. It is also noted that the strength of the tilt oscillations depends on the disc thickness, and so thick discs (and tori) would display weak oscillations.

Despite demonstrating that the location of the disc break is resolved (see Figure 3.8), the interaction of the rings in our simulations with tearing is not. Each interface between adjacent rings is only described by a small number of particles (necessarily, as this is a region of relatively low density). As a result, the interaction of neighbouring rings is governed by poorly resolved particles with relatively large smoothing lengths. A method to improve this is not clear as increasing the global resolution will increase the sharpness of the break (e.g. Figure 3.8) but will not lead to more particles between rings (as low density is defined by few particles). Because the initial tear and breaking of the disc is resolved, this only affects our estimated mass accretion rate profiles (Figure 3.10) and the width of the rings formed in our simulations (Figure 3.12).

With the exception of the points at $a = 0.1$, the simulations in Figure 3.11 have the same dependence on the spin as in the lower black line, representing the comparison between the sound crossing and precession timescales where the radius is proportional to $a^{2/3}$. At low spin the location of the break does not appear to be related to either criterion, but the breaking behaviour here is different (e.g. see Figure 3.11, top panels). At such a low spin the external torque is much weaker, and the breaking here is potentially caused by non-linear effects other than the Lense-Thirring torque — e.g. Lodato and Price (2010) demonstrated disc breaking in a strongly misaligned disc but with no external potential or forcing.

Perhaps the main caveat of our simulations is that we use an α viscosity to model the discs. Whilst a comparison between a purely hydrodynamical disc (with no explicit viscosity) and one where the viscosity is controlled by the MRI has shown that the behaviour of the disc is largely controlled by the hydrodynamic evolution (Sorathia, Krolik, and Hawley, 2013), for a complete picture of the disc evolution we should include magnetic fields to self-consistently generate a turbulent viscosity through the MRI. However it is not yet clear how the MRI will respond in the presence of a warp, especially at large angles. Additionally, we assume that our discs are vertically isothermal. Heating of the disc due to warping may further complicate this picture (Ogilvie, 2003).

For tilt oscillations and efficient wave transport to occur we require $H/R > \alpha$ (Papaloizou and Lin, 1995; Ivanov and Illarionov, 1997). Black hole accretion discs are often expected to be geometrically thin and have $\alpha \sim 0.1$ (King, Pringle, and Livio, 2007). However, for discs which are accreting either at very sub-Eddington ($\lesssim 0.1L_{\text{Edd}}$) or near-Eddington ($\gtrsim L_{\text{Edd}}$) rates, the disc may become geometrically thick (Narayan and Quataert, 2005). So the simulations presented here may be most relevant to the low luminosity state of X-ray binaries where the disc can be thick and α may be significantly smaller than its usual outburst value (Smak, 1984; Meyer and Meyer-Hofmeister, 1984). They are also relevant to AGN accreting at rates greater than Eddington and to the discs formed in tidal disruption events where the initial star orbit can be highly misaligned and the disrupted material infalling at super-Eddington rates (see Chapter 6).

Finally we must take into account that the outer edge of the simulations presented here ($160R_g$) is much smaller than the actual size of the disc (e.g. in X-ray binaries, 1000s of R_g are expected). In simulations which have a smaller outer radius than simulated here, the breaking/tearing behaviour is not converged. As the outer radius of the disc decreases to the tearing radius, the inner gas that should tear does not have enough of a mass anchor to tear off and solid body precession is instead observed — we emphasise that this behaviour is the result of the choice of outer boundary and is thus a numerical rather than physical effect in this case. The outer radius chosen in these simulations is large enough to allow converged tearing behaviour. As breaking/tearing cannot occur outside the breaking radius (i.e. where the internal viscous stresses are much stronger than the external torque *and* the disc can communicate the precession faster than it occurs) the outer disc is only

expected to warp slowly over time. We thus expect the tearing and breaking behaviour observed in these simulations to be consistent with discs that have a larger outer boundary.

3.5 Conclusion

In this work we have re-examined the Bardeen-Petterson effect in 3D using hydrodynamical simulations of accretion discs subject to Lense-Thirring precession, in the regime where warps propagate in a wavelike manner ($\alpha \lesssim H/R$). Our detailed conclusions are as follows:

1. The Bardeen-Petterson picture of an aligned inner disc smoothly connected to a misaligned outer disc occurs only at low inclinations and only when Einstein precession is not accounted for. Using high resolution calculations, we find both steady state oscillations in the disc tilt (when Einstein precession is included) and that discs break when they are relatively thin and highly misaligned to the black hole spin.
2. We recover steady tilt oscillations for the first time in a 3D hydrodynamics code, as predicted by LOP02. However, as the 1D code developed by LOP02 assumes that mass accretion is negligible, discrepancies remain between the predicted tilt profile and our 3D results.
3. Tilt oscillations are also present at higher inclinations (15°), showing that non-linear effects do not necessarily damp this behaviour.
4. Disc ‘tearing’ or ‘breaking’, rather than a smooth transition between spin-aligned and spin-misaligned parts of the disc, appears to be an inevitable outcome for accretion discs inclined to the black hole spin by more than a few degrees. This occurs regardless of whether the propagation of bending waves is governed by pressure forces or viscous stresses.
5. Tearing of the disc leads to rings that precess effectively independently. As in the diffusive regime, this can lead to direct cancellation of angular momentum and hence faster accretion. The main difference in the wavelike regime is that the rings are wider, with the width determined by the ratio of precession to sound crossing time rather than the disc scale-height.
6. The Bardeen-Petterson effect cannot occur in discs where the viscous time is comparable to the alignment time. In this case the disc material is accreted misaligned. Hence it is possible to have discs that are misaligned with respect to the black hole spin even in the absence of tilt oscillations, but this can only occur at high \dot{M} (i.e. for thick discs).
7. Mass accretion rates can be enhanced by an order of magnitude or more when the disc is inclined with respect to the black hole spin. This occurs regardless of whether the disc is thick or thin.

Declaration for Chapter 4

Declaration by Candidate

The work presented in § 4.1 to § 4.4 has been published in Nealon, Nixon, Price, and King (2016).

I declare that my contribution of the work presented in this chapter involved the following:

The writing of the publication, conducting the 3D simulations and developing the visualisation method to compare to previous simulations.

Extent of contribution by candidate: 80%

Co-author names: Chris Nixon, Daniel Price and Andrew King

Nature of co-author's contribution: Providing codes used for simulations, guidance in parameter choice and comments and feedback during the writing of the paper.

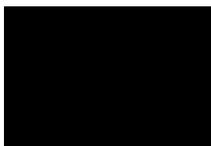
I have renumbered sections, equations and figures of the published paper in order to generate a consistent presentation within the thesis.

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the student's and co-authors' contributions to this work.

Student signature: _____

Date: November 20, 2017

Main Supervisor signature:



Date: November 20, 2017

Chapter 4

Apsidal precession, disc breaking and viscosity

The most recent simulation considering the behaviour of a tilted disc around a rotating black hole was conducted by Krolik and Hawley (2015), shown in Figure 1.12. This simulation used a grid-based numerical approach, focussing on the evolution of the disc inclined at 12° to the black hole spin. After ~ 25 orbits at $R = 10R_g$, the authors confirm the Bardeen-Petterson Effect: they find their disc has reached a steady state where the inner disc is aligned to the black hole spin and there is a smooth transition to the outermost, misaligned disc. However, these simulations did not take into account apsidal precession, an important effect discussed in the previous chapter. The relevance of this precession term, the viscosity description and the validity of the disc breaking approximation is investigated by focussing on a hydrodynamic comparison with their simulation.

First we explore the initial conditions needed to recreate their simulation. Using PHANTOM, the hydrodynamic evolution of this accretion disc is compared directly to their fully MHD driven simulation. In this initial simulation — as in theirs — the effect of apsidal precession is neglected so that any difference must be due to the viscosity formulation. In order to quantify these differences, the evolution of the disc is displayed in the same way as in their paper.

This simulation is then repeated but including the effects of apsidal precession — in Chapter 3 this effect was argued to be critical to the disc evolution. The differences between this second simulation and the initial simulation demonstrates the effect of either neglecting or considering the apsidal precession. This comparison demonstrates precisely how important it is to take these effects into account, even if it is at extra computational expense.

Finally, we consider disc breaking. The original simulation by Krolik and Hawley (2015) does not show any evidence of disc breaking, in line with the approximation discussed in the previous chapter. To confirm that disc breaking is possible, we repeat their simulation with a much higher inclination of 30° .

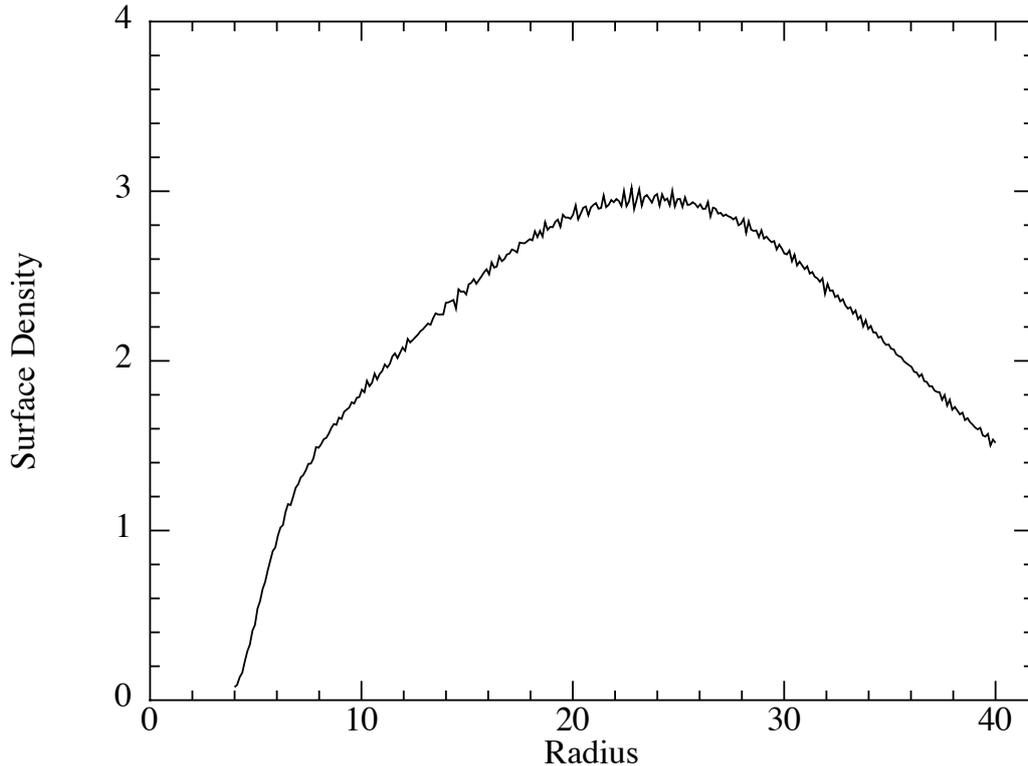


FIGURE 4.1: Surface density profile (in units of c^4/G^2M) of the initial condition used for all simulations, chosen to match KH15.

4.1 Simulations

4.1.1 Disc parameters

We match our initial disc setup to the recent simulation by KH15 for the purpose of comparison. We thus adopt an isothermal disc with sound speed set to $c_s(R) = c_{s,in}R^{-q}$ and a surface density profile $\Sigma(R) = \Sigma_{in}R^{-p}$. We conducted simulations with 10^5 , 10^6 and 10^7 particles to check for convergence, and the results are presented with 10^7 particles.

As in KH15, the disc spans $R_{in} = 4$ to $R_{out} = 40$, with aspect ratio $H/R_{in} = 0.06$ and central density $\Sigma_{in} = 0.1504$ at the inner edge. To match the surface density profile after relaxation and the scale height, we set $p = -1$ and $q = 0$. As we model the physical viscosity using the method described in Lodato and Price (2010), this results in an α viscosity that varies slightly with R .

KH15 confirm their disc is in the bending wave regime by measuring both waves travelling at $0.5c_s$ and the viscosity parameter α directly from their simulation. We seek to enforce this in our simulations by setting α such that $\alpha \approx 0.12$ at $R = 5$ and $\alpha \approx 0.049$ at $R = 10$, very similar to the values of 0.1 and 0.05 measured at the same locations in KH15. We note that this comparison is not exact because in our hydrodynamic simulation we are setting the α values using the artificial viscosity included to capture shocks (e.g. Equation 2.68) with a forced dependence on the radius that comes from the q value used, while Krolik

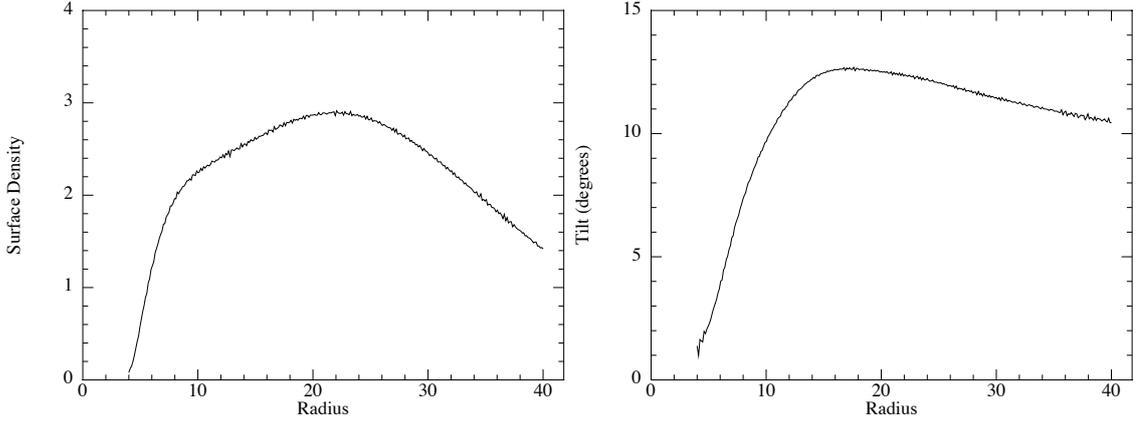


FIGURE 4.2: Surface density in units of $c^4/G^2 M$ (left) and tilt (right) after 25 orbits when apsidal precession is neglected.

and Hawley (2015) measure theirs directly from the MHD induced turbulence (and so no such dependence is forced). With this description, $H/R \lesssim \alpha$ when $R \gtrsim 7$, suggesting that most of the disc will be in the bending wave regime.

The spin parameter a is set from the requirement that the Lense-Thirring precession frequency is $1/15 \times$ the orbital frequency at $R = 10$ (KH15). As discussed in Sorathia, Krolik, and Hawley (2013), this constraint is made for numerical convenience as it can only be achieved with a non-physical value of $a \geq 1.05$ — a higher spin means precession affects occur faster for the same dynamical time when compared to a low spin case, so the simulation does not have to be run for as long. In these simulations, we chose the spin to be the maximum physical value of $a = 1.0$ (although in reality the spin is likely to saturate at a value somewhat lower than this). We note that this (minor) discrepancy appears in our results in that the black hole torque appears to affect the disc more slowly than observed in KH15.

4.1.2 Initial conditions

We arrange the particles in the disc using a Monte Carlo placement method, with the disc aligned to the black hole spin. The disc is then allowed to relax for 12.5 orbits at $R = 10$ in this plane (12.4 orbits were conducted in KH15). Figure 4.1 shows the surface density profile at this time, showing that most of the mass is in the outer regions. Our disc shows the same features as KH15 (see their Figure 1) but has a slightly lower mass. We use this disc as the initial condition for each of our subsequent simulations, to allow direct comparison with KH15. We produce the same surface density profile for the discs that include apsidal precession by repeating this process with the full effective GR potential.

The disc inclination $\beta(R)$ is the angle between the local angular momentum vector of the disc and the z axis (defined by the black hole spin). For each simulation, the particles are rotated through a constant inclination angle β , so the disc is tilted but not warped. The black hole torque is then applied using either the Keplerian or effective GR potential

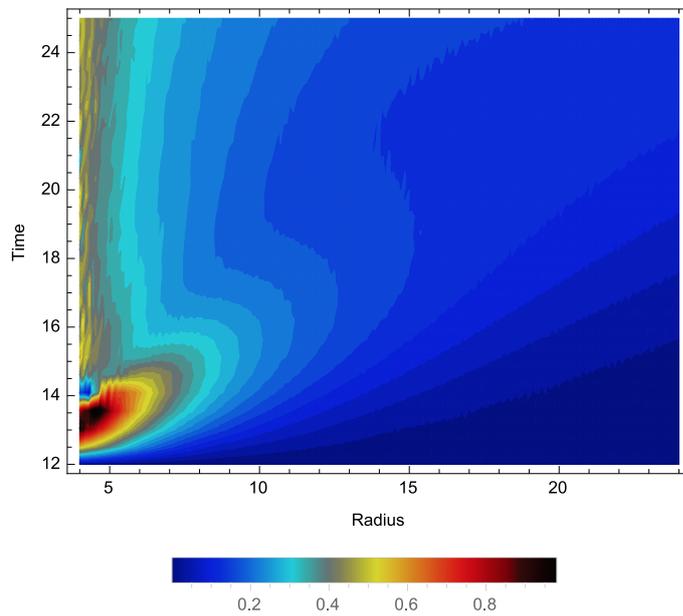


FIGURE 4.3: Twist, measured in radians/ π , as a function of time and radius ignoring apsidal precession. This figure, assuming the same precession physics, should be compared with Figure 2 of KH15.

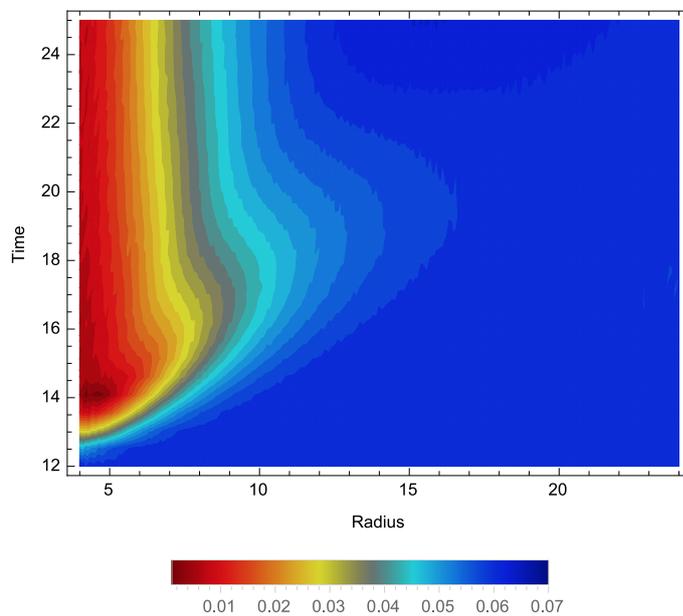


FIGURE 4.4: Tilt, measured in radians/ π , as a function of time and radius ignoring apsidal precession. This figure, assuming the same precession physics, should be compared with Figure 3 of KH15. Note KH15 define the tilt to be the negative of our definition.

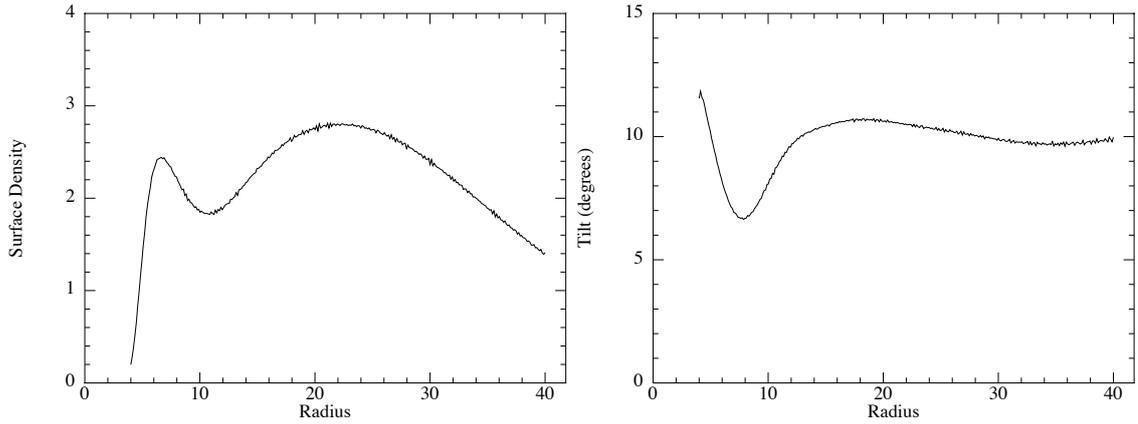


FIGURE 4.5: Surface density in units of $c^4/G^2 M$ (left) and tilt (right) after 25 orbits when apsidal precession is included.

as appropriate (for details see Chapter 3). We use units such that $G = M = 1$ and $R_g = GM/c^2 = 1$ and define our timescale with orbits at $R = 10$.

4.2 Results

We conduct two simulations with the only difference being the modelling of the precession frequencies (Nealon, Price, and Nixon, 2015). Both simulations implement Lense–Thirring precession, but the first simulation uses a Keplerian potential (as in KH15) while the second simulation includes the effects of apsidal precession by using the effective GR potential. The discs are inclined at $\beta = 12^\circ$ and run for an additional 12.5 orbits from the initial condition described in § 4.1.2, as described in KH15.

4.2.1 Apsidal precession neglected

The right panel of Figure 4.2 shows the tilt as a function of radius for the simulation with a Keplerian potential, i.e. where apsidal precession is neglected. The start of the simulation is marked by a sharp disturbance in the tilt and surface density profiles as the black hole torque is applied, but this wave damps within 4.5 orbits and the rest of the disc evolution is gradual (this initial disturbance occurs in all of the simulations reported). Our disc evolves in the same manner described by KH15; the inner edge aligns with the spin of the black hole, and there is a smooth transition to the outer region which remains misaligned. The left of Figure 4.2 shows the surface density profile of our disc. Similar to KH15, there is little evolution of the surface density profile when apsidal precession is neglected. To allow direct comparison, we also show figures plotted in the same way as in KH15 such that Figures 4.3 and 4.4 are directly comparable to Figures 2 and 3 of their MHD simulation. The maximum values of the tilt and twist measured from our simulation are within 15% of theirs and the evolution of our disc is in qualitative agreement. The comparison between

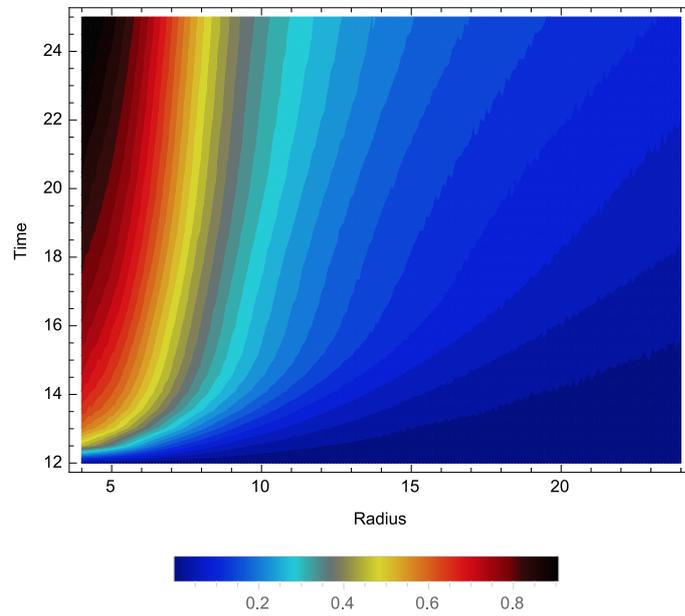


FIGURE 4.6: Twist, measured in radians/ π , as a function of time and radius in a disc including apsidal precession.

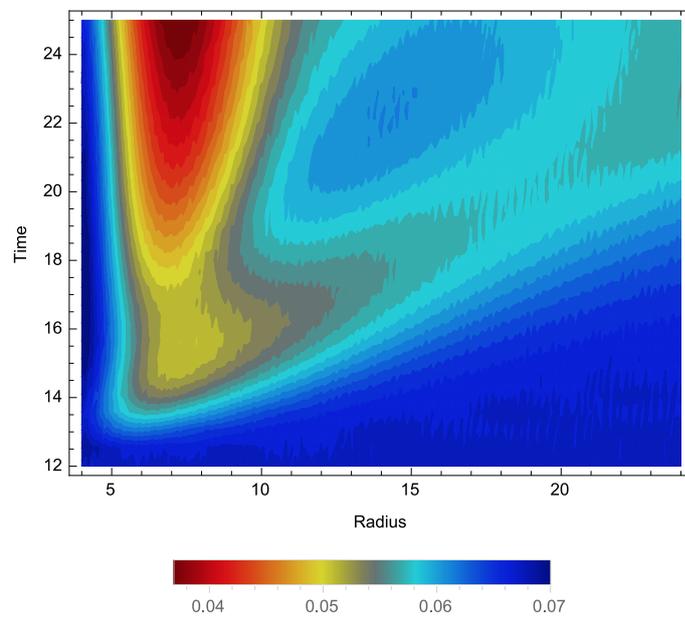


FIGURE 4.7: Tilt, measured in radians/ π , as a function of time and radius in a disc including apsidal precession.

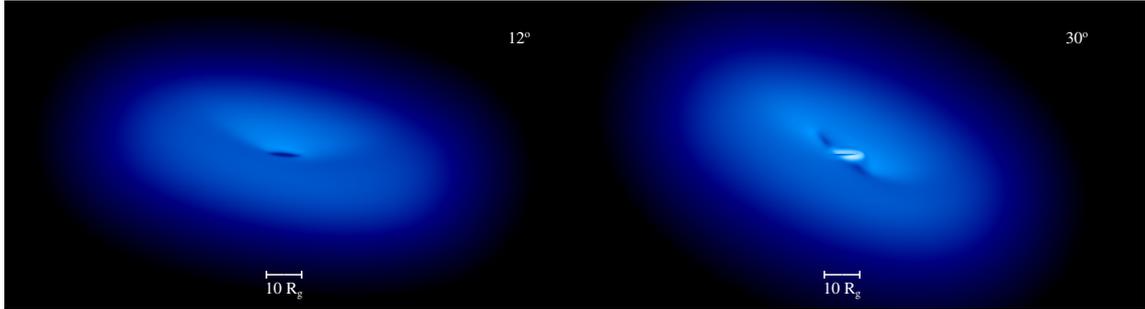


FIGURE 4.8: Density rendering of warped accretion discs inclined at 12° (left) and 30° (right) to the black hole spin (spin axis is vertical in the plane of the page, i.e. along the z axis). Evidence of disc breaking is seen only in the high inclination case.

our α -disc simulations neglecting apsidal precession and the magnetohydrodynamical simulation of KH15 confirms that an α -disc simulation gives similar results to an MHD one. That is, in at least this case MHD discs and pure α -discs show equivalent results (Sorathia, Krolik, and Hawley 2013 found a similar result).

4.2.2 Apsidal precession included

Figures 4.5, 4.6 and 4.7 show our results for discs with the same initial setup as § 4.2.1, but this time with the full effective GR potential, so that apsidal precession is explicitly present. Here, in contrast to simulations which neglect it, we see that the inner edge of the disc remains misaligned. The tilt of the disc then decreases to a minimum around $R \sim 9.5$ before increasing with radius. The outer edge of this disc differs from that of the previous case (also seen in the tilt comparison), suggesting that a larger disc is required to prevent the outer boundary affecting the disc evolution.

The marked differences in the simulation results with and without apsidal precession show that including apsidal precession is crucial to give even qualitatively correct results.

4.2.3 Breaking discs

Nixon, King, and Price (2012) showed that one can estimate the radius at which a disc may break by equating the Lense–Thirring precession torque with the internal viscous torque. This initial investigation neglected the additional viscous torque arising from a warp. Using this simplified criterion for $\beta = 12^\circ$ would give $R_{\text{break}} \lesssim 20$, suggesting that the disc should have broken, contrary to the numerical results of both ourselves and KH15. But this conclusion is wrong: neither of these simulations is in the diffusive regime where this criterion would apply. The relevant criterion for the wavelike regime involves

a comparison of the sound crossing and precession timescales (Nealon, Price, and Nixon, 2015) and shows that the disc should not break, as found in the simulations.¹

We address the question of whether the disc can break by simulating a disc with the same initial mass distribution, but this time with an inclination of 30° , again for 12.5 orbits and using the effective GR potential. As the disc evolves the inner edge tilt decreases from the initial 30° but does not fully align with the black hole. The tilt then increases sharply around $R \sim 10.5$, suggesting a break in the disc, before decreasing gradually at larger radii. Figure 4.8 shows a density rendering of the 30° disc at the end of the simulation, compared to the warped 12° disc. For the high inclination disc, the separated innermost disc is only slightly misaligned, whilst the inner edge of the outer disc is inclined at $\sim 30^\circ$.

4.3 Discussion

The striking differences that occur between simulations of warped wavelike discs that include or ignore the effects of apsidal precession confirm that it has a strong impact on the disc evolution. The inner edge of the disc is altered completely (from aligned to misaligned) and the oscillatory behaviour observed is consistent with previous analytical (Ivanov and Illarionov, 1997; Lubow, Ogilvie, and Pringle, 2002) and numerical (Nealon, Price, and Nixon, 2015) studies. We therefore conclude that GR effects like this must be modelled accurately in simulations. This is achievable by either an appropriate post-Newtonian approximation, as used here (Nealon, Price, and Nixon, 2015), or a GR treatment (e.g. Fragile et al., 2007; Morales Teixeira et al., 2014).

The comparison between our hydrodynamic simulation and the MHD simulation of KH15 confirms that the disc dynamics can indeed be captured by hydrodynamics with an α viscosity. To date, Morales Teixeira et al. (2014) have completed the thinnest disc simulation with both GR and MHD taken into account, with an aspect ratio of $H/R \sim 0.08$. However, this simulation was still in the bending wave regime as the viscosity parameter measured from the simulation was $\alpha \sim 0.01$. In contrast, SPH simulations have already studied the diffusive regime, with $H/R \simeq 0.01$ and $\alpha \simeq 0.1$ (Nixon et al., 2012). GRMHD simulations to date have been unable to simulate discs which would have been expected to undergo disc breaking or tearing. For the disc parameters used in the simulations presented in this paper, we show that disc breaking is possible when the disc is inclined at $\beta \gtrsim 30^\circ$.

¹Even in the diffusive regime, Doğan et al. (2015) showed that at small inclinations or when α is small, it is unreasonable to neglect vertical viscosity when estimating the internal torque in the disc. As these discs have low inclination and $\alpha \sim 0.03$, it would be crucial to include the vertical torque in the breaking criterion if this simulation was diffusive (Doğan et al., 2015).

4.4 Conclusion

We have conducted hydrodynamical simulations of accretion discs in the bending-wave regime with an α viscosity. By comparing with the MHD simulation of KH15, we confirmed that hydrodynamical simulations using an α viscosity capture the dominant evolution of warped accretion discs, showing results that are remarkably similar to the MHD simulations. We have shown that modelling the apsidal precession in the disc strongly affects its evolution. Simulations that do not take apsidal precession into account cannot give the correct disc evolution, which has nonzero disc tilt at the inner edge, and stable tilt oscillations with radius in the central disc regions. Finally, as expected, for the disc parameters chosen by KH15 we find no breaking, but demonstrate that a disc with the same parameters but a larger inclination does break.

Chapter 5

Quasi-periodic oscillations

Despite being first observed more than 45 years ago, the complete physical processes that drive quasi-periodic oscillations (QPOs) is still unknown. These rapid oscillations in the flux observed from low mass X-ray binaries (LMXBs), active galactic nuclei and even cataclysmic variables are revealed in power density spectra (PDS). Observations have identified that there are several sub-types of QPO, distinguishable by different timing properties and each associated with a different disc state.

These observations lead to two important conclusions. First, QPOs must be sourced from the innermost part of the accretion disc where general relativistic effects are strongest. Second, the nature of QPOs (simultaneously observable frequencies that are distinct from one another) suggests that the disc is not a smooth, continuous structure as suggested by current accretion disc theory. Although numerous models exist to explain QPOs, we propose that disc tearing provides a natural mechanism to generate the inhomogeneous disc that is implied by observations. The major strength of this model is that it only requires that we relax the assumption that the disc is aligned to the black hole spin. The simulations in previous chapters have identified precessing structures generated at tens of R_g , as proposed from observations of QPOs. We thus seek to understand the potential connection between QPOs and disc tearing.

This chapter focuses on conducting a simulation that demonstrates tearing and generates a PDS that is comparable to observations. Practically, this requires that the simulations resolve the scale height of the disc, have a radius larger than the breaking radius and that the simulation be conducted for as long as possible. We thus choose parameters for the simulations that attempt to satisfy the above requirements and in order to compare to observations, physical parameters related to the LMXB XTE J1550-564.

5.1 Connecting tearing discs and QPOs

The main goal of this chapter is to consider the connection between disc tearing and LFQPOs. This has challenges in both the numerical modelling of the disc and how the simulation is related back to observations, so we break the chapter into two main sections. The first concerns a physically motivated simulation that demonstrates disc tearing. Once we have simulated a disc that meets this criteria, we explore how we may measure a PDS from this simulation. Finally, we consider the accuracy of our simulations and whether we can make any comment on the different sub-types of LFQPOs. We split this process into the following set of questions:

1. Can a disc with the parameters of a LMXB (i.e. spin, black hole mass and viscosity) undergo disc tearing?
2. Our previous simulations have not demonstrated sustained tearing. Can this simulation show that tearing is a repeatable (or alternatively, one-off) event?
3. Using the dynamics of the gas, do the structures formed precess with a frequency consistent with observed LFQPOs?
4. As PHANTOM currently does not generate something like this directly, how can we measure a PDS from our simulations?
5. Can we simulate the disc for long enough to have resolved tearing behaviour *and* enough precessions to measure a PDS with any significance?
6. How do the results of our PDS compare to observations of QPOs?
7. How do the approximations we have made affect our findings (e.g. our use of a post-Newtonian approximation)?
8. Although we focus on the most commonly observed type-C QPOs, is there any scope to relate tearing rings to type-A or -B QPOs?

As each simulation informs the subsequent simulations, this chapter is necessarily laid out in chronological order. In § 5.2 - 5.5 we begin by quantifying how accurate our post-Newtonian approximation is for calculating the nodal, apsidal and orbital frequencies (in comparison to the Kerr metric), describing the approximate method we use to calculate the luminosity in the disc, how we can use this to generate a PDS and the initial conditions that are common to each simulation. In § 5.6 - 5.7 we describe the results from each set of simulations in turn that use these methods. In § 5.8 we introduce a method to measure the precession of the gas and a new method to generate a PDS, both using the existing simulation data. Finally, in § 5.9 we discuss the high resolution simulations that are continuing during the submission of this work and in § 5.11 our conclusions.

5.2 Precession frequencies

The Stella and Vietri (1998) model associates the three fundamental frequencies experienced by an eccentric, tilted orbit with the three observed types of QPO. Here, the fastest high-frequency QPO (HFQPO) is generated by the orbital frequency near the innermost stable orbit. In the limit for a test particle in the Kerr metric, this is given by (e.g. Kato, 1990)

$$\nu_\phi = \frac{1}{2\pi} \frac{GM}{R^3} \frac{1}{1 + a(R_g/R)^{3/2}}. \quad (5.1)$$

The lower HFQPO is associated with the difference between the radial oscillation and orbital frequencies, given at the same radius as

$$\nu_r = \nu_\phi \left(1 - \sqrt{1 - \frac{6R_g}{R} + 8a \left(\frac{R_g}{R} \right)^{3/2} - 3a^2 \left(\frac{R_g}{R} \right)^2} \right). \quad (5.2)$$

Finally, the LFQPO is represented by the Lense-Thirring precession frequency (i.e. the nodal precession frequency), given by the difference in the orbital and vertical oscillation frequencies to be

$$\nu_{LT} = \nu_\phi \left(1 - \sqrt{1 - 4a \left(\frac{R_g}{R} \right)^{3/2} + 3a^2 \left(\frac{R_g}{R} \right)^2} \right). \quad (5.3)$$

Because we use a post-Newtonian approximation, these frequencies in our simulations are not reproduced exactly. Instead, these frequencies are determined by our choice of potential. From Chapter 3, using the Einstein potential the equivalent expressions are

$$\nu_{\phi,E} = \frac{1}{2\pi} \sqrt{\frac{GM}{R^3} \left(1 + \frac{6R_g}{R} \right) - \frac{2aR_g^2 c \sqrt{GM}}{R^{9/2}}}, \quad (5.4)$$

$$\nu_{r,E} = \nu_{\phi,E} \left(1 - \frac{1}{2\pi} \sqrt{\nu_{\phi,E}^2 - \frac{6GMR_g}{R^4} + \frac{3aR_g^2 c \sqrt{GM}}{R^{9/2}}} \right), \quad (5.5)$$

$$\nu_{LT,E} = \nu_{\phi,E} \left(1 - \frac{1}{2\pi} \sqrt{\nu_{\phi,E}^2 - \frac{4aR_g^2 c \sqrt{GM}}{R^{9/2}}} \right). \quad (5.6)$$

Figure 5.1 shows a comparison of these three frequencies in both the free particle approximation and from the Einstein potential that is used in subsequent simulations. When applying the relativistic precession model (RPM) it is assumed that the emission for all three QPOs is sourced from the same radius. With the set of three equations then, the mass (encapsulated in R_g) and the spin of the black hole can be measured precisely if three

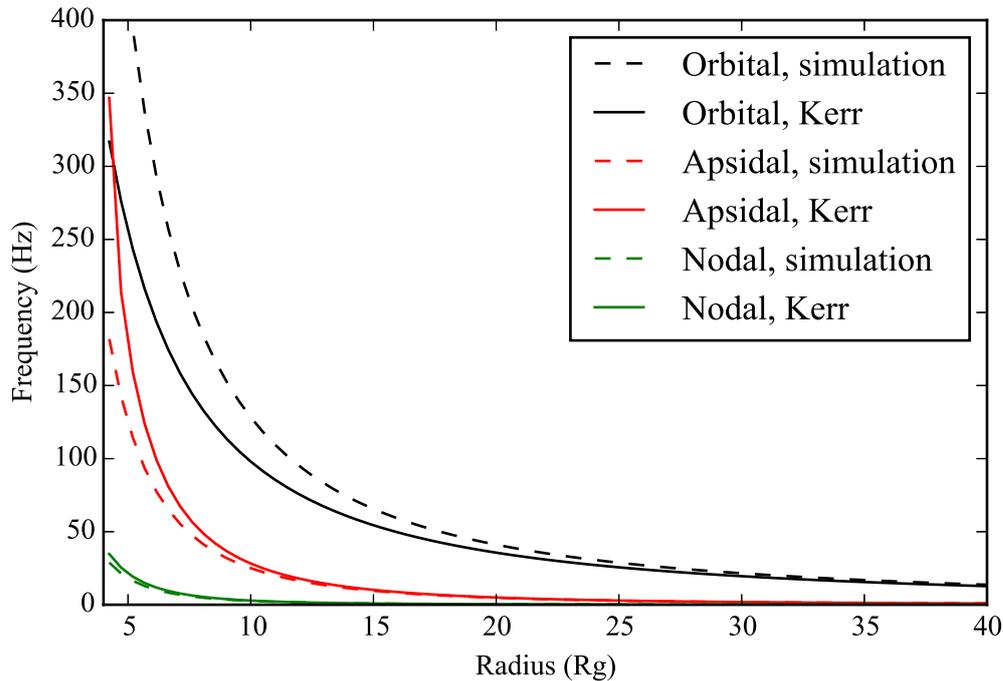


FIGURE 5.1: Comparison of the orbital, apsidal and nodal precession frequencies in the limit of a test particle in the Kerr metric and for the post-Newtonian potential used in PHANTOM for a $10M_{\odot}$ black hole with $a = 0.5$. Outside of $\sim 10R_g$ the potential we use compares well to expected frequencies, but any frequencies measured from inside of this should be treated carefully.

simultaneous QPOs are observed. This process is outlined in Motta et al. (2014a), with precision up to $a \pm 0.01$, but can also be completed when only two simultaneous QPOs are observed by using observational estimates of the mass (Motta et al., 2014b). We note that the precession frequencies for a globally precessing torus differ from these (except in the limit when the ring is extremely narrow), such that any observations that fit the above set of frequencies cannot be sourced from an extended torus. In this case, the precession frequency is dictated by the weighted average of the precession frequencies of the gas in the torus. Recently, Franchini, Motta, and Lodato (2016) argued that the rigid precession model was equivalent to the test particle approximation at the innermost stable orbit.

5.3 Luminosity in the disc

Building on simulations in previous chapters, here the requirement is to generate a power density spectrum (PDS) from a light curve. Observationally these are calculated by measuring the net flux from the system as a light curve and then taking a Fourier transform. The frequency of the oscillations (if roughly constant) is then revealed in local peaks of the PDS. To mimic this from our simulations requires two independent steps; the material in the simulation needs to be luminous (generating the flux), followed by ray-tracing taking

into account obscuring material to measure the flux that made it to the observer. Currently, PHANTOM does not have either of these features, and their full implementation is not in the scope of this work. Instead, we consider how these may be approximated using existing features in Phantom; the first consideration is the emission from the disc.

We approximate the luminosity emitted by the disc with the energy that is liberated during the heating of shocks and through the compression of gas - that is, the sum of the shock energy and compressive work. This approximation represents the energy that is discarded when the disc is modelled isothermally, so is equivalent to making the disc luminous rather than allowing the gas to heat up. A major motivation for this implementation is that it is straight-forward because the framework to calculate these values already exists in PHANTOM. In SPH form, this energy is represented by the the terms on the right hand side in Equation 2.48.

This energy is then summed across all the particles in each time-step (only including positive contributions from the compressive work), giving the net luminosity as a function of time. Summing the energy this way inherently assumes that the disc is optically thin so any bright part of the disc (whether ‘visible’ along an observers line of sight or not) is represented in the light curve. This method additionally assumes that cooling is instant, as no delay is implemented between heating of the gas and the luminosity increasing. The validity of this opacity assumption is investigated later in § 5.8.2. Figure 5.2 shows the emitted luminosity using this method (right panel) alongside the density in the same simulation (left panel). The strength of this approximation is that it efficiently traces out the location of the precessing rings (as they are always experiencing shocks due to their interactions with the neighbouring rings) and the developing innermost disc without much contribution from the outer, undisturbed disc. As these are the features that we want to focus on, the method appears justified.

The simulations in this chapter are assumed to be isothermal (as opposed to *locally* isothermal as in the previous chapters). Simulations completed by Nixon, King, and Price (2012) tested the effect of this assumption on the structures that form during tearing by conducting the same simulation with an isothermal and an adiabatic equation of state (their figures 7 and 8). In the case of the adiabatic simulation the rings that form are less well defined and more gas is present between the rings in a tenuous corona. For typical parameters in LMXB systems we anticipate behaviour somewhere between these two approximations, and thus more gas to form a corona-like structure in the tearing region than we demonstrate here.

5.4 Generating a power density spectrum

The light curve data is output at each time-step as calculated by PHANTOM. Because PHANTOM adjusts the time-step as it runs to be the longest possible, these are not necessarily equally spaced. As a result, we use the Lomb-Scargle method (Lomb, 1976; Scargle,

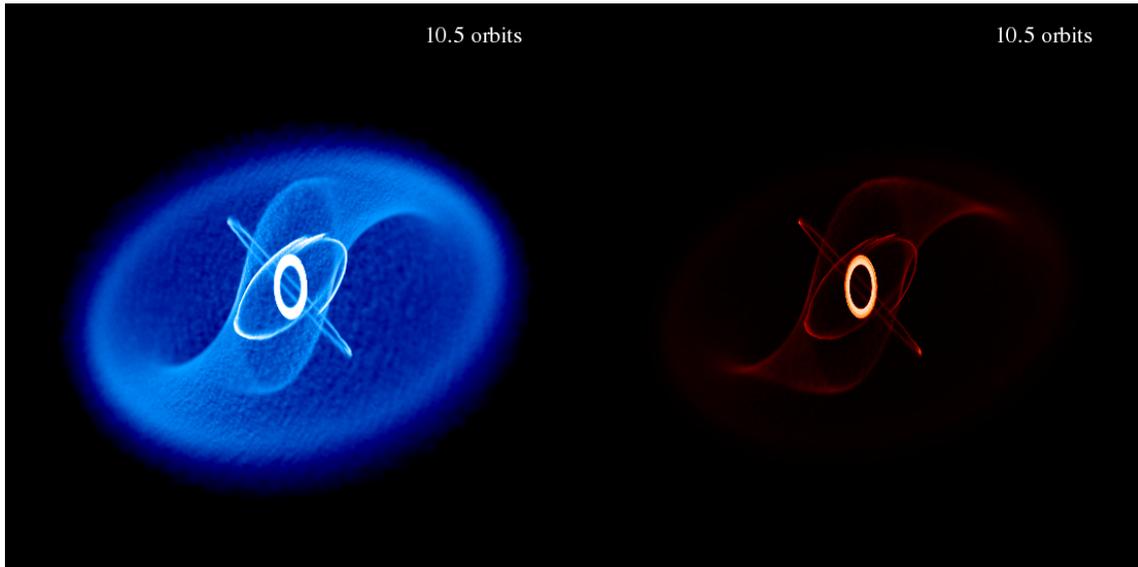


FIGURE 5.2: Column density (in blue) and luminosity approximation (in red) in a test simulation, where we assume that the dissipated energy is radiated instantly in an optically thin approximation. The strength of this model is that it traces the location and interaction of the rings but ignores the outer disc.

1982) to calculate the PDS, as this method is designed for unevenly spaced data. The specific code used comes from Press et al. (1993) and a test is documented in Appendix B. Here the luminosity as a function of time is used as the input, and so the power that is output represents the frequencies of any oscillatory behaviour in the light curve. As tearing disc behaviour is also associated with the mass accretion rate, we also consider PDS generated from this metric calculated in the same way as the light curve.

In the case that a feature is measured in the PDS, we need to quantify its significance to identify whether it is likely to be physical or not. Observational work quantifies features by fitting Lorentzian (or Gaussian, as required) profiles to features in the PDS and measuring the significance of these features (e.g. see Belloni, Psaltis, and van der Klis, 2002) — often achieved using the publicly available XSPEC package. This process assumes that the data has been retrieved observationally but of course, this is not true for our data. For a first measure we instead conservatively choose to calculate the significance using the null hypothesis straight from the generated PDS. This significance may be easily calculated from the PDS because the Lomb-Scargle method is already normalised, such that the false-alarm probability that any given peak with a particular power z is given by

$$P(> z) = 1 - (1 - e^{-z})^M, \quad (5.7)$$

where M is the number of independently sampled frequencies. This relation provides an estimate for the significance level of a peak z in the PDs. We note this method applied to this problem will only provide an indication of whether features are significant or not and that a proper treatment, using statistical methods (e.g. Bayesian statistics), is ultimately required.

No.	a	α	β°	H/R	R_{in}	R_{out}	p, q	t_f (sec)
S1	0.1 - 0.9	0.03 - 0.3	30, 60	0.02, 0.05	ISCO	30	1.5, 0.75	0.5
S2	0.1 - 0.9	0.03 - 0.3	30, 60	0.02, 0.05	30	100	1.5, 0.75	0.5
S3	0.1 - 0.9	0.03 - 0.3	30, 60	0.02, 0.05	ISCO	30	0, 0.75	0.5
S4	0.5	0.1	60	0.01	4.233	100	0, 0.75	4.0
S5	0.5	0.1	60	0.01	4.233	100	$p, 0$	20.0
S6	0.5	0.1	60	0.01	4.233	100	$p, 0$	25.0
S7	0.5	0.1	60	0.01	4.233	100	$p, 0$?
S8	0.5	0.1	60	0.01	6.0	100	$p, 0$?

TABLE 5.1: Summary of parameters used in the simulations described in § 5.6 - 5.9. Where multiple values are listed where we have performed a parameter sweep. In simulations 4-7, the p value is determined by the solution to the steady state solution described in § 5.7.1. Simulations 7 and 8 are not yet completed.

5.5 Initial conditions

The simulations presented here have three competing requirements: the resolution, size of the disc and length of the simulation. From Chapter 3, resolving at least the scale height of the disc was demonstrated to be crucial to recover tearing behaviour. However, the outer edge of the disc needs to be sufficiently larger than the breaking radius such that material being torn does not pull the outer edge of the disc with it (and precesses as a solid body). These two requirements directly conflict each other: for a given number of particles, increasing the size of the disc decreases the resolution and we are limited by the number of particles. Finally, in order to measure the frequency in a PDS from a structure in the simulation it must complete as many full precessions as possible. Indeed, PDS derived from observations make use of minutes if not hours of data which is computationally prohibitive. This final requirement competes with the first two; running the simulation for a long time means we need a large outer edge (to avoid any problems from the outer boundary) but having a large disc that has enough particles to be resolved is difficult. Finding a compromise between these three requirements while generating the tearing behaviour to be linked with QPOs motivates the parameters of the simulations in this chapter.

In order to compare to observations, we chose parameters appropriate to the LMXB J1550-564. This object displays both HF and LFQPOs and has been observed numerous times (see Figure 1.14 and Table 1.1 for a summary). We thus chose a nominal spin of $a = 0.5$ and a black hole mass of $10M_\odot$ (Steiner et al., 2011). The outer radius of the disc ideally would extend to the same order as the binary separation ($\sim 10^5 - 10^6 R_g$), but this is far too large to simulate with current resources. Instead, we consider material inside $R_{\text{out}} = 100R_g$ and ensure that the duration of our simulation is not compromised by effects from this boundary. As discussed in Chapter 3, we do not anticipate that this will change the tearing behaviour dramatically.

The subsequent (more than 200) simulations use different combinations of surface density profiles (controlled by p, q), aspect ratios (H/R) and inclinations (β). These parameters are summarised in Table 5.1 with simulations identified using S1-S8 throughout the

text. Associated light curves and PDS not presented here are included in Appendix B for completeness.

5.6 Initial parameter sweeps

The initial set of simulations were ambitiously designed to examine the effect of changing spin, disc thickness, viscosity and inclination. In an attempt to make the simulations computationally efficient, the disc was split by radius into an inner and outer section. The inner section went from the innermost stable orbit for the given spin to an outer radius of $30R_g$, and the outer disc continued from here through to an outermost radius of $100R_g$. For a given number of particles this method allowed much greater resolution than simulating the disc as a whole and (particularly the outer disc simulations) were comparably faster.

Inner disc simulations (S1)

An example of the inner disc simulations is shown in Figure 5.3. In most of these simulations, strong warping of the disc leads to the formation of an almost aligned, high density ring. The panels of Figure 5.3 visually demonstrates that this ring experiences apsidal precession and this structure is quite luminous (recall Figure 5.2). Despite experiencing strong warping, these discs did not display any tearing. The breaking radius for these discs with $a = 0.5$ and $\beta = 60^\circ$ is $\approx 40R_g$, larger than the outer radius. The lack of tearing confirms previous findings that the location of the outer edge relative to the breaking radius is critical to modelling tearing correctly, suggesting that global disc simulations are required. The PDS generated from the mass accretion rate and the light curves of these simulations did not show any significant features (Figure B.3 displays a representative sample).

Outer disc simulations (S2)

To make the outer disc appear to be an extension of the inner disc, the surface density profile was constructed to be continuous with the inner disc (shown in Figure 5.4) by changing the definition of the radius where the sound speed and surface density profile initial values were chosen in the code. In all of these simulations, the inner edge of the disc warped but did not experience any tearing. Additionally, any consistency between the inner and outer disc simulations was lost as the disc evolved: the surface density profiles, tilt and twist at the shared boundary did not evolve similarly as anticipated. As in the matching inner disc simulations, the PDS generated did not show significant features (Figure B.4).

A shared feature of the inner and outer disc simulations was a steep decrease in the luminosity of the disc at the start of each simulation (also seen in Figure 5.6 in the first 0.1 seconds). This sharp drop in luminosity represents material being accreted as the initial conditions relax because the surface density profiles chosen do not represent that of a steadily accreting disc. In these simulations, the profile of the surface density and sound

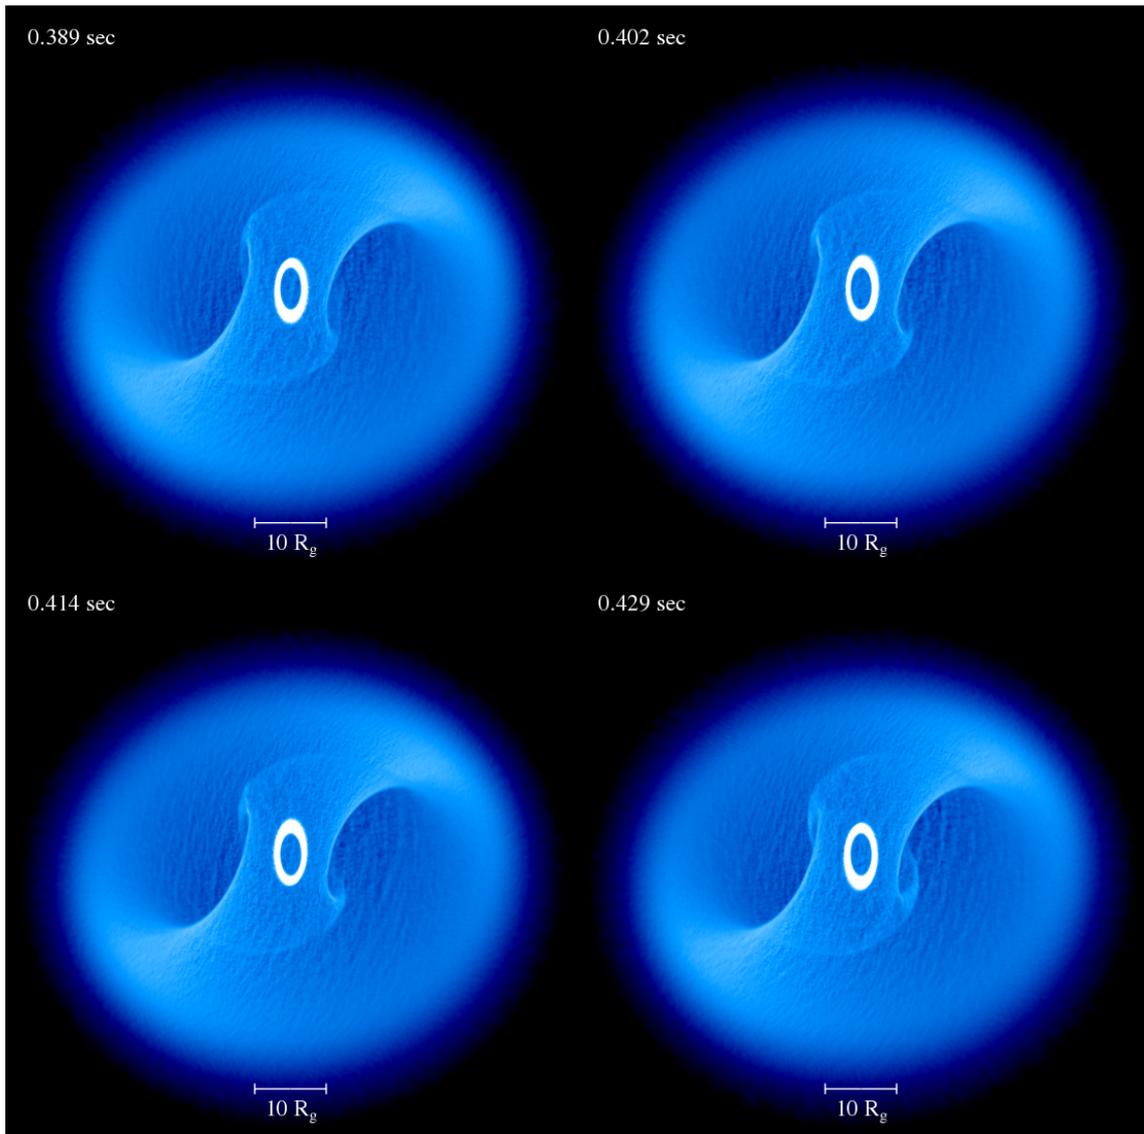


FIGURE 5.3: Apsidal precession identified in our inner disc simulations, where this simulation has $a = 0.9$, $\alpha = 0.3$, $H/R = 0.02$ and $\beta = 60^\circ$ and the view is looking down the spin vector of the black hole. Capturing this motion suggests that apsidal precession may also be important here but that high local resolution is required.

speed were chosen to give a constant α throughout the disc, but this is not the same as modelling a steadily accreting disc and thus caused the initial decrease in the light curve. In other words, our numerically motivated choice for the surface density profile was responsible for the sharp drop in luminosity and it is not a physical effect. In our future simulations we instead consider a surface density profile that corresponds to a constant mass accretion rate, which should not suffer from this feature.

Second parameter sweep (S3)

After the previous two parameter sweeps were completed, a coding error in the implementation of the luminosity sum was discovered. As the outer disc simulations did not show

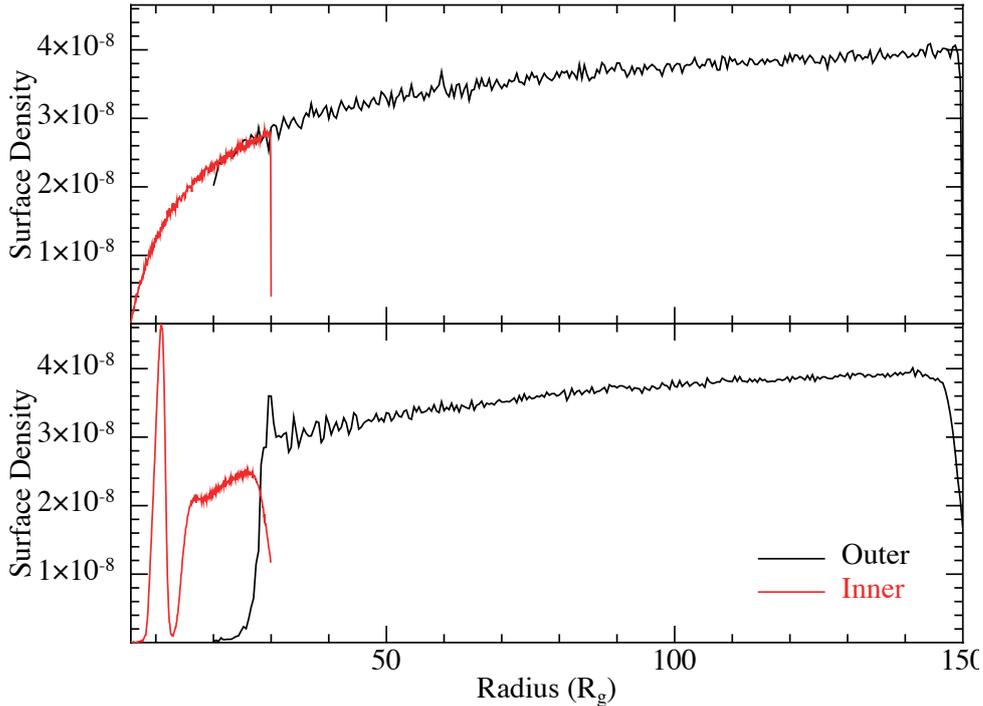


FIGURE 5.4: The evolution of the surface density between our matched inner and outer disc simulations (§ 5.6) confirms that global, radially extended disc simulations must be used to investigate disc tearing. The upper panel represents the initial condition and the lower panel the end of the simulation, where the discs are no longer consistent.

any interesting physical behaviour dynamically (i.e. the inner edge warped slightly but did not agree with the the behaviour of the outer parts of the inner disc simulations, see Figure 5.4), only the inner disc simulations were repeated. In an effort to reduce the initial jump in the light curve, the profile of the surface density was altered to one suitable for a constant mass accretion rate, \dot{M} . By recalling that $\dot{M} = \nu\Sigma$, but $\nu = \alpha c_s H$, $\alpha \propto \langle h \rangle / H$, and $c_s \propto R^{-q}$,

$$\dot{M} \propto \Sigma \langle h \rangle c_s \propto R^{-p} R^{(p-q)/3+1/2} R^{-q} = R^{-2/3p-4/3q+1/2}. \quad (5.8)$$

This constrains $p = 3 - 2q$ for a constant mass accretion rate. Using tests of aligned discs, for these simulations we found that any reasonable choice that matched this criteria effectively eliminated the initial step decrease in the light curve. The combination of $p = 0$ and $q = 0.75$ provided the slowest mass accretion rate and so was chosen for these simulations. These simulations showed the same behaviour as the previous inner disc simulations (1), including the apsidal precession at high spins. Despite the corrected luminosity implementation, the PDS generated did not show any obvious features (Figure B.5). Confirming further the need for large disc simulations, these are considered next.

5.7 Large disc simulations

Rather than conduct a parameter sweep of large disc simulations, we instead focus on a single simulation with associated features in the PDS before expanding the parameter space. For these simulations, using $R_{\text{out}} = 100R_g$ was found to have both a resolved disc using 10^6 particles and be large enough to be able to simulate for a few seconds of physical time. Decreasing this outer radius did not increase the resolution significantly but increasing it restricted the final run time considerably.

First large disc simulation (S4)

Figure 5.5 shows the column density evolution of the first large disc simulation, Figure 5.6 the luminosity and Figure 5.7 the associated PDS. A warp propagates outwards from the inner region and a small ring is immediately formed, less than $10R_g$ wide with an inclination of $\approx 30^\circ$. This ring is repeatedly destroyed and reformed as the warp propagates to the outer edge. At 1.5 seconds, the inner ring has stabilised and is aligned to the black hole spin. Narrow rings are subsequently ripped off the outer disc, precess briefly and then collapse onto the inner ring. At 2 seconds a large ring rips off and completes a full precession before aligning and increasing the inner disc. The final 0.5 seconds of the simulation show the inner ring to be steady and no indication of further tearing events. The evolution of this disc demonstrates tearing in an ‘inward-outwards’ sense — a ring is ripped off the the outer disc, precesses independently, accretes and this process repeats at a larger radius.

The steep drop in the initial part of the light curve in Figure 5.6 was previously removed by choosing a more physically motivated surface density profile. In this case, the simulation was mistakenly initialised with $p = 1.5$, $q = 0.75$ and so the steep decrease returned. The prominent peaks in the early part of the light curve occur with a separation of ≈ 0.3 seconds. Each of these peaks corresponds to the increase in shocks when a ring is being torn off, suggesting that these rings are ripped off the outer disc at a particular rate (but not at the same radius each time, so the rings do not have the same precession frequency). However, after the initial tearing (till about 1.5 seconds) the behaviour in the light curve changes — material is being ripped from the outer disc at larger radii, and so the frequency with which they are ripped decreases. This change in tearing behaviour raises an important question about the initial condition: these simulations have been initialised with material inside the breaking radius that *has not yet torn*. Whilst this may occur in a physical system if heating and cooling are taken into account, our simulations do not. Rather, we would expect material to start misaligned at a large radius and be torn as it moves inwards through the breaking radius.

The PDS shown in Figure 5.7 identifies an unexpected HF feature at 500 Hz and an associated harmonic at 250 Hz. Assuming that this is caused by free particle motion, we can estimate that this frequency must be generated by orbital material at a radius of $\approx 5R_g$. This measurement agrees well with theory, and so we tentatively identify a HF

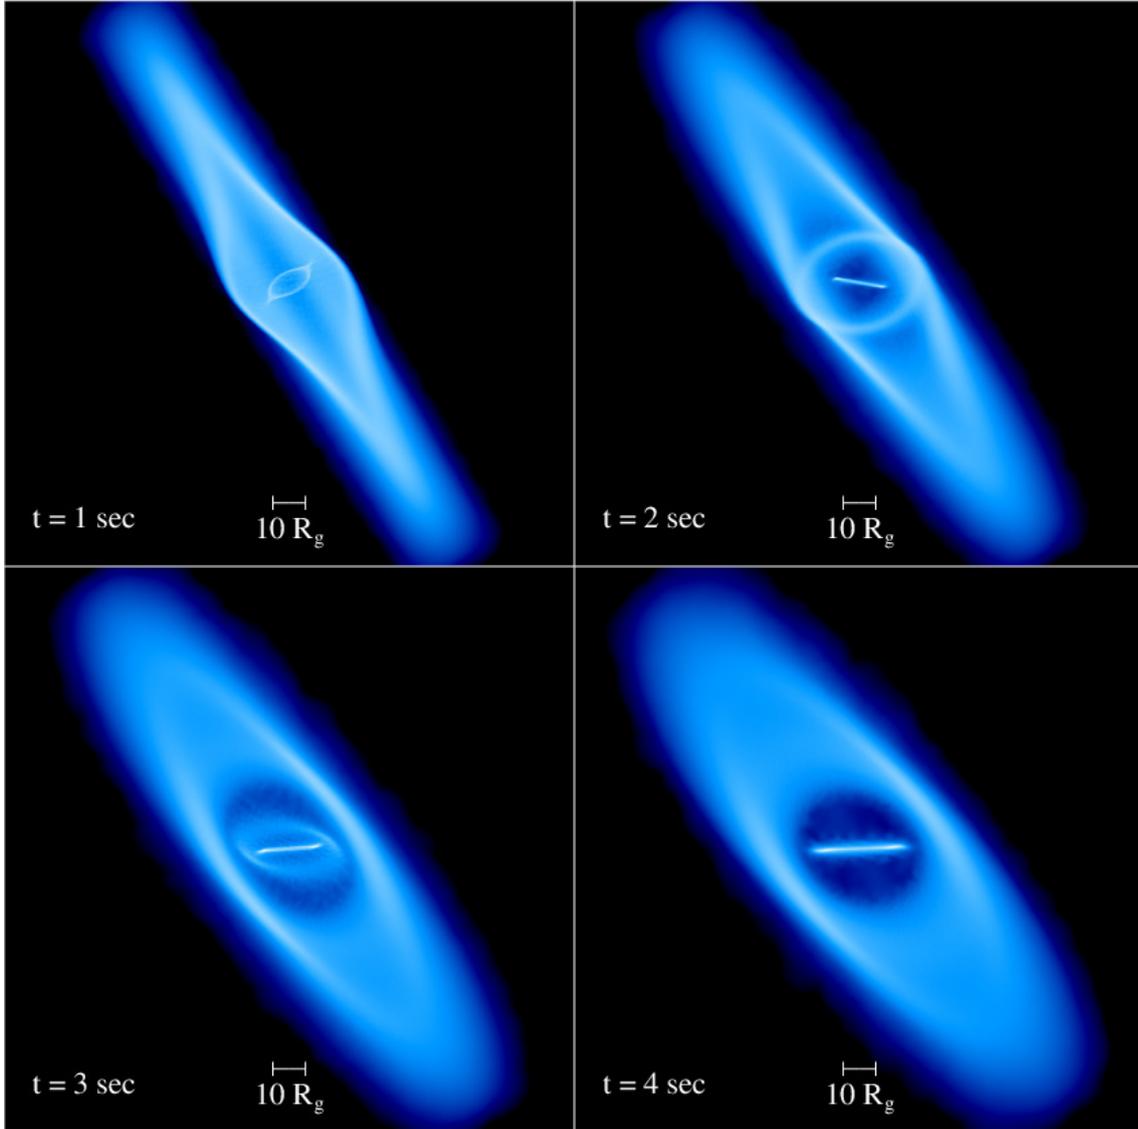


FIGURE 5.5: Column density rendering of the first large disc simulation (S4), with the black hole spin vertical on the figure. As in our previous tearing simulations, tearing occurs from the inside edge of the disc to larger and larger radii (in an inside-out fashion). We conclude that this behaviour is an artefact of our initial conditions and thus use different initial conditions for subsequent simulations. Times are scaled to a black hole mass of $10M_{\odot}$.

QPO from these simulations. The validity of the identification of this feature is discussed in later sections.

This simulation highlights the need for longer duration simulations. Although we would expect HF features to have a higher power than LF features in a given amount of time, it was not necessarily anticipated that HF features would be demonstrated in these simulations. The focus here however is on LF features, so for future simulations we consider how we can increase the duration of the simulations to ensure that the LF features are repeated enough periods to be represented in the PDS.

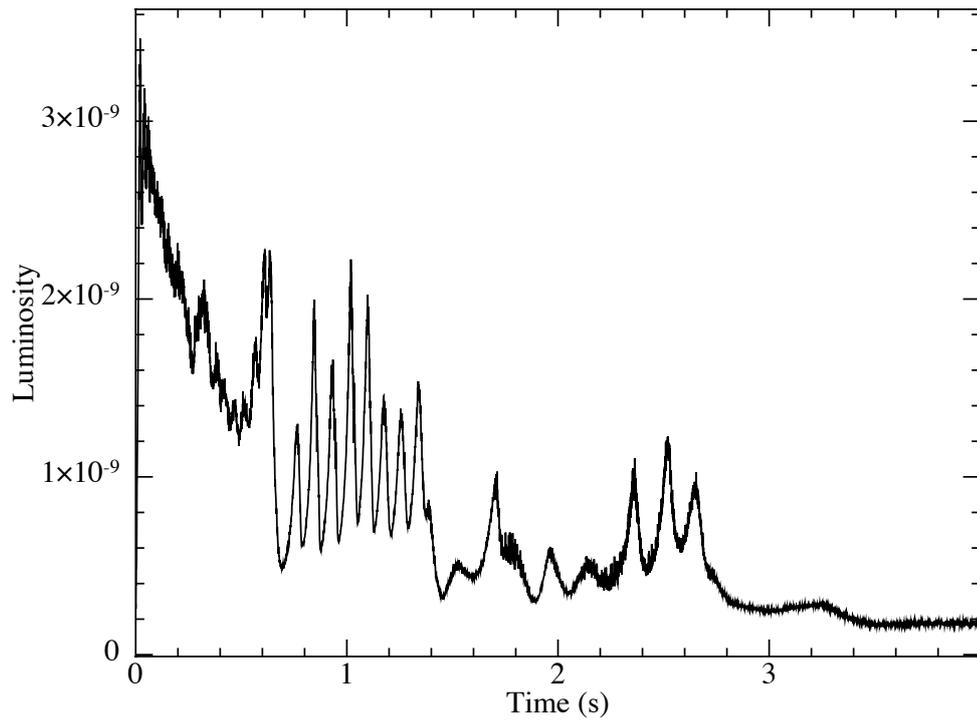


FIGURE 5.6: Simulated light curve derived from Figure 5.5 (S4). Here each peak is associated with a ring tearing in the inner region (e.g. first panel, Figure 5.5).

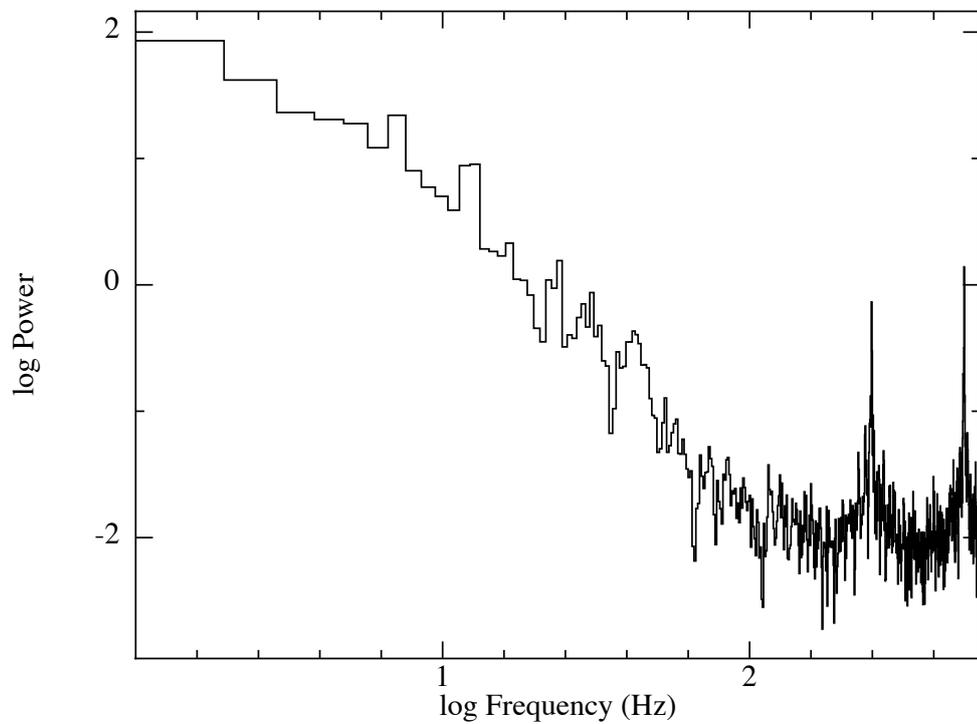


FIGURE 5.7: PDS derived from the full light curve from Figure 5.5 (S4). A HF feature is identified around 500 Hz with a harmonic at 250 Hz but no clear LF features are apparent.

5.7.1 Steady disc simulations

To increase the length of the simulation, we consider adding mass at large radii to mimic the effect of material accreting from a radius larger than the outer edge of our disc. This is achieved by injecting mass near the outer edge of the disc, with some material given a negative radial velocity (accreting) and some given a positive radial velocity (to conserve momentum). To represent a constant mass accretion rate, this material is added with the rate that mass is lost at the inner radius. To achieve this we utilised setup routines developed by Chris Nixon that calculate the surface density profile for a constant mass accretion rate and the rate of injection required to achieve this.

Recalling that the requirement for a constant mass accretion rate is $\dot{M} = \nu\Sigma$, the mass accretion rate must be consistent with both the $\nu(R)$ and $\Sigma(R)$ profiles. A given q parameter sets both the sound speed and scale height at the inner edge of the Σ profile. A first guess is made for \dot{M} and ν is calculated on a 1D grid using $\nu = \alpha\Omega H^2$, assuming Keplerian rotation. This allows the profile of Σ to be estimated, but because an arbitrary guess was made for \dot{M} it is not scaled correctly. From the unscaled Σ profile the mass of the disc is estimated and then used to scale both Σ and \dot{M} to the required mass of the disc. The scaled Σ profile is then allowed to relax on the 1D grid for a prescribed amount of the viscous time at the outer edge of the disc (we use 10%) and then both \dot{M} and Σ are again rescaled to the required disc mass. Once the Σ profile is defined the positions and velocities of the particles are set using a Monte Carlo approach. In previous chapters the contributions from the pressure gradients and external forces were set using analytical expressions, but here they must be interpolated from the results of the 1D code. Finally, the positions and velocities are rotated through the inclination angle β as described previously in § 2.2.1.

For our simulations, a compromise of $q = 0$ was found to be a relatively fast setup and a reasonable surface density profile. As the surface density profile is disrupted as soon as tearing begins, its actual profile is not critical to the evolution of the disc.

Second large disc simulation (S5)

We implement a large disc simulation with two changes from the previous simulation (S4). Firstly, mass is injected at $R_{\text{inject}} = 70R_g$ to ensure a constant mass accretion rate. Secondly, the inner edge of the disc is started at a larger radius of $40R_g$, larger than the estimate of the breaking radius. The combination of these two changes means that the simulation can be conducted for ≈ 20 seconds (keeping in mind effects from the outer boundary), but the first 10 of these only show material moving inwards to the tearing region. The length of this simulation corresponds to approximately 64 orbits at the outer edge or 0.73 of a precession at the outer edge. Even conducting the simulation for this long, the results at later times need to take into account the location of the outer disc.

Simulating for any longer would certainly require a larger outer radius, compromising the already marginal resolution.

Figure 5.8 shows the column density evolution of this simulation. In the first 10 seconds the outer disc precesses to a more face-on orientation and the inner edge of the disc moves towards $\approx 30R_g$. At ≈ 8 seconds, the first indication of tearing appears as the surface density decreases around $40R_g$. At 9 seconds, the ring has fully separated, precessing as a complete ring for 1.7 seconds. By 11.7 seconds, this ring separates into two smaller rings, and as the inner one precesses faster they have opposing inclinations by 12.2 seconds. The inner ring splits, forming two smaller, inner rings that quickly precess separately and then collapse into an inner, aligned disc. The outer of the original rings splits into two rings that precess together and reform into a large, precessing ring by 14.1 seconds. At 15.5 seconds this ring has completed a full precession and interacts strongly with the outermost disc before completely separating into two rings that collapse and align with the inner region before 18 seconds. This inner disc continues undisturbed for the rest of the simulation, and although no more tearing is observed the first signs of another tearing event are present by ≈ 22 seconds.

As the disc evolves in Figure 5.8, the particles that are injected in the original disc plane become increasingly visible as the disc precesses around but the region they are injected into does not. While these particles are injected with the intent to feed the outer disc so that the simulation can run for a longer time, their presence confuses the simulation. However, the simulation is not run for a full viscous time (i.e. the time it would take the entire disc to accrete). This means that by the end of the simulation, the majority of the disc has not been accreted. But this was one of the initial reasons for including the injecting particles; to try and keep enough mass outside the tearing radius to allow tearing to continue. Including these particles does not change the simulation as much as we thought, but including them confuses it. As a result, we do not include these in future simulations.

Figure 5.9 shows the light curve for this simulation, with the large peak around 11 seconds representing the first large ring tearing off the outer disc and the second peak a large ring merging event. When compared to the light curve that has an ‘outward-inwards’ tearing behaviour in Figure 5.6, the contrast is stark. In this case, each time a large ring separates into smaller rings there is a large, broad peak in the light curve rather than the numerous, distinct peaks seen in the previous simulation. The large initial drop in luminosity has been successfully avoided by using the steady state initial condition.

The lack of modulating features in the light curve leads to a featureless PDS, shown in Figure 5.10. Given that a tearing ring is observed in this simulation, the lack of features in the PDS suggests that the amplitude of the light curve features is not high enough or there are not enough repeated features (irrespective of their amplitude) for the PDS to identify a frequency on the simulation timescale. However, the HF feature is maintained (500 Hz and its associated harmonic at 250 Hz). The PDS shown in Figure 5.10 also shows

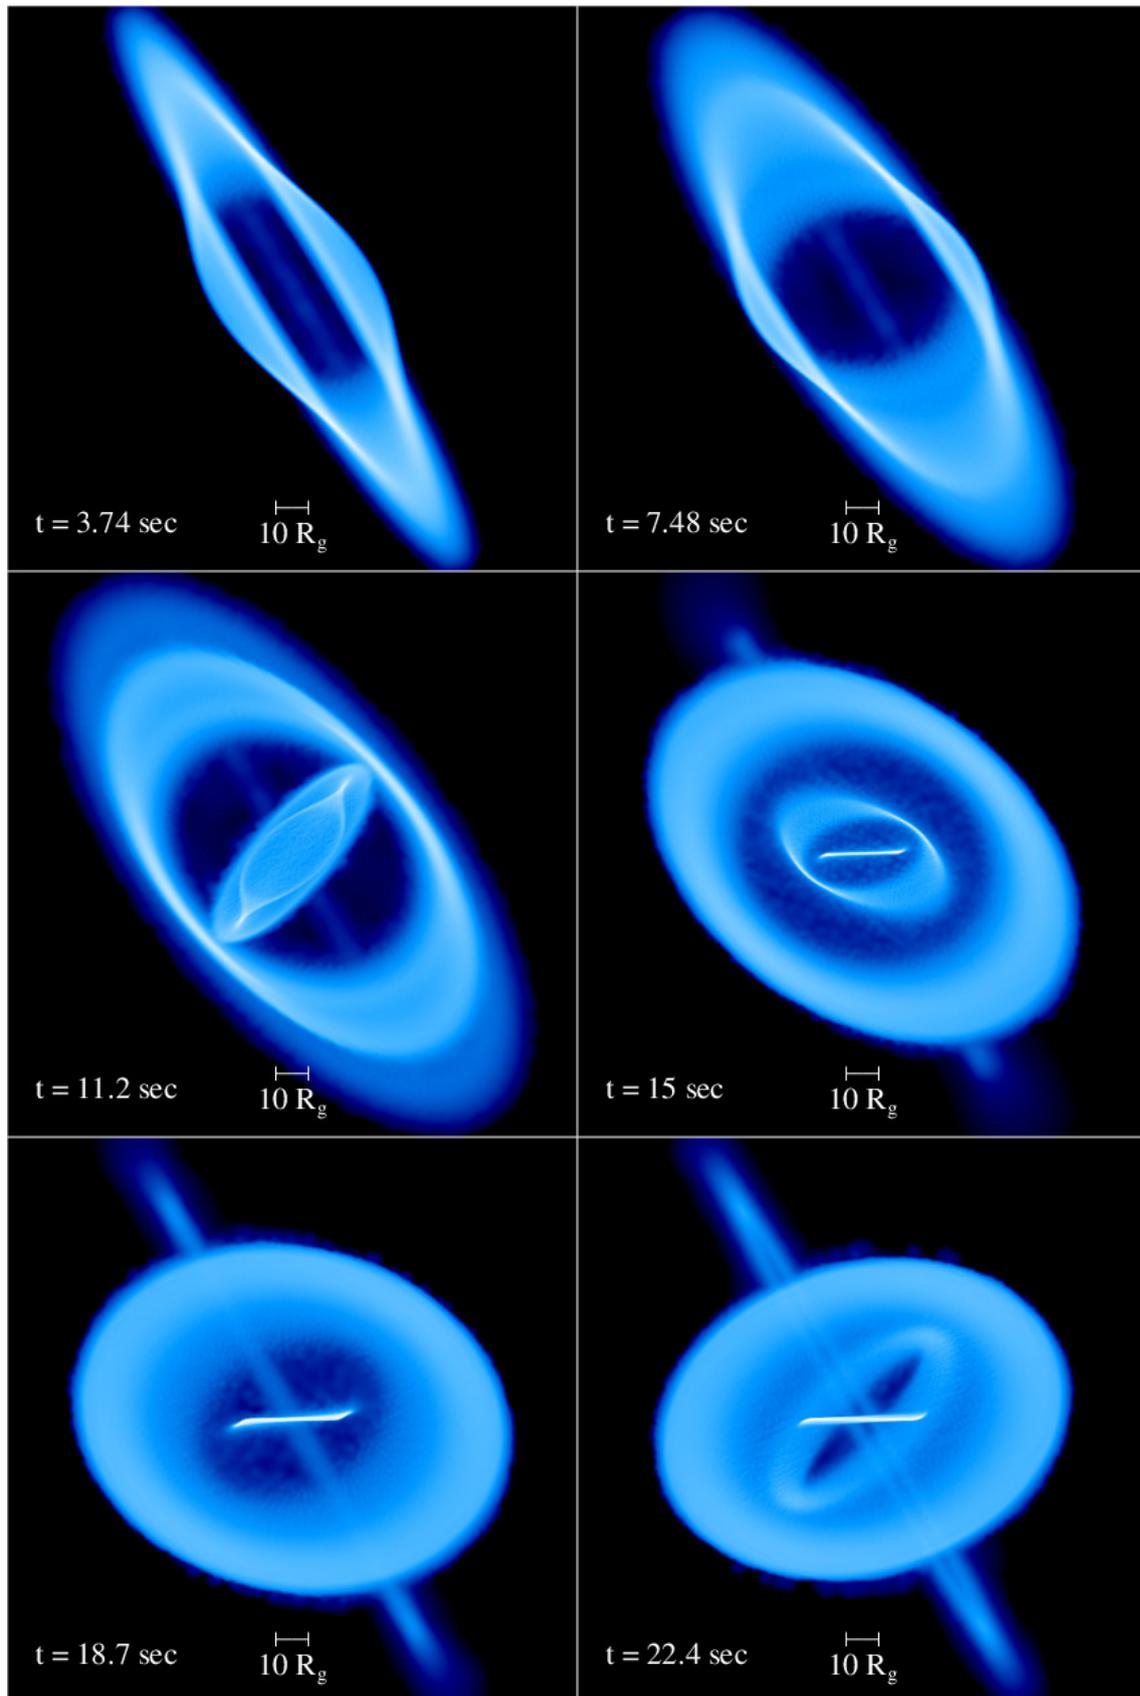


FIGURE 5.8: Column density rendering of the second large disc simulation (S5), with 10^6 particles and the black hole spin vertical on this figure. Particles start with an inner radius of $40R_g$ and move inwards through the tearing radius during the first 10 seconds of the simulation. Major tearing events are observed at 9 and 12 seconds, with a final one developing at 25 seconds. The ring of particles extending from the top-left to bottom-right in each panel are the injected particles — these particles are not included in subsequent simulations.

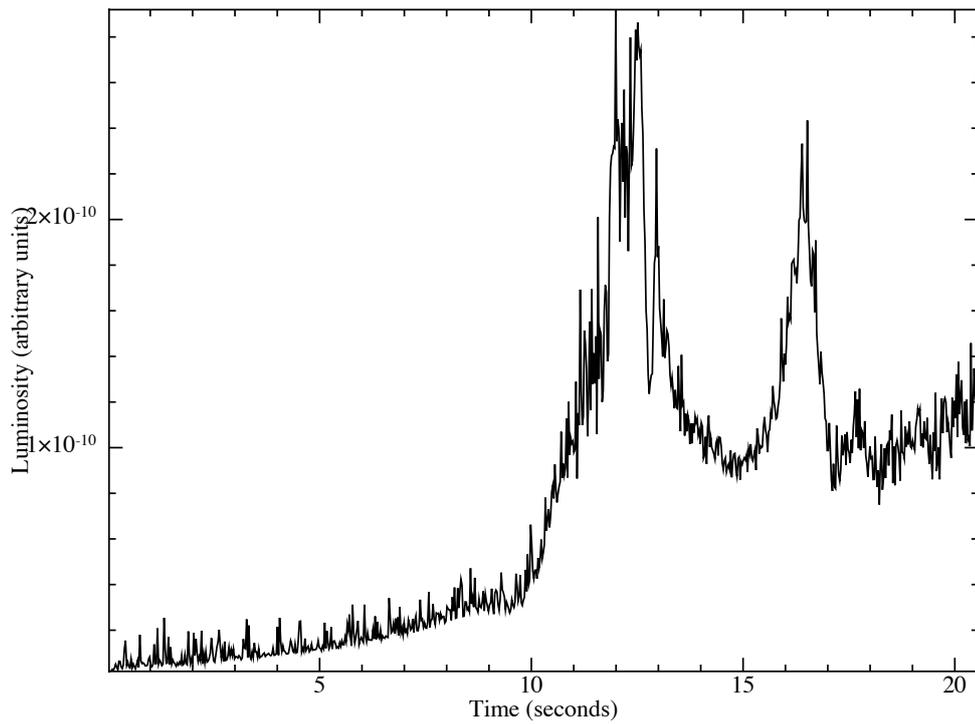


FIGURE 5.9: Simulated light curve from Figure 5.8 (S5). Here the two peaks are associated with the onset of tearing and when the inner ring present at 15 seconds merges with the innermost disc, corresponding to the shocks experienced by the gas during these events.

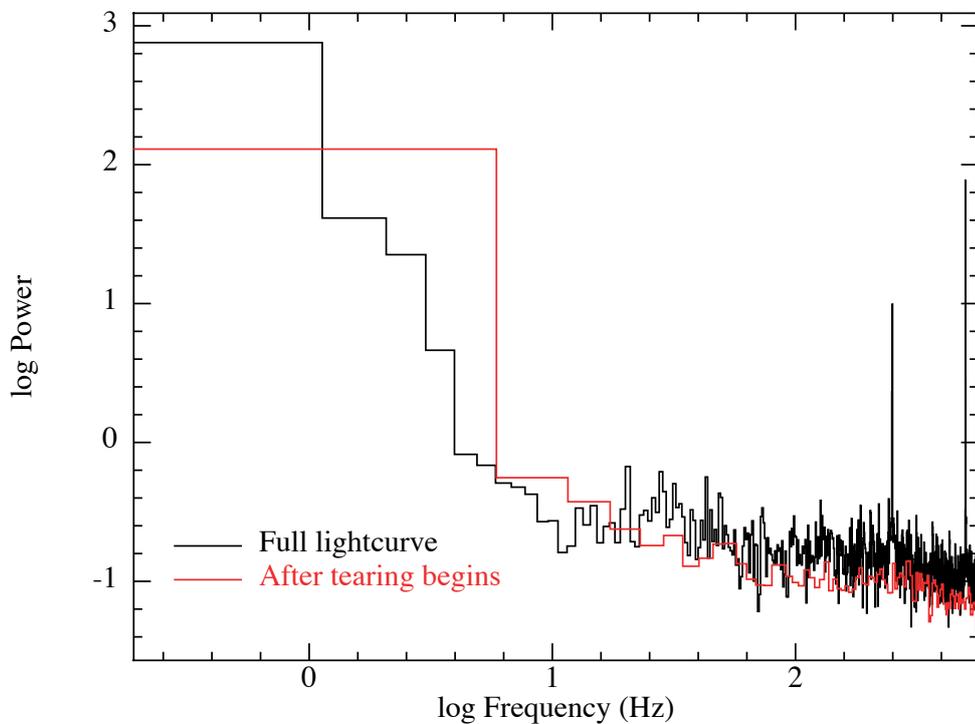


FIGURE 5.10: PDS derived from the lightcurve of Figure 5.8 (S5). As the black line shows the PDS from the entire light curve and the red shows the PDS from after the onset of tearing, we conclude that the high frequency features found in the PDS are not associated with tearing or precessing gas.

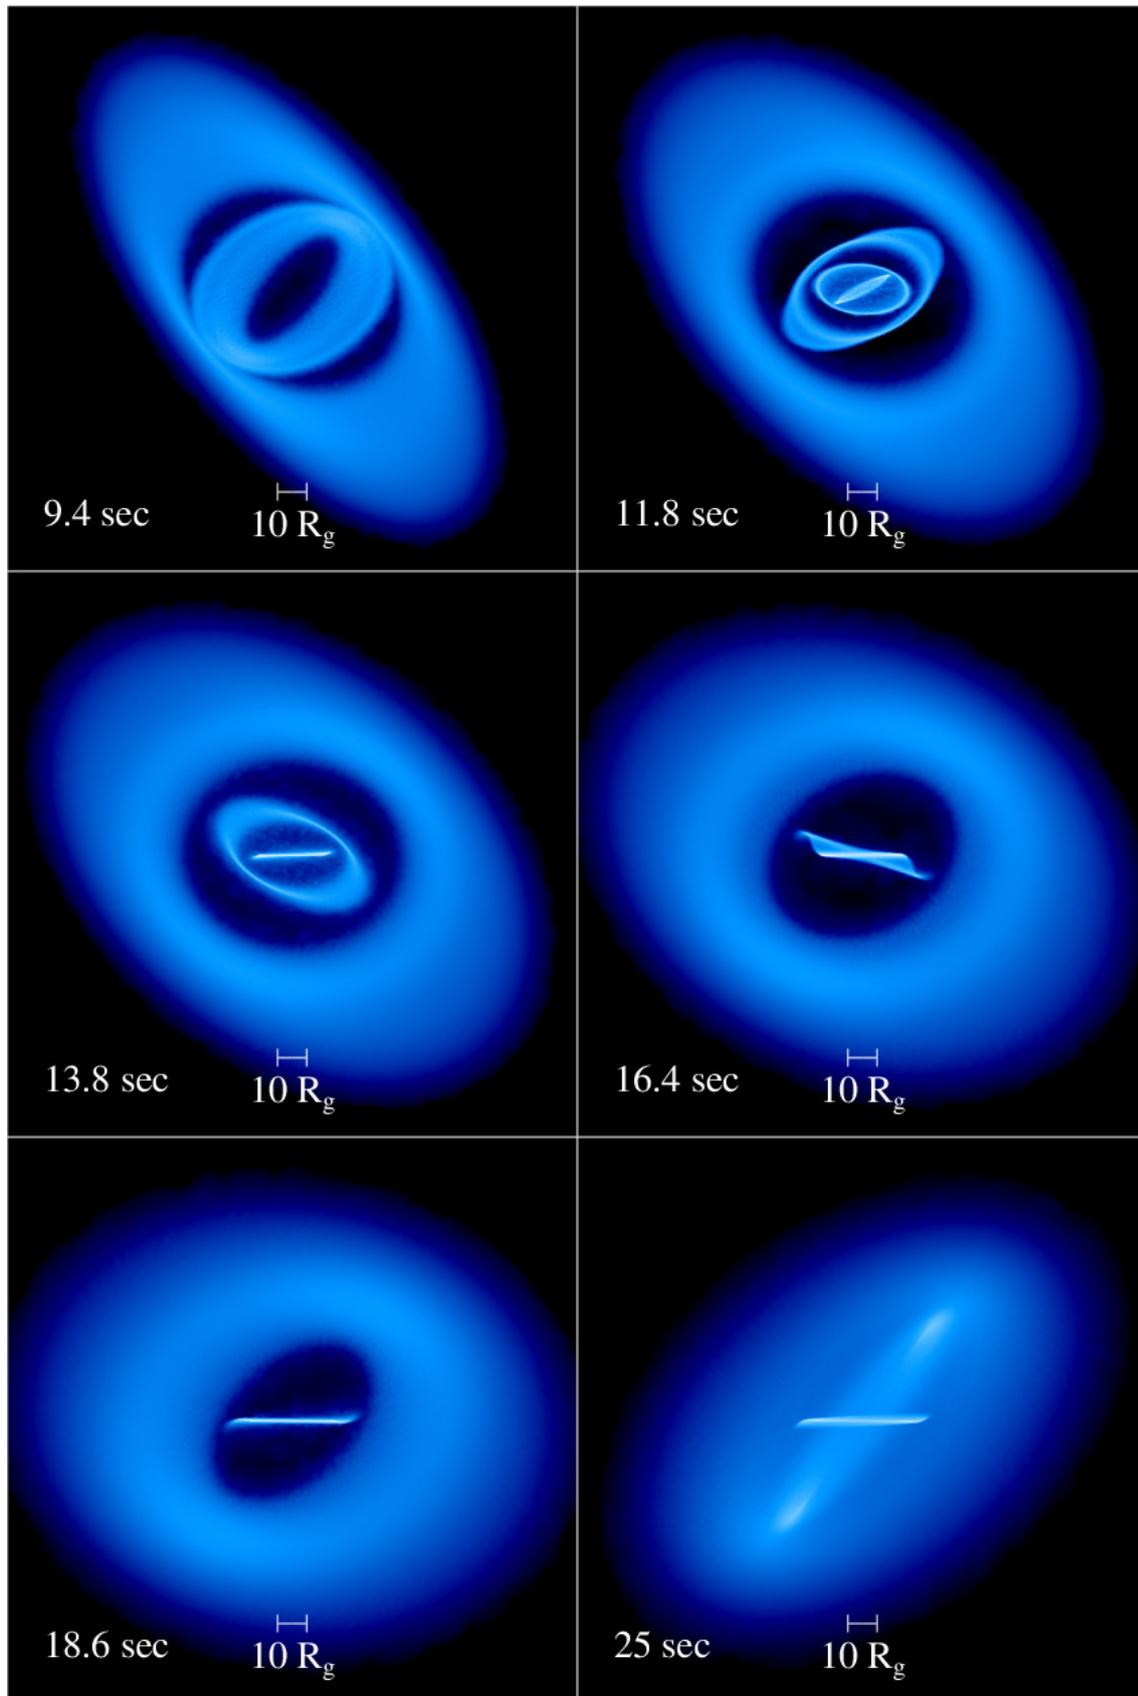


FIGURE 5.11: Density rendering of our final large disc simulation (S6), with 10^6 particles and the black hole spin vertical on this figure. Particles start with an inner radius of $40R_g$ and move in during the first 10 seconds of the simulation. Major tearing events are observed at 9 and 12 seconds, with a final one developing at 25 seconds. This simulation confirms that tearing is a repeatable process, as material continues to be fed from the outer disc.

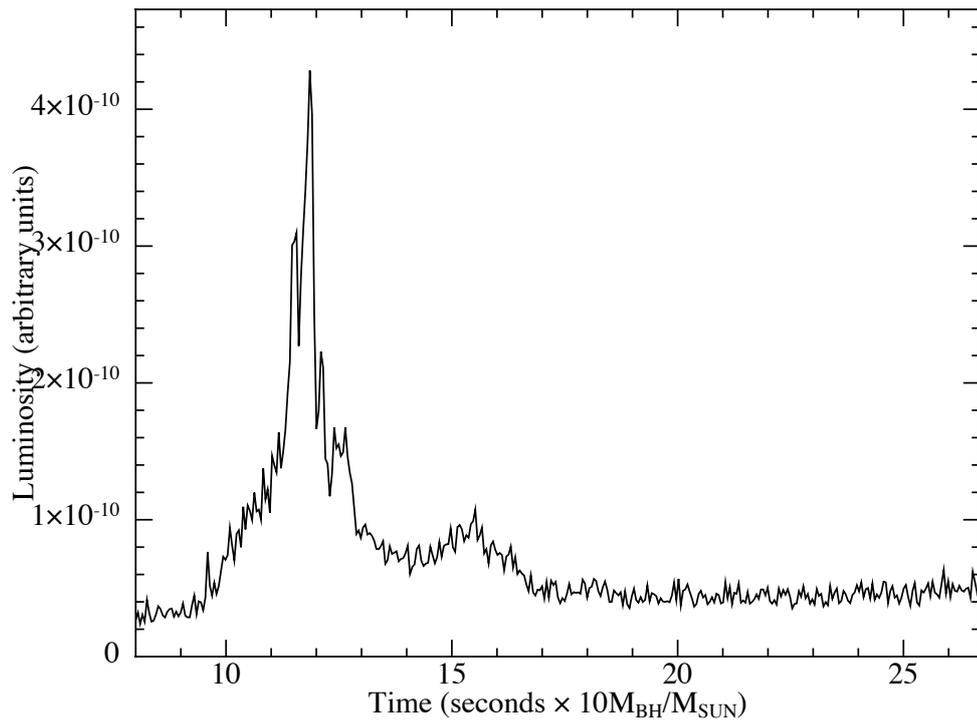


FIGURE 5.12: Simulated light curve derived from Figure 5.11 (S6). As this curve is qualitatively similar to Figure 5.9 only the region after tearing has commenced is shown. Again, the peaks are associated with strong shocks in the gas as a ring is formed or destroyed.

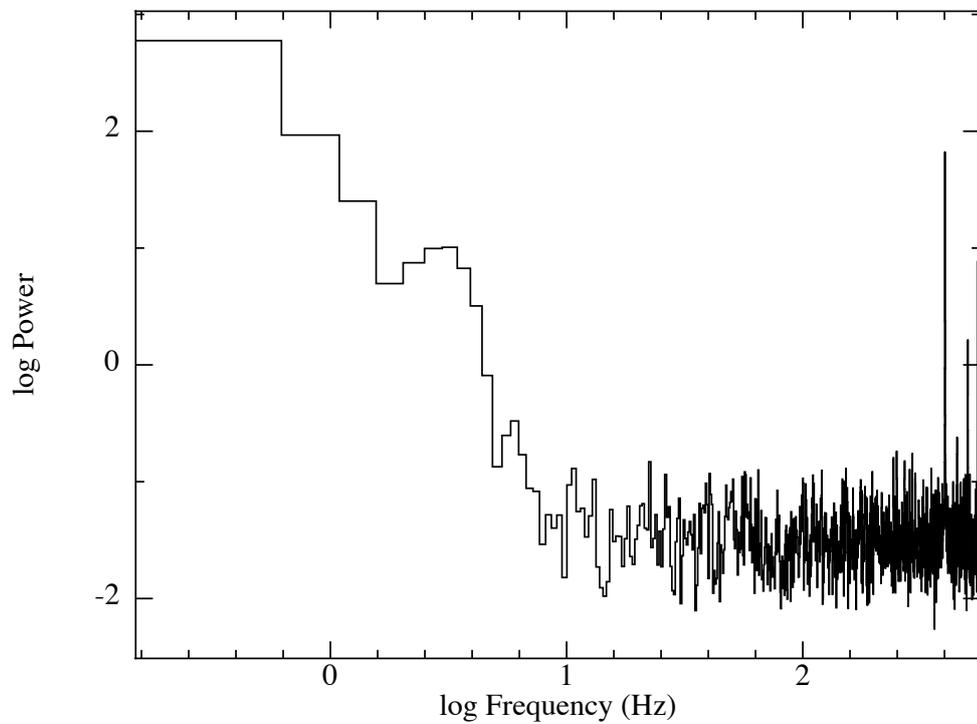


FIGURE 5.13: PDS derived from the final large disc in Figure 5.11 (S6). Again HF features around 500 Hz are identified (although we dismiss this, see Figure 5.14) as well as a potential broad feature around 1.5 Hz.

the PDS calculated from two different time frames in the light curve; the total simulation and the region only after tearing has commenced. A comparison of these shows that the HF feature is generated mainly during the start of the simulation — before any significant material has reached the inner region to develop an inner disc. This suggests that the HF feature is produced by particles that are not part of the ring or disc structures (i.e. orbits of individual particles that are not resolved).

Third large disc simulation (S6)

The final large disc simulation at this resolution brings together the initial conditions of the previous simulation but does not include particle injection. Figure 5.11 shows the column density evolution of this disc, similar in most regards to S5, with the initial signs of a second tearing event at 25 seconds. This simulation demonstrates the physical behaviour we are hoping to show: tearing events are capable of producing disconnected precessing structures and material around the inner edge and these tearing events repeat as long as there is material in the outer disc. Figure 5.12 and Figure 5.13 show the corresponding light curve and PDS, again similar to the previous simulation. Again, a HF feature is identified but in this case a potential LF feature is apparent ≈ 1.5 Hz.

The distinct HF feature at ~ 500 Hz is so rapid, it can only be caused by the orbital frequency (see Figure 5.1). However, solving Equation 5.4 indicates this must be sourced from particles $\sim 5R_g$, well inside of the structure observed in the simulations (which has an inner edge at $\approx 10R_g$). Further investigation, shown in Figure 5.14, identifies that this is probably caused by low resolution particles inside the inner disc. This figure identifies the edge of the innermost disc of gas ending at about $10R_g$, but individual particles exist inside this region as they are accreted through the ISCO (at $4.233R_g$). As they are individual particles, they are low resolution, causing them to be highly dissipative and hence luminous (four orders of magnitude more luminous than the gas in the surrounding ring). The luminosity inside the inner ring is thus due to numerical rather than physical reasons. Because they are just about to be accreted, their motion is strongly restricted to be in tight orbits, hence the narrow feature identified in the powerspectra. This fits with the PDS analysis (Figure 5.10) suggesting that the HF feature was generated before there was significant material in the inner region: individual particles were being accreted from the inner edge of the outer disc more than $20R_g$ from the black hole. As solitary particles, they were particularly low resolution and thus bright. These particles present a difficult problem: at higher resolution there will be fewer of them inside the inner disc, but they will be as luminous (if not more) because the local resolution will be lower. A potential solution to this would be to cut out the luminosity generated by particles with a large smoothing length, but this would have to be considered for future implementations. We thus disregard the HF features measured from these simulations as a numerical rather than physical feature.

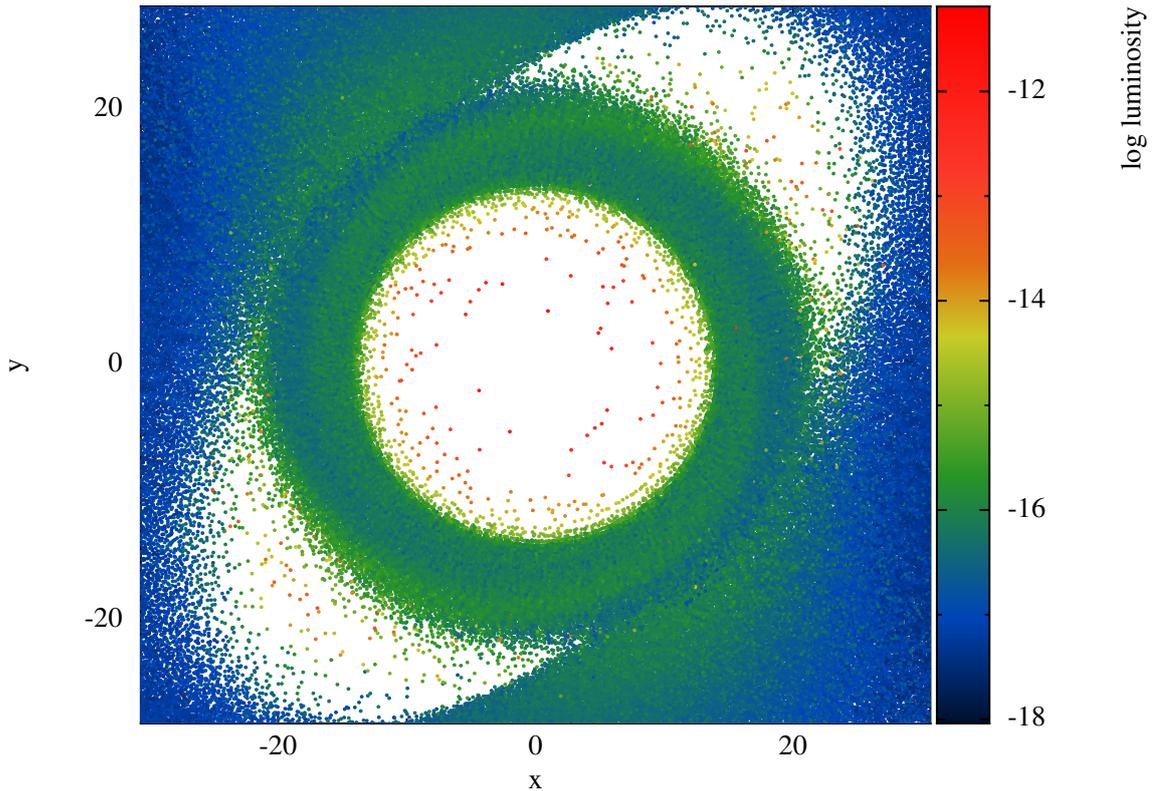


FIGURE 5.14: The source of our HF features in the PDS is particles in the innermost region that are artificially bright due to poor resolution. The particles have been coloured with the logarithm of the luminosity in a distinct colour scheme for clarity, showing the luminosity distribution in the inner disc. The particles inside the main ring structure are three orders of magnitude more luminous than the particles in the ring, and contribute disproportionately to the luminosity distribution.

The LF feature tentatively identified in Figure 5.13 cannot be straightforwardly associated with nodal precession. Firstly, although previous simulations have identified shocks in the disc due to apsidal precession, it is not obvious that nodal precession must lead to shocks — the inner structure could be nodally precessing and not luminous in our naive optically thin luminosity approximation. Secondly, even if this is a signature of nodal precession, there are not enough precessions for this feature to be significant. For the S6 simulation only 5 precessions are visually evident. However, each time there is a tearing event the size of the inner disc and hence the precession frequency changes, and so even these precession events occur with varying frequency. Extending the simulation even further to try and measure more precessions becomes hard to justify, as the outer disc boundary is already beginning to affect the simulation with the current duration (≈ 25 seconds).

5.8 A different approach to measuring the precession

The previous sections have demonstrated at least two limitations of the current luminosity approximation outlined in § 5.3. First, the innermost part of the disc is populated with particles that have an artificially high luminosity because they are poorly resolved. While these particles are strongly identified in the PDS, they are not connected to an extended

structure in the simulation. Future use of the luminosity approximation must mitigate contributions from such particles.

Second, it appears that the luminosity approximation does not necessarily demonstrate what we are hoping to measure. As it highlights shocks, it easily picks up the tearing and reformation of rings and discs. However, we are trying to measure the sustained precession of an inner disc — a structure that is not necessarily undergoing shocks. Indeed, the inner disc is only luminous when material is feeding it and hence increasing the outer radius. To measure a consistent frequency in the PDS the innermost disc should have a constant outer radius, and so we would want to measure it when it is not being fed and thus not luminous. Apical precession, as identified in the high resolution inner disc simulations (e.g. Figure 5.2) is identified by this luminosity approximation but only when the inner part of the disc is very well resolved and even then, it is not present in the PDS we measure from these simulations.

Finally, the validity of the opacity approximation is not clear, especially as the disc is assumed to be isothermal. The difficulty with these two assumptions is illustrated when considering neighbouring precessing rings. With the current approximations, these rings will be non-luminous when their angular momenta are similar, but as the inner one precesses around faster the luminosity will increase until they have opposing angular momenta (and shocks are occurring). As the inner ring continues to precess the luminosity will decrease until they have similar angular momenta again. Whilst this suggests a mechanism for the disc to have oscillations in the luminosity, this scenario is not physical and is unique to highly resolved isothermal tearing simulations. In a physical system, the gas should heat up as it shocks, generating a corona and invalidating the approximation that the disc is optically thin. In this case, fluctuations in the luminosity can only be generated by taking into account obscuration by less luminous material. This method has the advantage that it is closer to the accepted emission mechanism for producing type-C QPOs (e.g. Ingram et al., 2016).

These three considerations motivate the development of a different method to measure precession in the simulations and to identify this in a PDS. We consider measuring the precession of the gas in the disc dynamically and modifying the luminosity approximation to try to take into account the above three considerations.

5.8.1 Measuring the precession dynamically

As outlined in § 2.2.5, the rate of precession in the disc may be measured directly from the angular momentum components of the particles. The precession frequency of the material is determined directly from the rate of the twist, γ , and when considered as a function of radius can identify if a structure is precessing as opposed to free particles. The obvious strength of this method is that because it uses the dynamics of the gas it is independent of the luminosity approximation. This means that it can be used to independently verify any

features that may be found in the PDS that is generated from our synthetic lightcurves, although it cannot yet be used to create a PDS currently.

As in previous chapters, the disc is discretised radially into shells and the twist measured as a function of time for each shell (an example of this from Chapter 3 can be found in Figure A.1). With the twist, $\gamma(r)$, in the disc at a particular radius given as a function of time, the precession frequency of the gas ν_{gas} at a particular radius is given by

$$\nu_{\text{gas}} = \frac{1}{2\pi} \frac{d\gamma(r)}{dt}. \quad (5.9)$$

The precession frequency was measured from the slope of the twist ($d\gamma(r)/dt$) using a centred differences scheme, requiring at least three sequential points to measure a positive slope. When a negative slope was measured (i.e. sampling was across distinct precession periods, this occurs twice in the example of Figure A.1) the result was discarded. This method returns the frequency as a function of time for each radius (i.e. azimuthally averaged radial bin) but in the following figures is shown as the frequency as a function of radius for each time-step for clarity. We stress here that the previous simulations have not changed or been repeated, rather we are conducting a different analysis.

Figure 5.15 shows the frequency and surface density from S6 at a few representative time-steps. The innermost, noisy region reveals the frequencies of individual particles at ≈ 10 Hz. The smooth curve at $\sim 20R_{\text{g}}$ documents the inner precessing structure, with frequencies of consistently $\lesssim 5$ Hz. The lack of a corresponding feature at these low frequencies in Figure 5.13 confirms that the current simulations are not able to identify precessing structures in a PDS. The resolution used, the luminosity approximation or the duration of the simulation may all be responsible for the lack of features.

5.8.2 Luminosity projection with obscuration

Although consideration of the angular momentum components confirms that the inner disc is precessing, it cannot be used to generate a PDS. In an effort to assess the viability of approximating the disc as optically thin, we consider measuring the luminosity of the disc in a way that takes into account obscuration. Motivated by oscillations in the light curve being generated by obscuration from precessing structures, this light curve would take into account when cooler, less luminous material passes between the inner bright disc and the observer.

Here we use the rendering function of the visualisation tool SPLASH (Price, 2007) to generate a light curve that includes the effects of obscuration. First, the luminosity was displayed with 3D surface rendering, producing a luminosity surface (discussed in depth in Price, 2007). The same process was used to render a density surface in previous chapters (e.g. Figure 3.9). In SPLASH this is achieved by ray tracing through the particles and tying the optical depth to the local density. This choice means that low luminosity regions

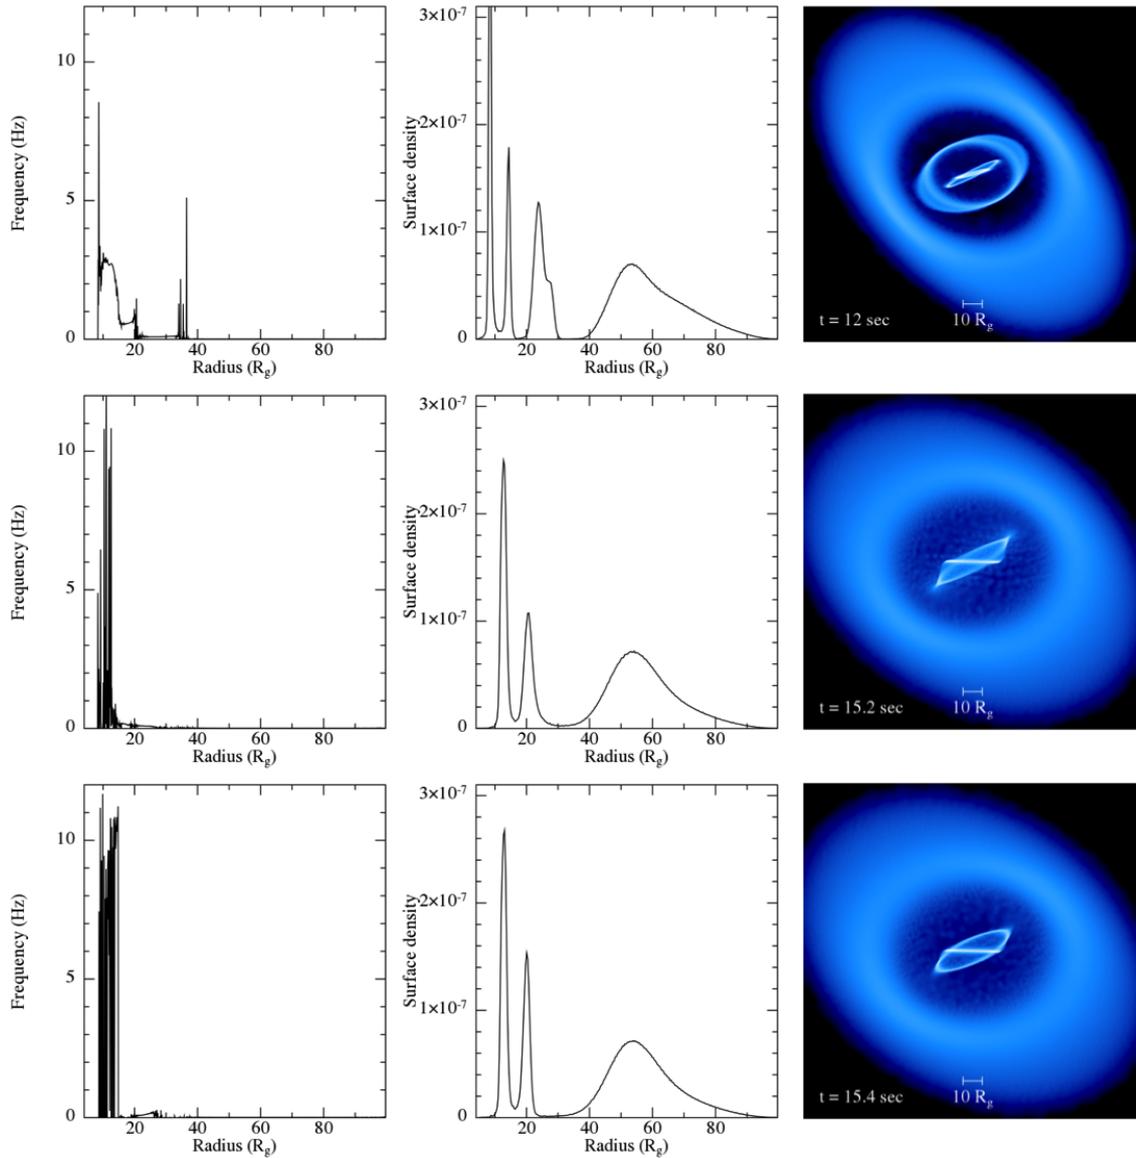


FIGURE 5.15: Comparing the precession frequency of the gas with the surface density profile and the visual structure in the disc. Precessing structures can be identified by sections of the disc precessing at a shared frequency (e.g. top panel). The higher frequencies ($\gtrsim 5$ Hz) with a discontinuous frequency profile are generally associated with gas between disc structures. In the final panel, these frequencies vary on short timescales despite little change in the disc structure. The lower frequencies (~ 1 Hz) are associated with a continuous frequency-radius profile, i.e. a radially extended structure. This comparison demonstrates that the PDS we generate with the luminosity approximation in § 5.3 cannot distinguish between precessing structures (e.g. rings) or precessing gas (i.e. the gas between the rings).

appear as transparent and high luminosity regions as opaque. Second, these visualisations were written to an ascii file as a pixel map using the ‘SPLASH to ascii’ feature. Because the visualisations are rendered, these pixel maps only include luminous regions that are naturally visible to the observer while less luminous regions or luminous regions on the far side of the disc are not included. Finally, the values in each pixel map were summed giving a total luminosity for each output file that is visualised. Thus the total luminosity, including viewing dependence and the effects of obscuration, could be generated without re-running the simulation. The difficulty with this method is its use of visualisation outputs from PHANTOM; these outputs are only generated every 0.02 seconds (of physically scaled time), so our PDS can only consider frequencies slower than this and so we cannot consider the HF features anymore. However, this places the upper limit on the frequencies we measure of 50 Hz so we can still investigate LFQPOs.

Third large disc (S6) with obscuration

To test the new approach proposed above, a number of subtly different variations of the light curve with obscuration from S6 were used to generate PDS. To quantify the effect of low resolution, artificially bright particles, a PDS with all particles was compared to a PDS excluding particles with smoothing length $h > 2.0$ (Figure B.6). This comparison still identifies that the low resolution, luminous particles contribute significantly to the PDS generated. From here onwards we use the smoothing length criteria to try to eliminate the contributions from these particles in the simulated PDS.

PDS using only well resolved particles ($h < 2.0$) were generated from different sections of the light curve and from different isolated regions of the disc, shown in Figures 5.16 and 5.17. From these light curves we conclude that the greatest contribution to the shape of the PDS comes from the material in the inner $25R_g$ during the first tearing episode. Importantly, this is *before* the formation of the first inner ring, confirming that the PDS features we measure from the full lightcurve are not due to the precession of the innermost structure. The PDS generated from the second half of the light curve, where this structure is present, does not show any clear features.

Figure 5.18 shows PDS generated from different viewing inclinations, finding little difference in the PDS for each inclination. If the features being measured are indeed a geometric effect, then a strong change would be expected between inclinations. Because this is not the case, we conclude that the luminous particles that are being identified must not be obscured at any point, independent of the viewing inclination.

From the three above approaches to calculating the PDS we have shown that the largest contribution to the shape of the PDS is from the inner $25R_g$, before an inner ring has formed and before precessing rings obscure the viewing of the central region. Particles in this region are luminous because they are lower resolution, despite our attempts to remove low resolution particles from the simulation by introducing a smoothing length

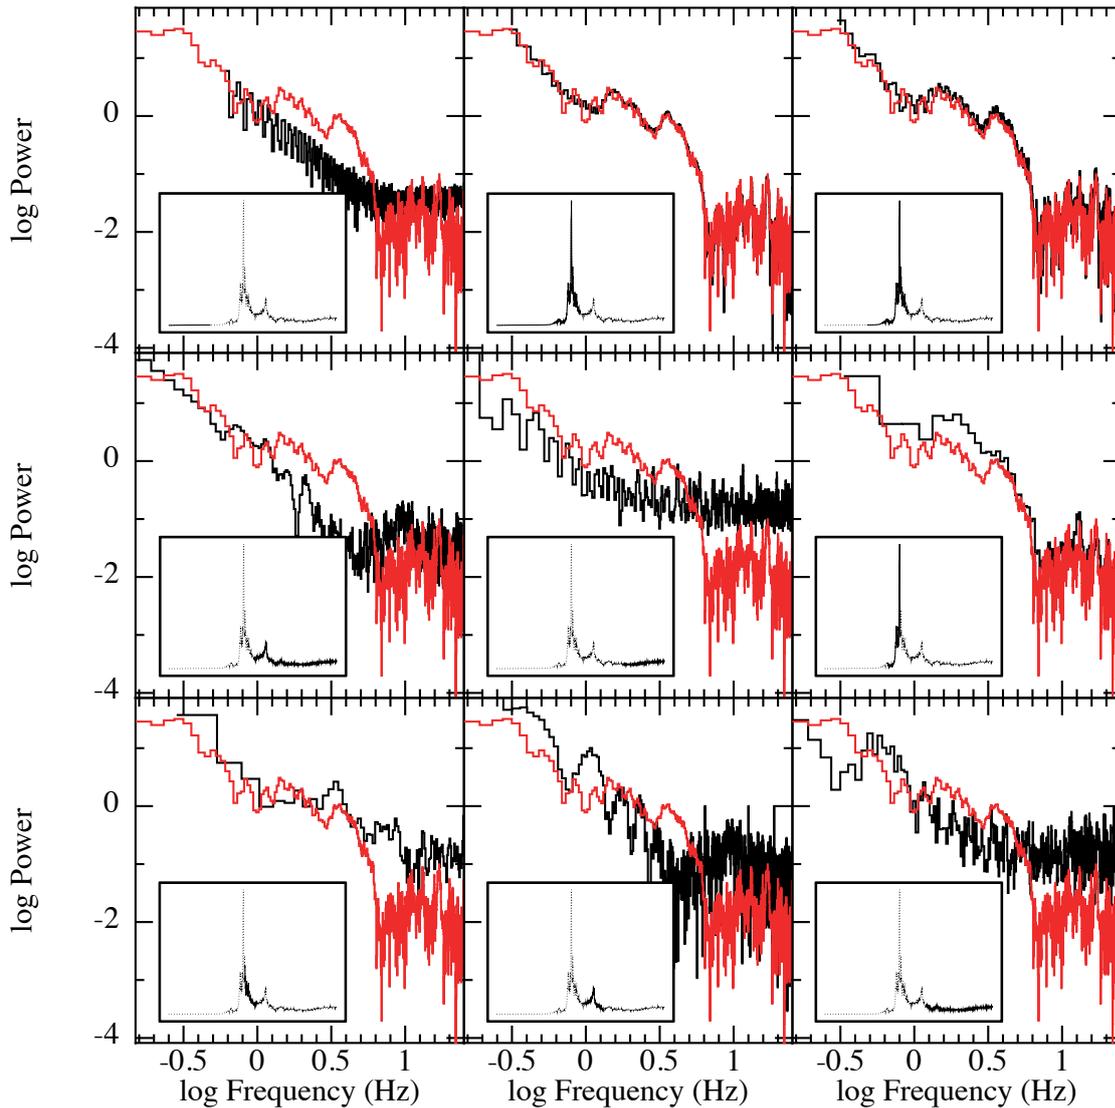


FIGURE 5.16: PDS from S6, including the effects of obscuration (see text), from different sections of the light curve (Figure 5.12) indicated in the insets. The red line represents the PDS from the full light curve, so the closer the match between this and the PDS in black the more this section of the light curve contributes to the full PDS. The best comparison in the upper, middle and upper, right panels suggest that the features in the PDS come mainly from the initial tearing event rather than subsequent precessions of rings.

restriction. Discarding the first half of the light curve (where these particles make their largest contribution) demonstrates no clear features in the PDS.

In a final attempt to mitigate the impact of the lowest resolution particles, we introduce an artificial bright source at the centre of the simulation (where the black hole is). This source is given a constant luminosity that is two orders of magnitude larger than the highest luminosity generated in the simulation (not including the particles discarded by the smoothing length restriction). We repeat the process of generating a light curve with obscuration using SPLASH and a pixel map as before. Here, the flux measured at each step can conceptually be thought of as light not blocked by the disc structures between the

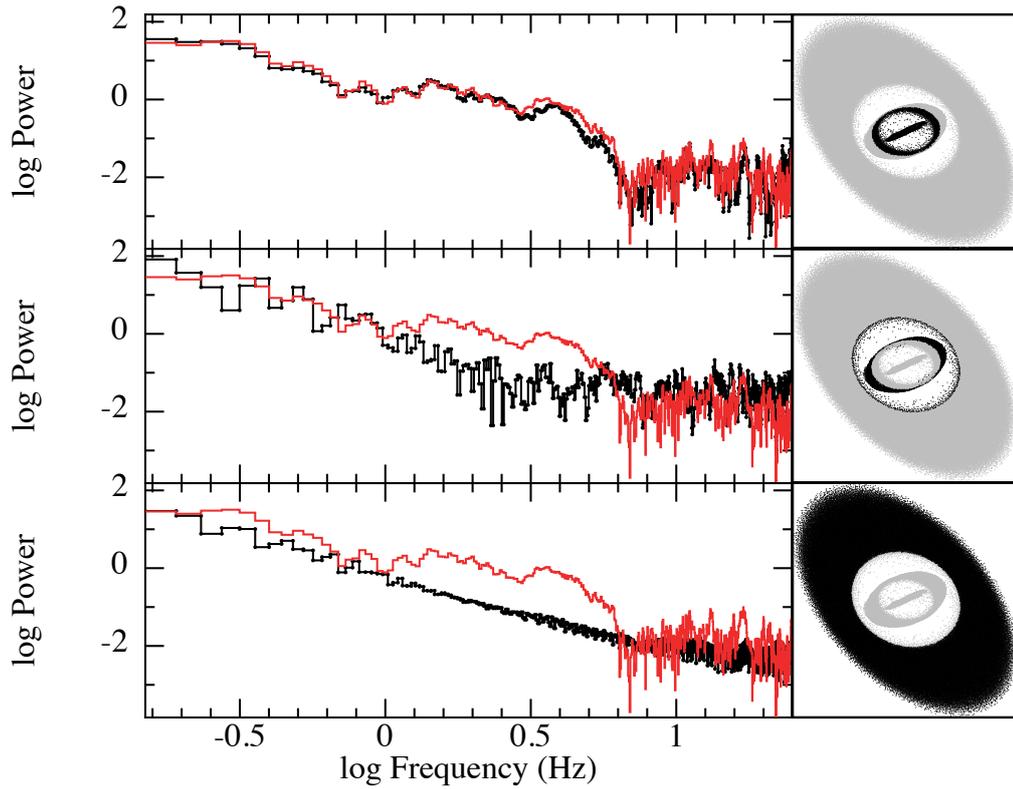


FIGURE 5.17: PDS from S6, including the effects of obscuration, from different sections of the disc indicated in the insets as the black sections of the disc. The red line represents the PDS from the full disc, so the closer the match between this and the PDS in black the more this section of the disc contributes to the full PDS. The better comparison in the upper panel suggests that the PDS features are mainly driven by gas in the inner regions.

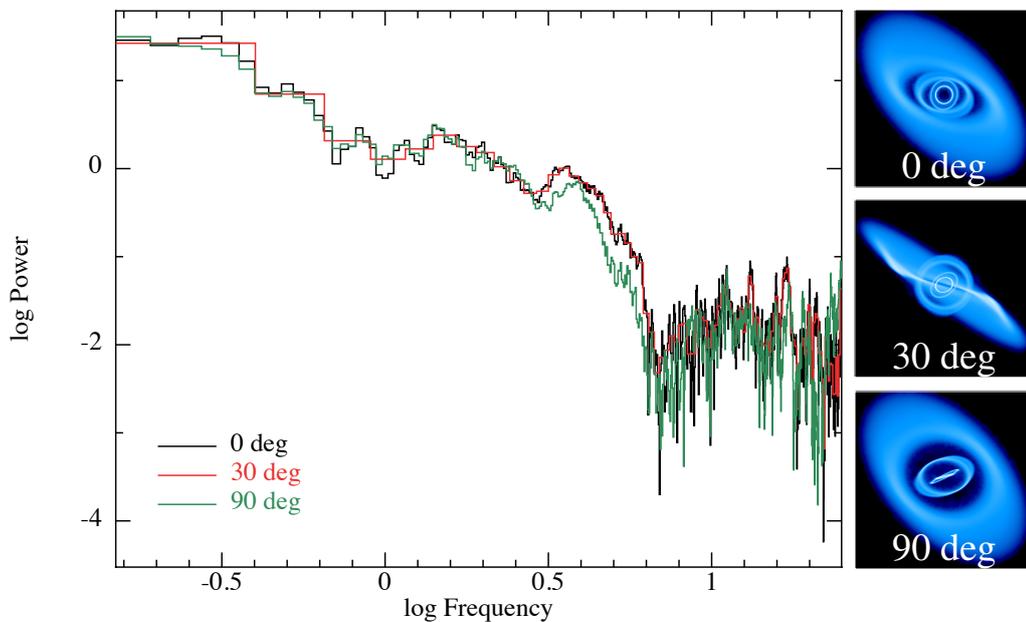


FIGURE 5.18: PDS from S6, including the effects of obscuration from different viewing inclinations of the disc indicated by the inset. Because changing the inclination results in little change in the calculated PDS, the features being measured by the PDS must not be a geometric effect — this is at odds with our understanding of both LFQPOs and our interpretation from our simulations.

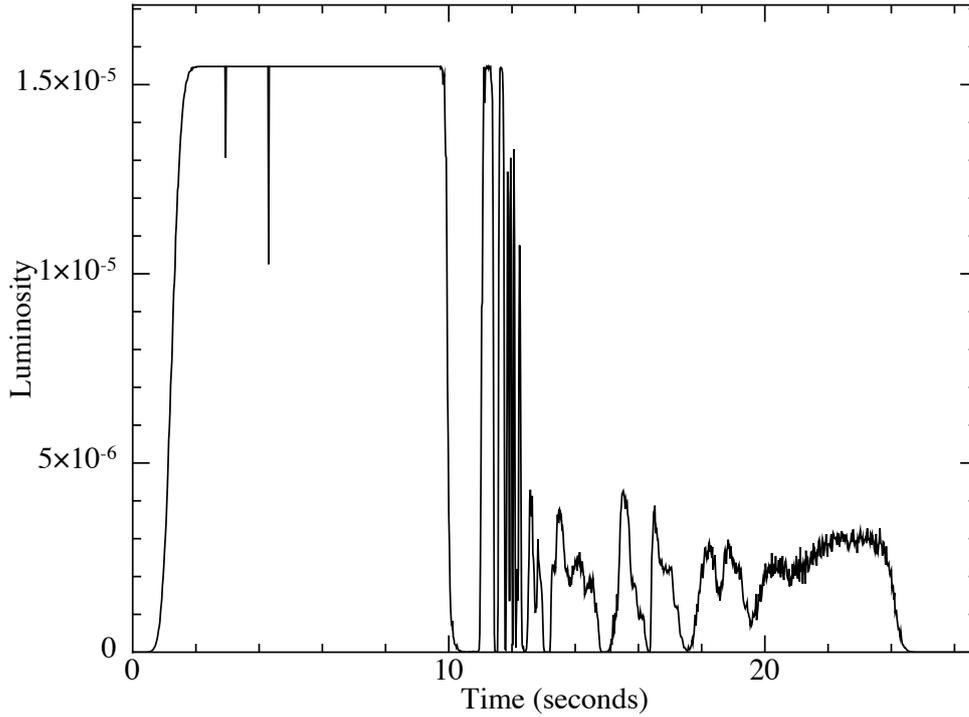


FIGURE 5.19: Simulated light curve from Figure 5.11 but including an artificial, bright point source at the centre of the simulation. A viewing point is assumed and the lightcurve is generated by measuring the obscuration of this bright source by the precessing rings (tearing does not occur in the first ten seconds so there is no obscuration). This method is similar to representations of LFQPO models (e.g. Ingram et al., 2016) and does not depend on the interactions between the rings.

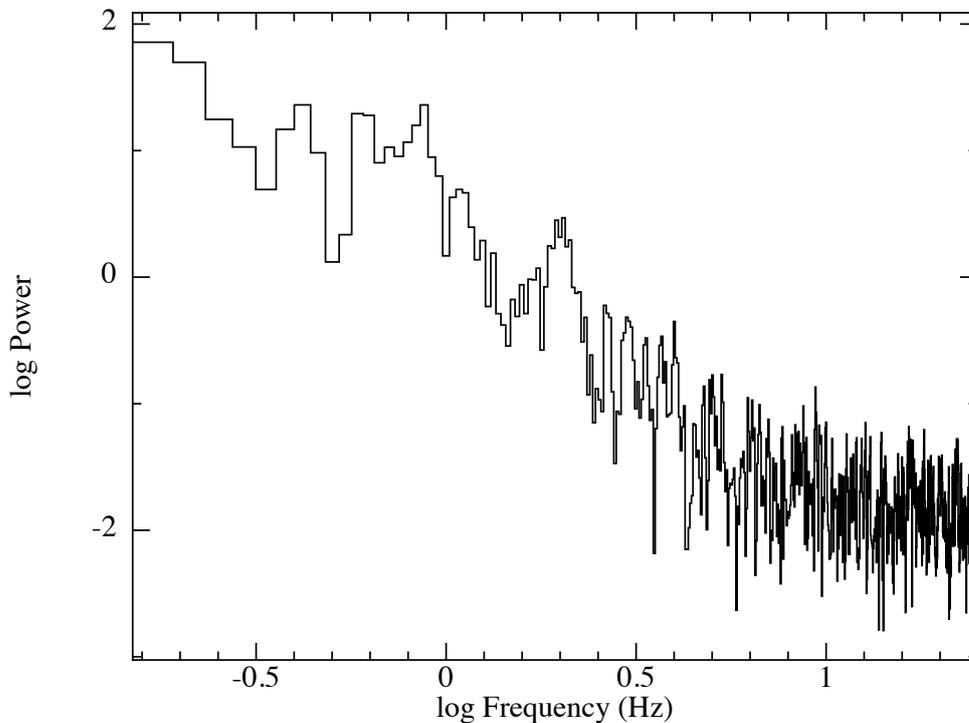


FIGURE 5.20: PDS generated from the light curve shown in Figure 5.19, which includes obscuration and an artificial luminosity source. As with all the other PDS generated with obscuration, the disc is viewed at an inclination with maximum obscuration. A tentative LF feature is identified ~ 2 Hz.

artificial bright source and the observer. Figure 5.19 shows the resultant light curve in this case; there is no tearing and thus no obscuration before 10 seconds but each subsequent dip in the light curve is caused by a ring passing in front of the artificial source.

We generate a PDS from this light curve, including our smoothing length restriction to avoid low resolution particles and only considering the second half of the light curve. Figure 5.20 shows that despite all the difficulties that we have tried to mitigate, the PDS is again inconclusive. We repeat the analysis described above by generating PDS from different sections of the light curve, different regions of the disc and at different viewing inclinations (Figures 5.21, 5.22 and 5.23). These results confirm that the features in the PDS are mainly generated from the innermost $25R_g$ only when tearing is occurring — the region that is predominantly sampled in Figure 5.20.

As Figure 5.20 does display our most physically motivated PDS, we try to use our previously described significance test to see if any of the potential features are statistically significant. The broad feature at ≈ 2 Hz in particular is consistent with our findings from the dynamical measurements, and visually it appears as though it may be significant. Because here our data are evenly sampled we assume that the number of independent frequencies sampled is equal to the original number of data points that were used to calculate the PDS (i.e. $M = \text{length of data sampled}$). We solve Equation 5.7 for the power that would be required for $p = 0.5, 0.05$ and 0.01 . These significance levels are included in Figure 5.24, showing the low-frequency region of Figure 5.20 with linear scales. Here we identify that the feature at ≈ 2 Hz is not significant while 0.42, 0.61 and 0.87 Hz are. These features at < 1 Hz may correspond to the precessing structures identified using the dynamic analysis in Figure 5.15. However, these features are only identified in a single simulation — these would need to be reproduced at higher resolution in order to be confirmed. Additionally, we have not tried to take into account the broad trend of increasing power at low frequencies (e.g. the fit that is in Figure 1.13), and so we additionally expect more noise at lower frequencies due to low sampling. As a result, we do not consider the features identified in Figure 5.24 as significant.

Underlying this entire analysis is the acknowledgement that the simulation may simply not be long enough to clearly measure precessing features or high enough resolution. For example, the light curve in Figure 5.19 only shows five distinct dips that are attributable to a precessing structure. The effect of the length of the simulation can only be quantified by simulating for longer, but we are already constrained by the strict balance between the outer edge, the resolution and the length of the simulation. We next consider increasing the resolution of our simulations to shift this constraint.

5.9 High resolution simulations (S7 and S8)

The previous discussion of low resolution particles prompts a resolution study, as at higher resolution these particles should have a smaller effect. Up to this point, the simulations

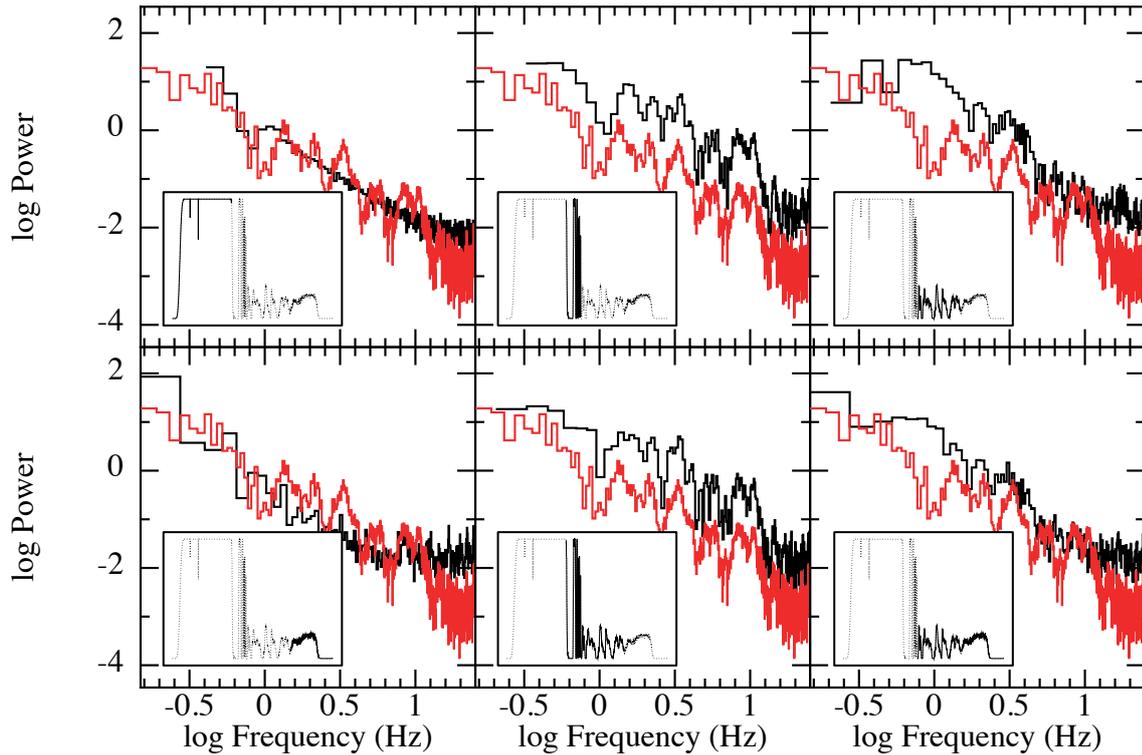


FIGURE 5.21: PDS from S6, including the effects of obscuration with an artificial luminosity source, from different sections of the light curve (Figure 5.19) indicated in the insets. The red line represents the PDS from the full light curve, so the closer the match between this and the PDS in black the more this section of the light curve contributes to the full PDS. The better match to sections of the lightcurve that involve tearing demonstrate that this method of measuring the lightcurve (e.g. upper, left panel compared to lower, middle panel) only takes into account the precessing gas (as opposed to Figure 5.16 which models the tearing gas instead).

made use of a maximum of 10^6 particles. In each simulation the tearing is only marginally resolved — that is, the simulations are conducted at the minimum resolution to guarantee tearing. This criteria means that resolution effects cannot be investigated by going to lower resolution, as the behaviour of the disc will be different and tearing will not occur. Importantly, although tearing is resolved with the current number of particles, this not necessarily true for the individual rings. As a result, at higher resolution the rings generated will be qualitatively similar but narrower and will last longer (as they do not interact with neighbouring rings as strongly). This will affect the frequencies measured from our simulation, and so is an important consideration.

S6, which does have structures precessing at $\approx 1-2$ Hz, is repeated at an increased resolution of 10^7 particles (S7). Simulating for the same physical time as the lower resolution simulation (e.g. ~ 20 seconds) is extremely computationally expensive with 10^7 particles. On the Swinburne g2 cluster with the current computational resources (16 cores) it was estimated this would take more than a year for the simulation to be completed. In an effort to conduct the simulation faster we made use of the Monarch cluster, where up to 24 cores are available. With these resources the simulation was still estimated to take more

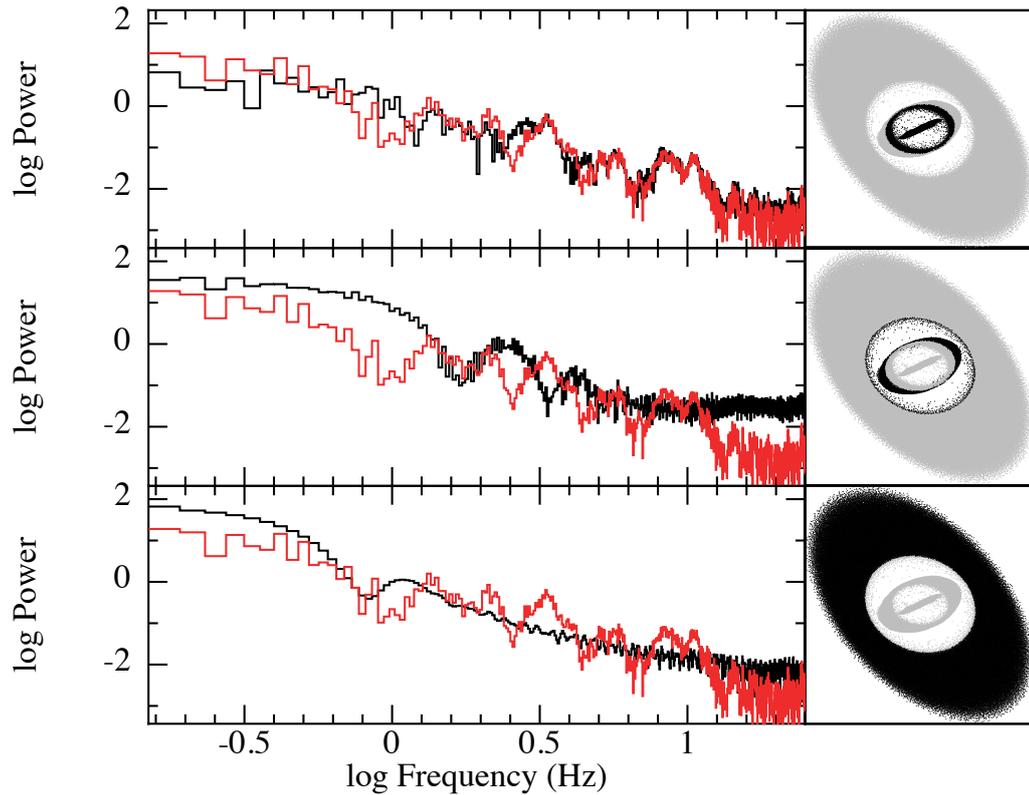


FIGURE 5.22: PDS from S6, including the effects of obscuration with an artificial luminosity source, from different sections of the disc indicated in the insets as the black sections of the disc. The red line represents the PDS from the full disc, so the closer the match between this and the PDS in black the more this section of the disc contributes to the full PDS. Here we again confirm that the gas in the innermost section drives the features in the PDS.

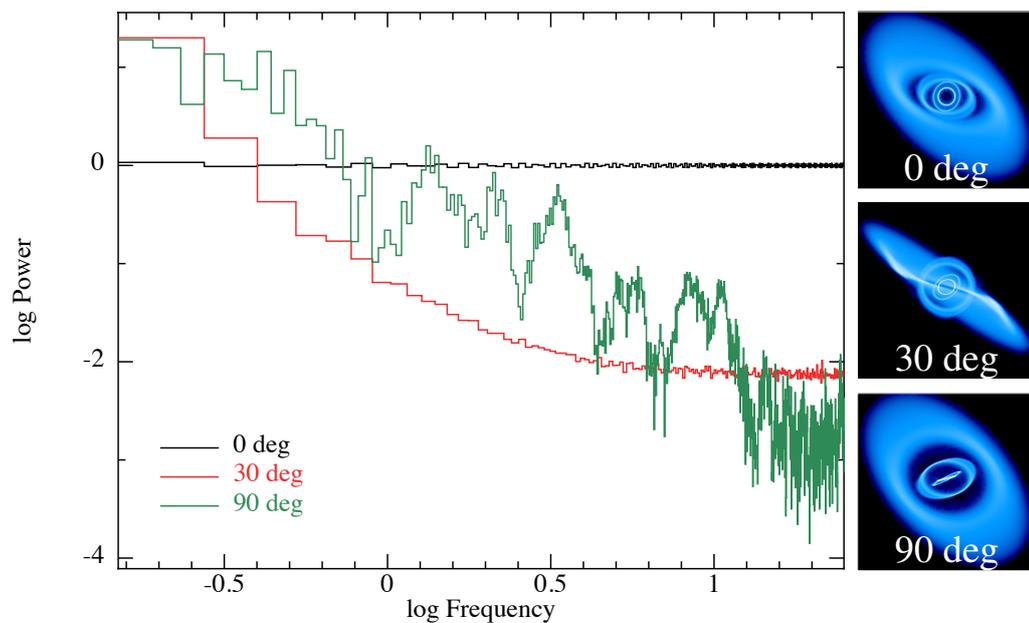


FIGURE 5.23: PDS from S6, including the effects of obscuration and an artificial luminosity source from different viewing inclinations of the disc indicated by the inset. The strong inclination viewing dependence of this method is more physically motivated than our previous method (e.g. Figure 5.18).

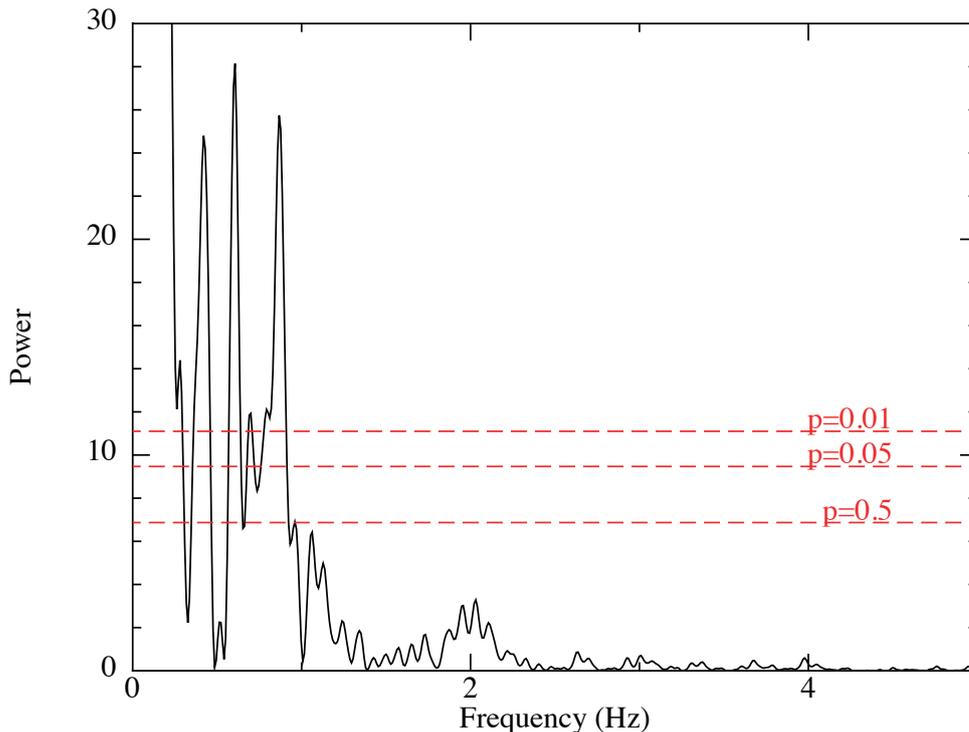


FIGURE 5.24: Significance of low-frequency features identified in the PDS in Figure 5.20. The significant features identified at 0.42, 0.61 and 0.87 Hz are tentatively associated with the structures precessing at <1 Hz, measured dynamically in Figure 5.15.

than 9 months and was not likely to reveal more precessions of the innermost material — rather we expect it to better resolve the interaction of the rings.

As the slowest part of the simulation is the orbits just outside the ISCO, a final simulation is initiated with the same parameters as S6 but a larger inner radius (S8). The lower resolution simulations suggest the inner truncated disc has an inner radius of $\approx 8R_g$ (e.g. Figure 5.15) and at higher resolution this is expected to move further inwards. Choosing an inner radius that compromises the inner edge of the disc will affect the frequencies measured, so a conservative value of $R_{\text{in}} = 6R_g$ was chosen for this final simulation. This simulation is predicted to take a shorter time than S7, but precisely how much shorter is difficult to gauge.

At the time of submission, these simulations were incomplete (they have approximately 8 seconds of tearing behaviour). In the context of the analysis that has been completed for the lower resolution cases, we note that these simulations are not guaranteed to generate a PDS with a significant LF feature.

5.10 Summary of QPO simulations

Here we summarise the findings from the (more than 200) simulations that we have conducted to try and show the connection between tearing discs and QPOs. Although the

high resolution simulations are continuing, we have identified the following important considerations for future work:

- Long duration, high resolution simulations appear unavoidable in order to demonstrate the connection between QPOs and tearing in discs.
- The general behaviour of the disc is demonstrated to be consistent with our previous simulation; material tears off the outer disc, cascading into precessing rings that interact and accrete onto the black hole. This process is demonstrated to be repeatable and the gas in the rings precesses with the frequencies expected from observations (although slightly lower than these).
- If the resolution is high enough in the inner parts, apsidal precession is demonstrated in the inner regions (although not yet measured in the PDS).
- The artificially high luminosity contribution from low resolution particles must be taken into account or preferably eliminated.
- Features in the PDS can be confirmed dynamically, using precession timescales estimated from the twist in the disc.
- A cooling prescription is important here as the tearing process is expected to generate hot gas with a range of cooling rates. From previous work the gas that cools slowly will form a corona while the gas that cools quickly will form a disc. Without a cooling prescription, this behaviour cannot be included.

Despite what we have learnt using the luminosity approximation, this work suggests that future investigations may require a more sophisticated method of generating and measuring flux in the disc (i.e. detailed radiation transfer calculations), which is currently not possible with PHANTOM.

5.11 Conclusions

Although the simulations shown in this chapter do show discrete, precessing disc structures that may be associated with LFQPOs, it is difficult to identify such features in the PDS due to the relatively short simulation time. While we are able to identify the precession of gas in the inner region after the initial tearing using a PDS, we are not able to confirm any significant features. Here we recall and respond to the questions posed in § 5.1:

1. We have demonstrated that a disc with the parameters of an observed LMXB and strongly inclined can undergo disc tearing. Here we assumed a large inclination to guarantee disc tearing, but using the previous analytical estimates tearing is permitted in these systems at lower inclinations as well. Additionally, although we cannot simulate the full radial extent of the disc (e.g. the order of the binary separation), we do not expect this to affect the tearing behaviour significantly.

2. Our simulations confirm that tearing from the outer disc is a repeatable phenomenon: material starts outside the tearing region, moves inwards on the viscous timescale and then a new ring is torn off the outer disc. As shown by simulations S5 and S6 (and by the high resolution simulations that are still running) the behaviour is repeatable as long as material continues to be supplied from the outer disc.
3. Dynamical measurements from the rings produced in our simulations confirm that they do precess with frequencies consistent with LFQPOs — we note however, that the range of frequencies observed from this particular LMXB mean that any frequency less than ~ 30 Hz is consistent with observations. Importantly, although the tearing from the outer disc is numerically resolved, the interactions between the rings are not. As a result, we expect the frequencies that we measure from higher resolution simulations to be different from what we have shown here.
4. We have explored a few different methods of generating a PDS. Initially, we considered the luminosity of the disc itself, approximated by the energy liberated through shock heating and compressive work. This was motivated by the ease with which it could be implemented into the code and our initial tests that indicated it highlighted the tearing, interacting rings but ignored the outer disc. However, while this method is a good indicator for when the rings are interacting, it does not model the precession of an individual ring well and had significant issues from poorly resolved particles. Our second method included a luminosity source at the centre of the simulation and considered the change in flux from this point as the precessing rings passed in front of it. This method has the advantage that it is more closely aligned with the accepted model for type-C QPOs. Additionally, we demonstrated that the PDS generated using this method only represented features due to precessing structures, measured features from the inner regions specifically and showed a strong viewing dependence.
5. Observations of QPOs in nature have minutes or more of data with which to sample the PDS. This is currently not possible with our computational resources, but we were able to generate ~ 10 seconds of tearing disc/precessing rings behaviour. Within this, we observe a handful of precessions of each given ring but note that this behaviour is not yet resolved — although each ring tearing off the outer disc is resolved, the interactions between the rings is not. We are currently conducting simulations with 10^7 particles in an effort to better resolve this and thus the frequencies with which the rings precess. However this will still not yield more than 5 or so precessions before the end of our simulation, and it is not clear that this is enough to generate a meaningful PDS.
6. From the PDS we are able to generate, there are no features we would consider as significant enough to compare to observations of LFQPOs. The potential reasons for this have been highlighted in the previous point.
7. Our use of a post-Newtonian approximation mainly affects the frequencies measured in the inner $\sim 10R_g$. Outside of this region, the approximation is close enough to the

analytically derived expression such that the Lense-Thirring precession is accurate and tearing should not be affected. Additionally, the apsidal precession demonstrated in our inner disc simulations shows that this effect is adequately modelled for our purposes. As HFQPOs are sourced from the innermost regions of the disc where the post-Newtonian approximation is less accurate, any measurement of these from our simulations must be treated carefully.

8. Although observations of QPOs motivate separate mechanisms for type-A, -B and -C QPOs, the simulations shown here generate many precessing structures that all have distinct properties (e.g. the rings are necessarily radially narrow). This tempts the question: are the different types of LFQPOs related to the different precessing structures driven by disc tearing? This remains a question for future work.

The most physically motivated PDS generated in this chapter did not contain any significant features. While this may suggest that tearing discs are not related to LFQPOs, we have discussed the limitations and numerical difficulties of the simulations we have presented. Until these are resolved — in particular, the number of precessions measured in our simulations — tearing discs remain a possible explanation for LFQPOs. In future work we intend to consider the results from our high resolution simulations (§ 5.9) as well as moving to a more physically motivated method of measuring the luminosity. As simulating on physical timescales of minutes continues to be untenable, we will also continue to consider different methods that may be used to generate a PDS that may be compared directly to observations.

Declaration for Chapter 6

Declaration by Candidate

The work presented in § 6.2.1 to § 6.3 and § 6.6 to § 6.9 has been submitted for publication.

I declare that my contribution of the work presented in this chapter involved the following:

The writing of the publication, solving for the mode growth, 3D simulations, adapting and using the 1D code to measure angular momentum transport and analysis of the 3D simulations.

Extent of contribution by candidate: 80%

Co-author names: Daniel Price, Clément Bonnerot and Giuseppe Lodato

Nature of co-author's contribution: Providing the initial conditions for the wide torus simulations, expertise in the analytical description for the development of the instability, 1D code to measure angular momentum transport, comments and feedback during the writing of the paper.

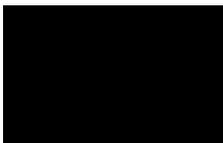
I have renumbered sections, equations and figures of the published paper in order to generate a consistent presentation within the thesis.

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the student's and co-authors' contributions to this work.

Student signature: _____

Date: November 20, 2017

Main Supervisor signature:



Date: November 20, 2017

Chapter 6

Thick accretion tori

We finally consider the evolution of geometrically thick structure around a compact rotating object, beginning by considering a large cloud of turbulent gas that is set to engulf a black hole. We show that when the cloud is inclined to the black hole spin it is possible to generate rings of precessing gas, as in the simulations of inclined accretion discs.

Second, we consider a torus formed during a tidal disruption event. Bonnerot et al. (2016) simulated the evolution of tidal disruption events taking into account apsidal precession with different cooling efficiencies. They confirmed that including the effects of apsidal precession allowed the gas to circularise rapidly, as the head of the stream of disrupted gas intersected the tail of the stream each time it passed through pericentre. These intersections lead to shocks and heating of the gas and, in the case where the gas was not able to cool fast enough, the formation of a torus.

We will demonstrate in this chapter that when this torus is not already accreting, the innermost region of such a remnant is unstable to the Papaloizou-Pringle instability (PPI). This instability generates non-axisymmetric perturbations in the density of the torus. Although accretion of this torus will ultimately be driven by the magnetorotational instability (MRI), the initial magnetic field in the torus is weak because it comes from the disrupted star. Previous work has suggested that in this case it will take ≈ 3 orbits for the MRI to saturate and result in significant accretion (see § 6.7). However, in the torus we are considering the PPI has a comparable growth rate and saturates by accretion through the inner edge. We are thus interested in whether this instability can drive accretion during the narrow window where the MRI is growing but not yet established.

To investigate the connection between the PPI and the initial accretion of tidal disruption remnants, we begin by considering the evolution of a radially narrow torus that is Papaloizou-Pringle unstable and confirm that the PPI is capable of driving angular momentum transport. We consider a circularised version of the tidal disruption remnant by Bonnerot et al. (2016) to investigate whether the PPI is capable of causing significant accretion before the MRI is likely to be established in these geometrically thick tori.

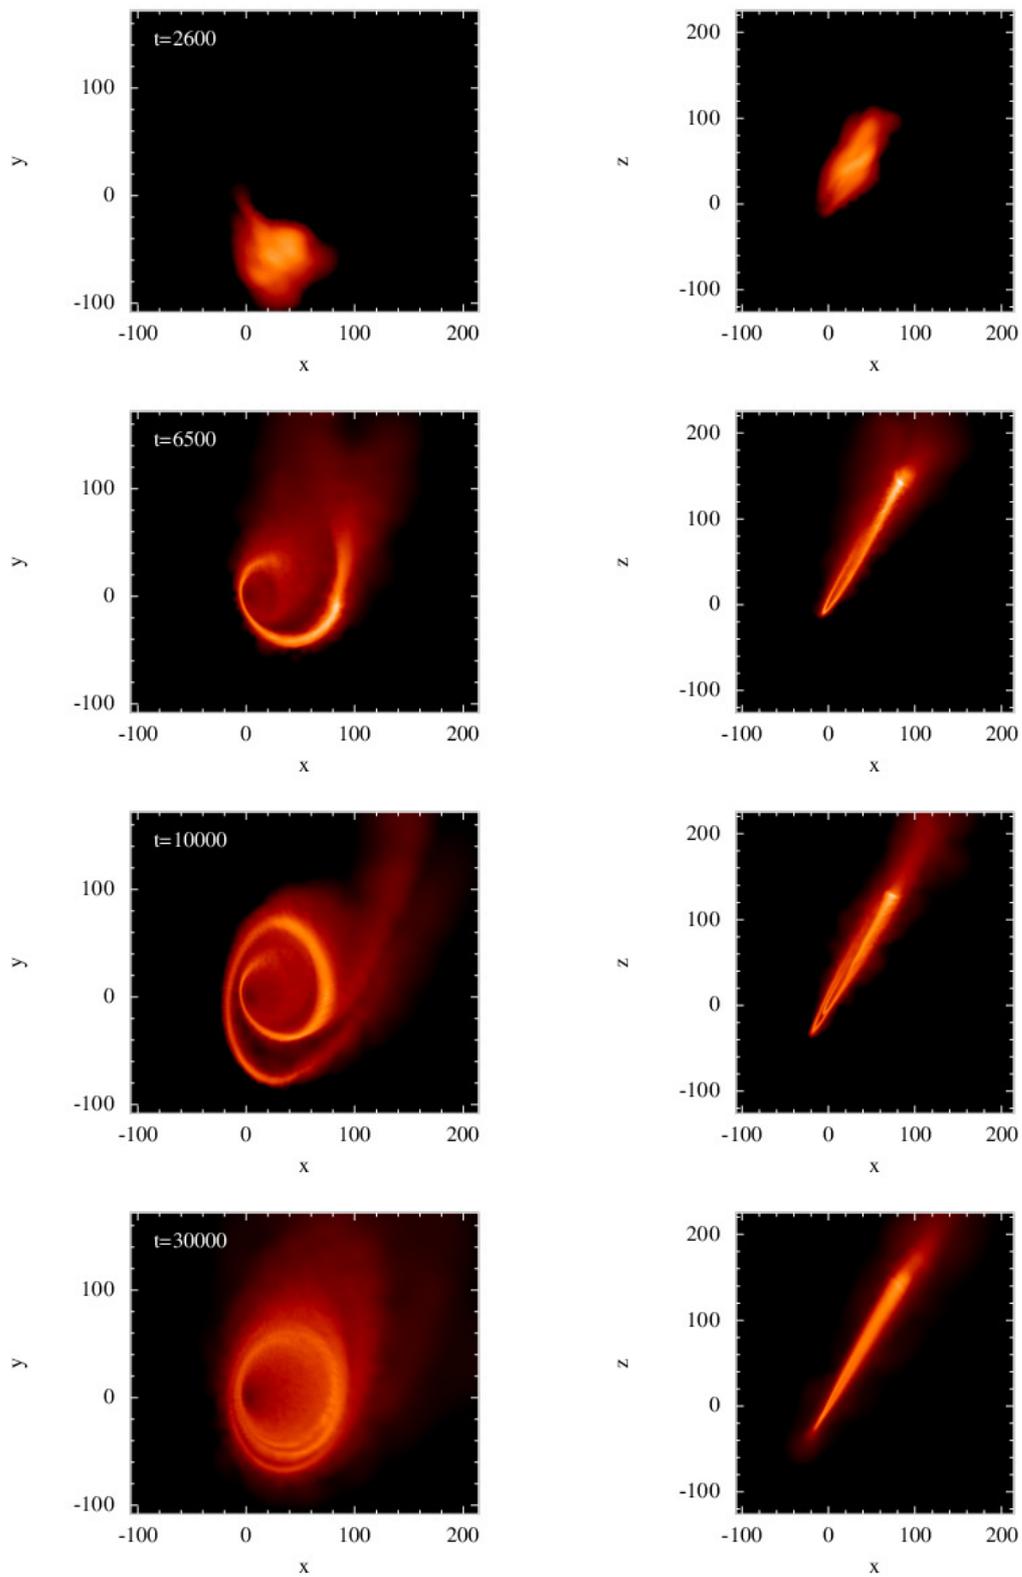


FIGURE 6.1: Non-rotating black hole engulfing a cloud inclined at 60° . The black hole is located at the origin and time is in units at R_g (i.e. an orbit at $R = 1$ takes 2π), distances are in code units. As this scenario is equivalent to approaching a rotating black hole at zero-inclination, direct comparison with Figure 6.2 demonstrates the effect of inclining the black hole spin.

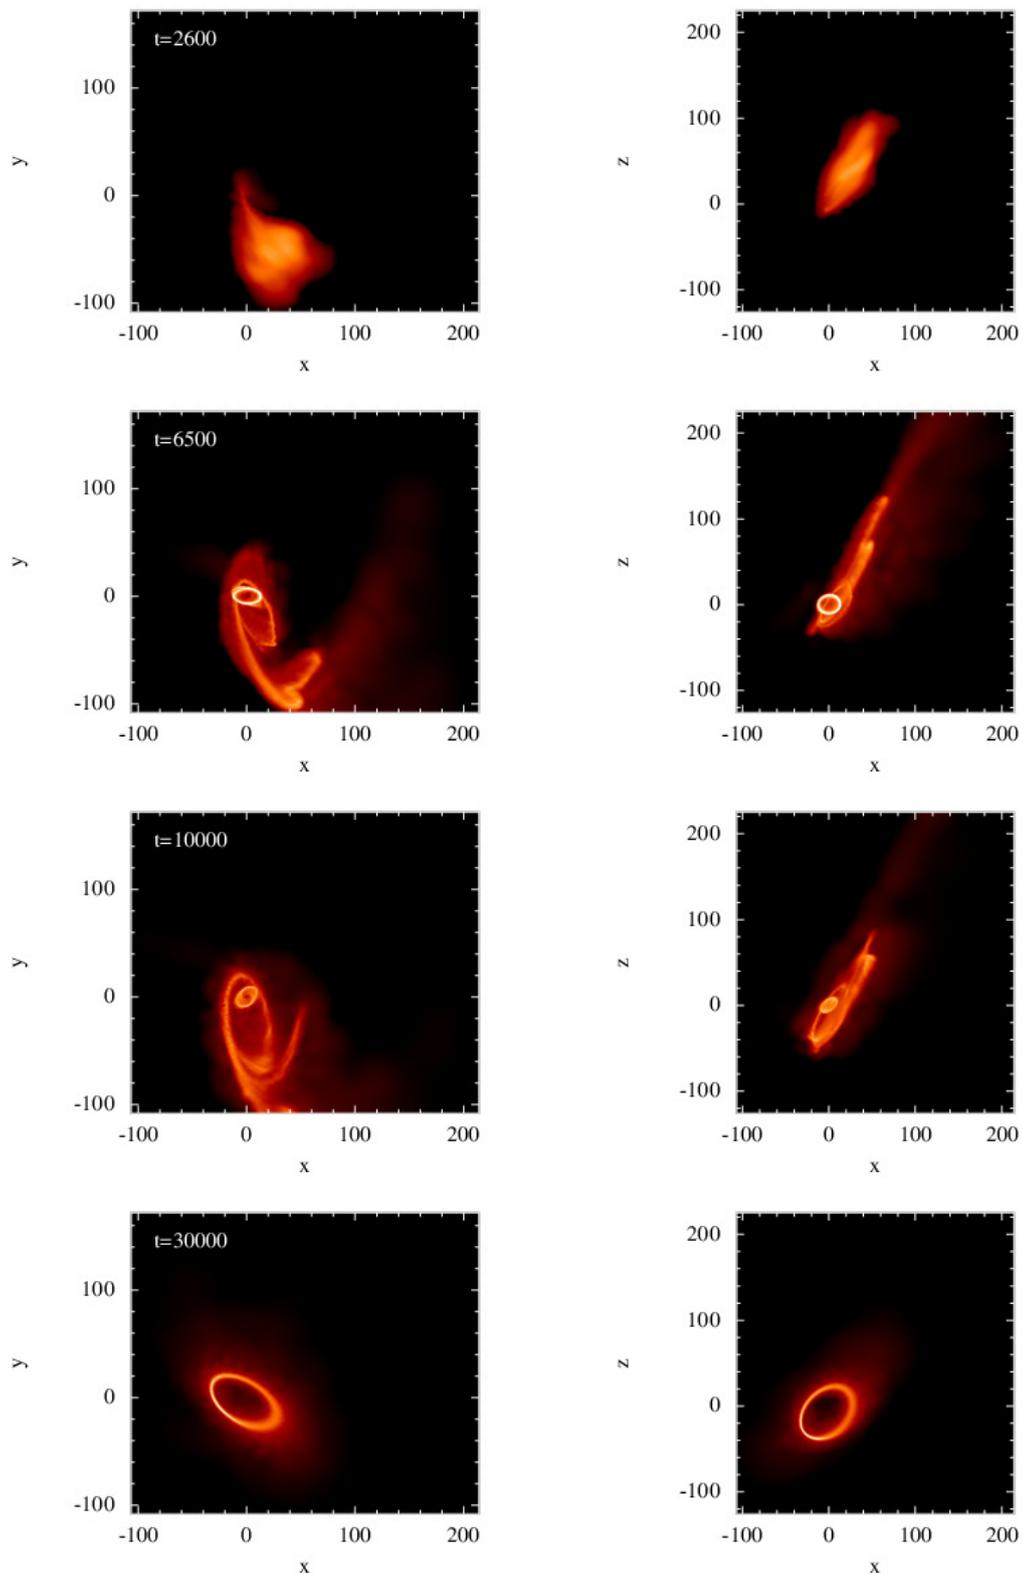


FIGURE 6.2: Same as Figure 6.1 except the black hole here is rotating. Inclining the black hole spin demonstrates the formation of rings, which should lead to the development of a truncated inner disc. This scenario produces more shocks such that the gas circularises faster, in agreement with Bonnerot et al. (2016).

6.1 The interaction of a cloud with a rotating black hole

To begin considering non-disc structures around rotating black holes, we investigated the effect of Lense-Thirring precession on a cloud of gas. These simulations were based on work by Lucas et al. (2013) who simulated misaligned streamers of stars around Sgr A*. Lucas et al. (2013) conducted SPH simulations of an elliptical cloud engulfing a rotating black hole, finding a possible source for the two misaligned rings of stars that surround the black hole at the centre of the galaxy.

We start with a similar setup to Lucas et al. (2013), but move the initial cloud closer to the black hole so that the effects of nodal precession are more pronounced. We thus set the pericentre of the orbit at $R_p = 10R_g$ and shrink the cloud to have a width of $20R_g$ and length $50R_g$. The gas in the cloud is modelled as adiabatic with no shock heating or compressive work (i.e. a co-moving isothermal frame) and gas is initialised as turbulent using the turbulent velocity fields available with PHANTOM. As in Lucas et al. (2013) we use 3×10^6 particles. In both cases we incline the cloud to the black hole spin, but in the non-rotating case with $a = 0$ the evolution would be equivalent at any inclination.

Figures 6.1 and 6.2 shows the evolution of these clouds. In the case of the non-rotating black hole the cloud circularises and forms a disc quickly. However, in the case with a rotating black hole there is evidence of rings in the second and third panels. These rings are formed as the incoming material precesses with a common rate, but their interactions are the same as in the tearing disc case explored earlier. Angular momentum is cancelled more efficiently than in the non-rotating case, resulting in a comparatively smaller final disc. Additionally, this disc is not aligned with the original inclination of the cloud. This comparison confirms that even in the case of accreting material that has not formed a disc, precessing rings of gas can still be formed from misaligned material. The interaction of these rings ultimately leads to the formation of an inner, aligned disc that feeds the black hole but this occurs on timescales much faster than in the aligned case.

The difficulty with these simulations is that a cloud like this would have been disrupted and formed a disc well before it gets to such a small radius. As misaligned gas moves through the tearing region, our simulations in Chapter 5 (e.g. Figure 5.11) demonstrate that a single, radially wide ring of gas will rip off, cascading into many smaller interacting rings. This behaviour is not demonstrated in these simulations because of our choice of initial conditions.

6.2 The Papaloizou-Pringle instability

We instead consider the simulation from Bonnerot et al. (2016) as a physically motivated scenario which generates a geometrically thick torus. The cross section from the remnant in this simulation was shown in Figure 1.21. The thick, puffy nature of this cross section is due to the combination of shock heating and numerical viscosity. As the shocks heat the

gas it cannot cool efficiently, increasing the aspect ratio (H/R) and puffing up the torus. Additionally, the numerical viscosity that accommodates these shocks leads to spreading, particularly at the edges of the torus. Because the numerical viscosity is resolution dependent, at higher resolution we expect the torus to have spread less and be more compact (while the aspect ratio should remain consistent). Although this is a purely hydrodynamic simulation, the weak initial magnetic field of the star means that the magneto-rotational instability (MRI) is not likely to be established yet such that the torus is not expected to be accreting.

The torus described above already meets some of the criteria for the Papaloizou-Pringle instability (PPI), but the final criteria is that the slope of the specific angular momentum profile, $l \propto r^q$, be shallow such that $q < 0.25$. We measure the slope of the specific angular momentum profile in the innermost 0.5 AU of the remnant from Bonnerot et al. (2016), finding that $0 \lesssim q \lesssim 0.1$. Thus the innermost region of the torus is Papaloizou-Pringle unstable in the case when it is not accreting and more compact than the result from Bonnerot et al. (2016) (as would be expected at higher resolution).

We recall for a given mode m , the growth rate ω of the PPI satisfies (Blaes and Glatzel, 1986)

$$\frac{(\omega + m\Omega(R_-))^2 + mg(R_-)/R_-}{(\omega + m\Omega(R_+))^2 + mg(R_+)/R_+} = \left(\frac{R_+}{R_-}\right)^{2m} \frac{(\omega - m\Omega(R_-))^2 + mg(R_-)/R_-}{(\omega - m\Omega(R_+))^2 + mg(R_+)/R_+}, \quad (6.1)$$

where R_- is the inner edge, R_+ the outer edge and $\Omega(R)$ and $g(R)$ have particular descriptions given in § 1.5.3. In the case of a radially thick torus with say, $R_- = 0.8R_0$ and $R_+ = 1.6R_0$, the above equation gives that the $m = 1$ grows fastest with $Im(\omega) \approx 0.5\Omega_0$.

To understand how angular momentum may be transported by the PPI we initially consider the evolution of a radially narrow torus for two reasons; for a given number of particles a radially narrow torus will have a higher resolution and we can use a radially narrow torus as a ring spreading test (recall Figure 1.1). We set $R_- = 0.9R_0$ and $R_+ = 1.1R_0$ for this thin torus, with an almost circular cross section (rather than the stretched out cross section in Figure 1.21). Figure 6.4 shows the solutions to Equation 6.1 for different modes m for this torus, suggesting that the fastest growing mode for our thin torus simulation is the $m = 3$ mode with a growth rate of $Im(\omega) \approx 0.3$ per orbit.

6.2.1 Growth of the PPI in a thin torus

We start by studying the development of non-axisymmetric perturbations in a radially thin, vertically thick torus. To compare as closely as possible to the description assumed to derive Equation 6.1, we consider a radially narrow torus with $G = M = 1$, $R_0 = 1.0$, $R_+ = 1.1$ and $R_- = 0.9$ (orbital times below are specified at R_0). The left-most panels of Figure 6.3 show the initial torus structure. The density and pressure of the torus are assigned by assuming a polytropic equation of state $P = A\rho^\gamma$ and (Papaloizou and Pringle,

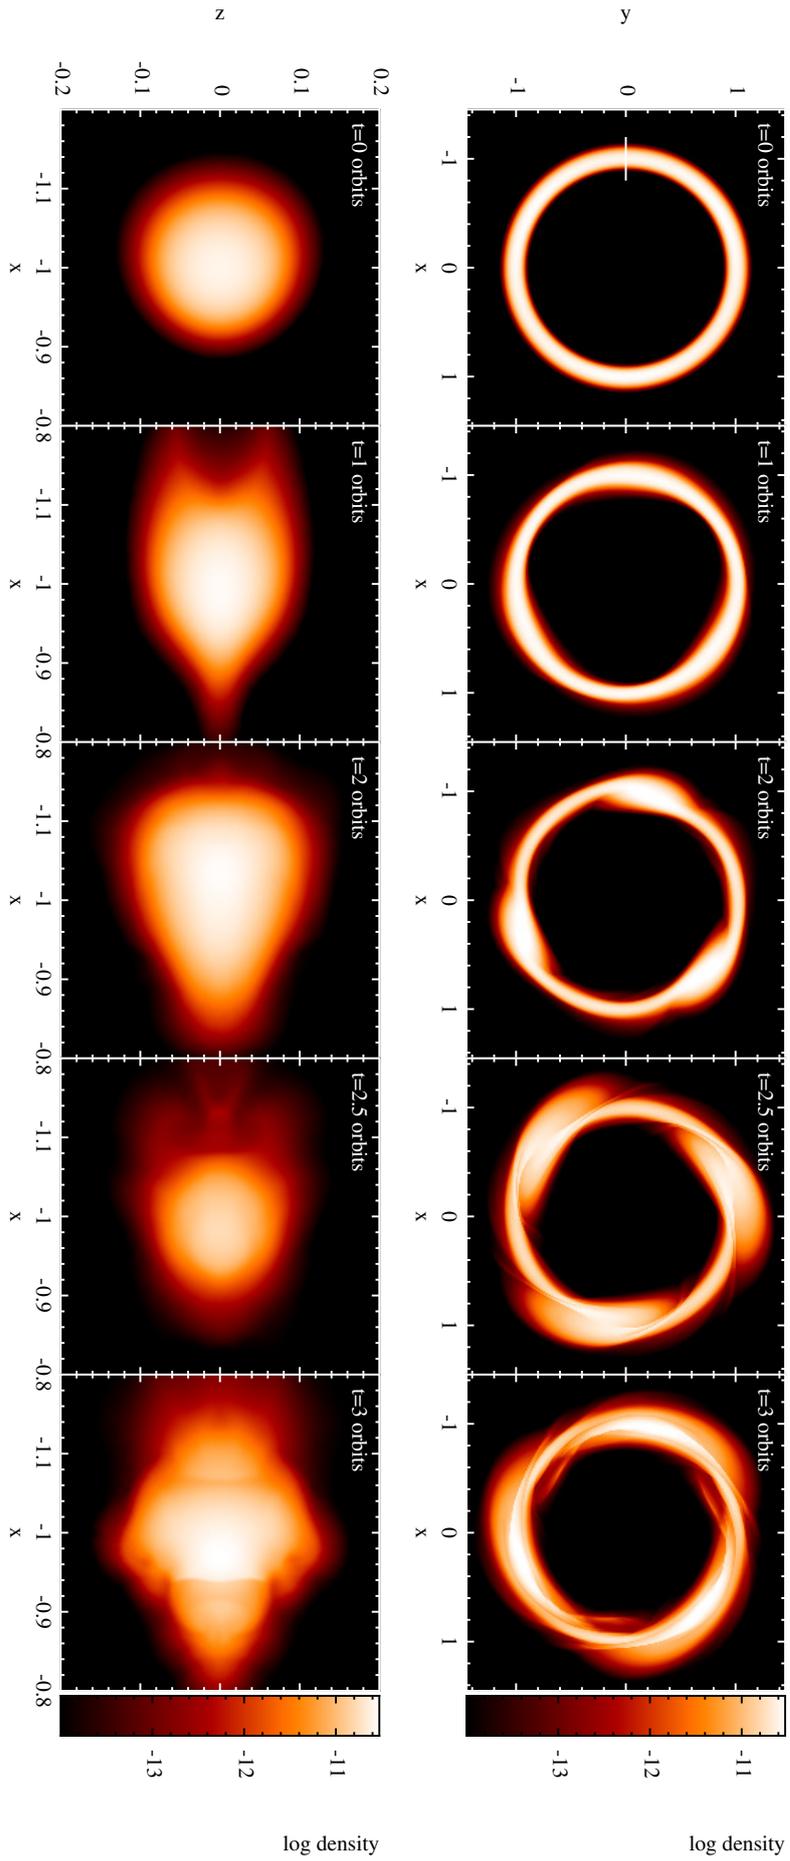


FIGURE 6.3: The Papaloizou-Pringle instability in action, showing the development of non-axisymmetric shocks (centre panel) that lead to radial spreading of an initially radially thin torus (bottom panels). Upper panel shows the x - y cross section and lower the x - z , with the white line in the top panel showing the location of the x - z cross section. Evolution is from a thin torus simulation seeded with an $m = 3$ mode.

1984)

$$\frac{P}{\rho} = \frac{GM}{(n+1)R_0} \left[\frac{R_0}{r} - \frac{1}{2} \left(\frac{R_0}{r \sin \theta} \right)^2 - \frac{1}{2d} \right], \quad (6.2)$$

where the maximum density

$$\rho_{\max} = \left[\frac{GM}{(n+1)AR_0} \left(\frac{d-1}{2d} \right) \right]^n, \quad (6.3)$$

is used to specify A . Here r is the radius in spherical coordinates, R_0 describes the cylindrical radius of maximum density, θ is the angle that describes the height out of the plane and $n = (\gamma - 1)^{-1}$ is the polytropic index. The factor $d = (R_+ + R_-)/(2R_0)$ determines the profile of the cross section, where values close to unity correspond to a circle. For our radially thin torus we specify an almost circular cross section given by $d = 1.01$. In code units, we chose a maximum density $\rho_{\max} = 2.5 \times 10^{-9}$, but the evolution of the torus is independent of this choice. We repeated the simulation using 3.0×10^5 , 2.3×10^6 , 1.6×10^7 and 1.25×10^8 particles. At the lowest resolution the vertical thickness is resolved by approximately two smoothing lengths and in the highest resolution case by ten.

Following Zurek and Benz (1986), the particles in the torus were initially given zero velocity and relaxed in an effective potential that accounted for both the gravitational potential and the pressure forces due to the initial specific angular momentum profile. The torus was allowed to relax in this potential until the potential energy stopped oscillating, corresponding to 5 orbits at R_0 . Subsequently the particles were given orbital velocities and then seeded with the fastest growing mode. Figure 6.4 shows the growth rates from Equation 6.1 for this choice of R_- and R_+ , from which we see that $m = 3$ is the fastest growing mode. Seeding was achieved with a small azimuthal perturbation in density ρ , given by $\rho = \rho_0 [1 + B \cos(m\phi)]$, where ρ_0 is the original density, B is the amplitude of the perturbation and ϕ is the azimuthal angle. To achieve this the particles were shifted in position by $\delta\phi = -B \sin(m\phi_0)/2$, where $B = 0.05$ and ϕ_0 is the original azimuthal angle (e.g. Price and Bate, 2007). We use a Keplerian gravitational potential with an accretion boundary of $R_{\text{acc}} = 0.1$ for the lower two resolutions and $R_{\text{acc}} = 0.2$ for the highest resolution simulation, within which particles are removed from the simulation. The central object has $M = 1$ in code units.

The time evolution shown in Figure 6.3 shows the formation of three over-densities corresponding to the $m = 3$ mode, which co-rotate with the fluid at the orbital velocity at R_0 . After 2.5 orbits (fourth panel from left) these over-densities reach their maximum radial width, with evidence of strong shocks at high resolution. The PPI saturates (fourth panel), and the over-densities become less prominent (fifth panel) as the torus continues to evolve and spread radially. The shocks generated by this instability are particularly clear in the x - z cross-section (bottom row) and are confirmed by plotting the divergence of the velocity field.

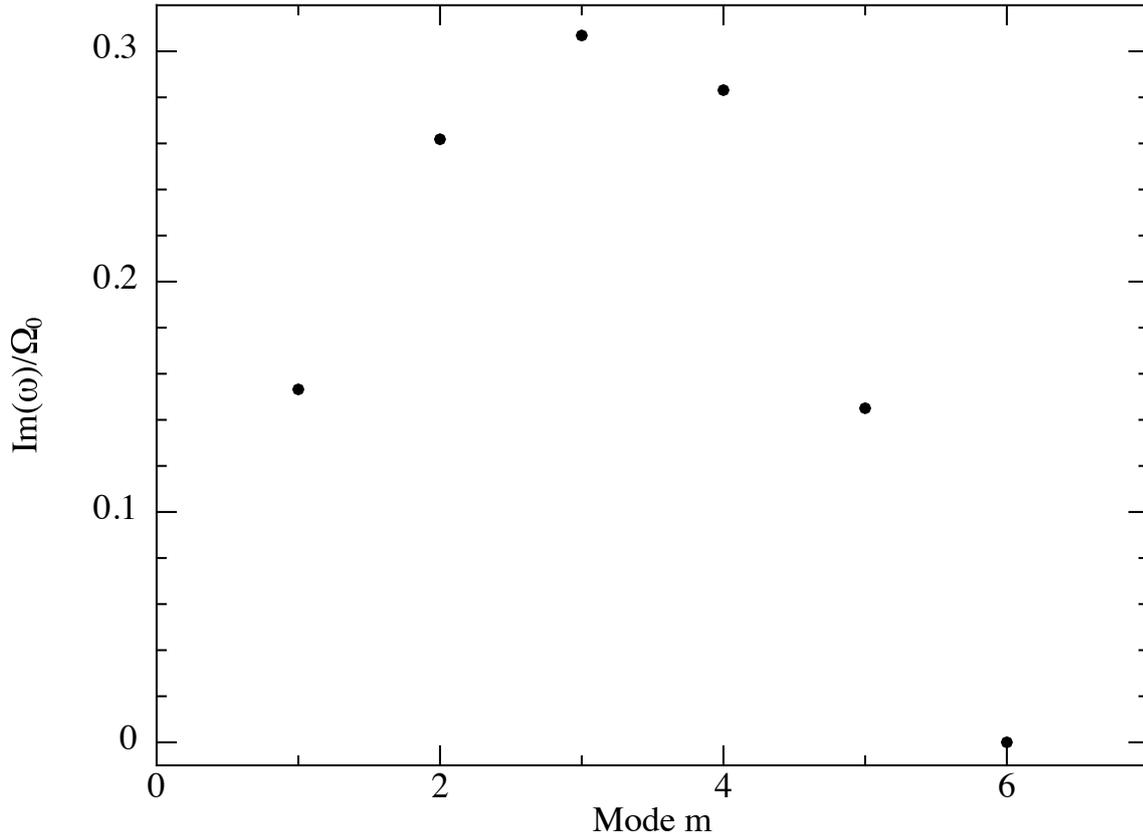


FIGURE 6.4: Growth rate calculated from Equation 6.1 for a few azimuthal wavenumbers m when $R_- = 0.9R_0$, $R_+ = 1.1R_0$. The fastest growing mode ($Im(\omega) \approx 0.3\Omega_0$ for $m = 3$) is used to seed the radially narrow torus in Figure 6.3.

To confirm that the growth of the density perturbations in our simulation matches Equation 6.1, we consider particles in the x - y cross-section, with $-0.05 < z < 0.05$. We compute the average density, ρ_{mean} , from all the particles in this ring. The properties of the particles in the torus are then radially averaged in a method analogous to the azimuthal averaging described in Lodato and Price (2010), where we divide the torus into $N = 45$ azimuthal bins such that each bin represents $\Delta\phi = (360/N)^\circ = 9^\circ$. The deviation from the mean density in each bin j as a function of azimuthal angle is calculated with $(\rho_j - \rho_{\text{mean}})/\rho_{\text{mean}}$. We measure the density variation at each timestep and identify the maximum as the location of the over-density — as these co-rotate with the fluid, this occurs at a different azimuthal angle at each timestep. Figure 6.5 shows the maximum density variation as a function of time for different numerical resolutions (see legend) compared to the expected growth rate (purple line). We estimate the average growth rate using a least squares fit between 0.1 and 3.0 orbits, measuring $Im(\omega)/\Omega_0 = 0.5 \pm 0.19$ from the simulations (with the uncertainty derived from the lowest resolution case). Taking into account the assumptions of compressibility and cylindrical flow used to derive Equation 6.1, we consider our measured growth rate to be consistent with the analytical prediction of $Im(\omega)/\Omega_0 \approx 0.3$. We additionally visually confirm that the motion of the over-densities is on the orbital timescale, which is consistent with the analytical prediction of $Re(\omega)/\Omega_0 = m$. The ‘dip’ observed around 1 orbit is caused by mixing of unstable modes with stable oscillations of

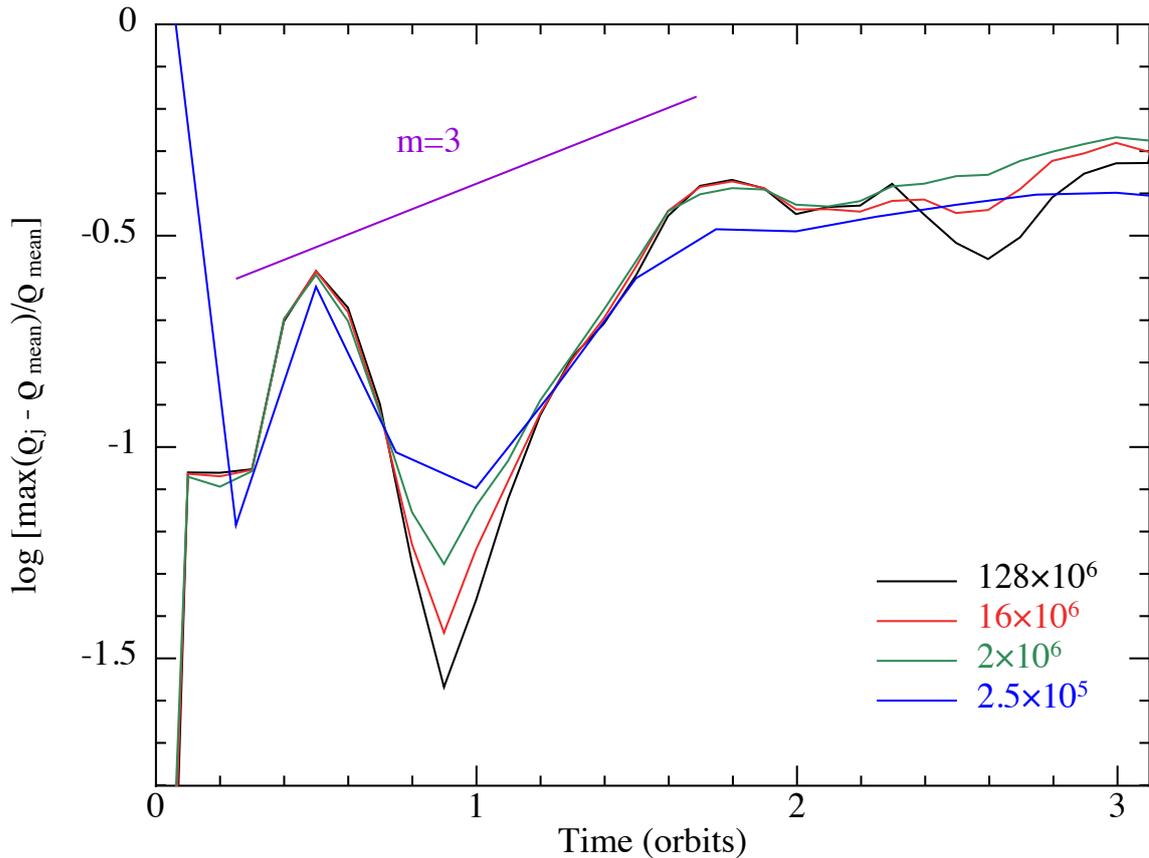


FIGURE 6.5: Resolution study showing the amplitude of the non-axisymmetric density perturbation as a function of time from the thin torus simulation. Purple line shows the expected mode growth from Equation 6.1 for the fastest growing ($m = 3$) mode. The instability grows on the orbital timescale, saturating after approximately 3 orbits. Comparing the average growth in the first two orbits we find numerical convergence in the growth rates.

the torus, and also affects our time averaged growth rate measurement. The amplitude of the perturbations saturates after 2.5 orbits in agreement with Figure 6.3. The linear growth phase (first two orbits) is converged for even moderately low resolution, but high resolution is needed for the saturation phase due to the interacting shocks.

6.3 Does the Papaloizou-Pringle instability lead to angular momentum transport?

The radial spreading that occurs in Figure 6.3 suggests that the PPI may be capable of generating angular momentum transport. Figures 6.6 and 6.7 quantify this, showing the surface density $\Sigma(R)$ as a function of radius in the thin ring and the measured effective α parameter.

To measure the viscosity, ν , in the thin torus, which has contributions from both the artificial viscosity in the code and any viscosity generated by the PPI, we use the same technique utilised in Lodato and Price (2010). That is, we match the surface density

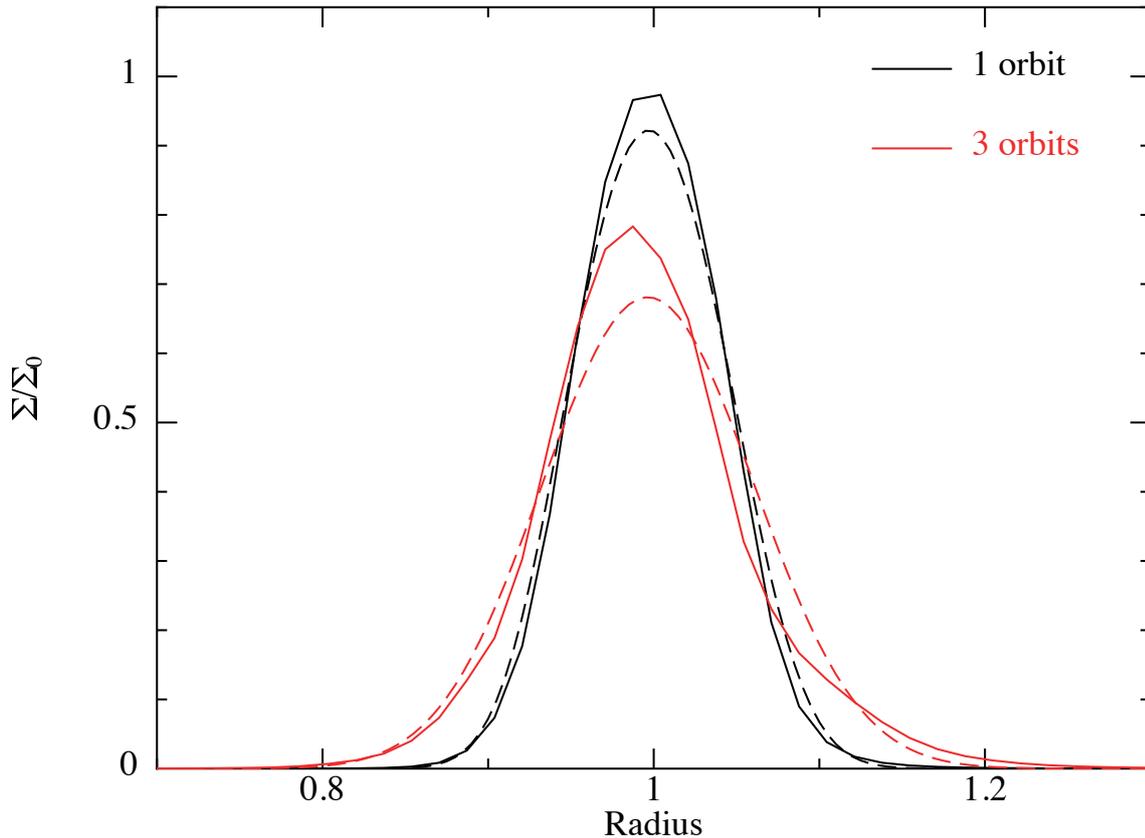


FIGURE 6.6: Evolution of the surface density from the thin ring simulation (dashed line), matched to the corresponding evolution of Equation 6.4 (solid line) at the same time. Spreading suggests angular momentum transport may be generated by the PPI in a radially thin torus, measured in Figure 6.7.

evolution in the 3D simulations to the solution of the 1D disc diffusion equation (Pringle, 1992)

$$\frac{\partial \Sigma}{\partial t} = \frac{3}{R} \frac{\partial}{\partial R} \left[R^{1/2} \frac{\partial}{\partial R} \left(\nu \Sigma R^{1/2} \right) \right]. \quad (6.4)$$

We measure the ‘effective α ’ using a root-finding algorithm, where we minimise the difference between the evolution of Σ (modelled by Equation 6.4) of the 1D code and our simulation at corresponding times using a finite difference method as outlined in Lodato and Price (2010). We follow the process outlined in Section 4.2.1 of Lodato and Price (2010) in order to quantify the difference between the 1D code and the simulation at a given time step, finding the best fit between surface density in the 1D code and the surface density in the 3D code, using bins within a radial range of ± 0.05 around the maximum in the radial surface density profiles.

We set $\Sigma = 0$ as the inner boundary condition for the 1D code at $R = 0.5$ as in Lodato and Price (2010) and set the same condition far away from the outer edge of the torus at $R = 5$. As the Σ profile of the relaxed ring cannot be described by a simple power law, the initial Σ profile was interpolated directly from the thin torus simulation. The aspect ratio H/R (where H is measured from the standard deviation of particle position in the

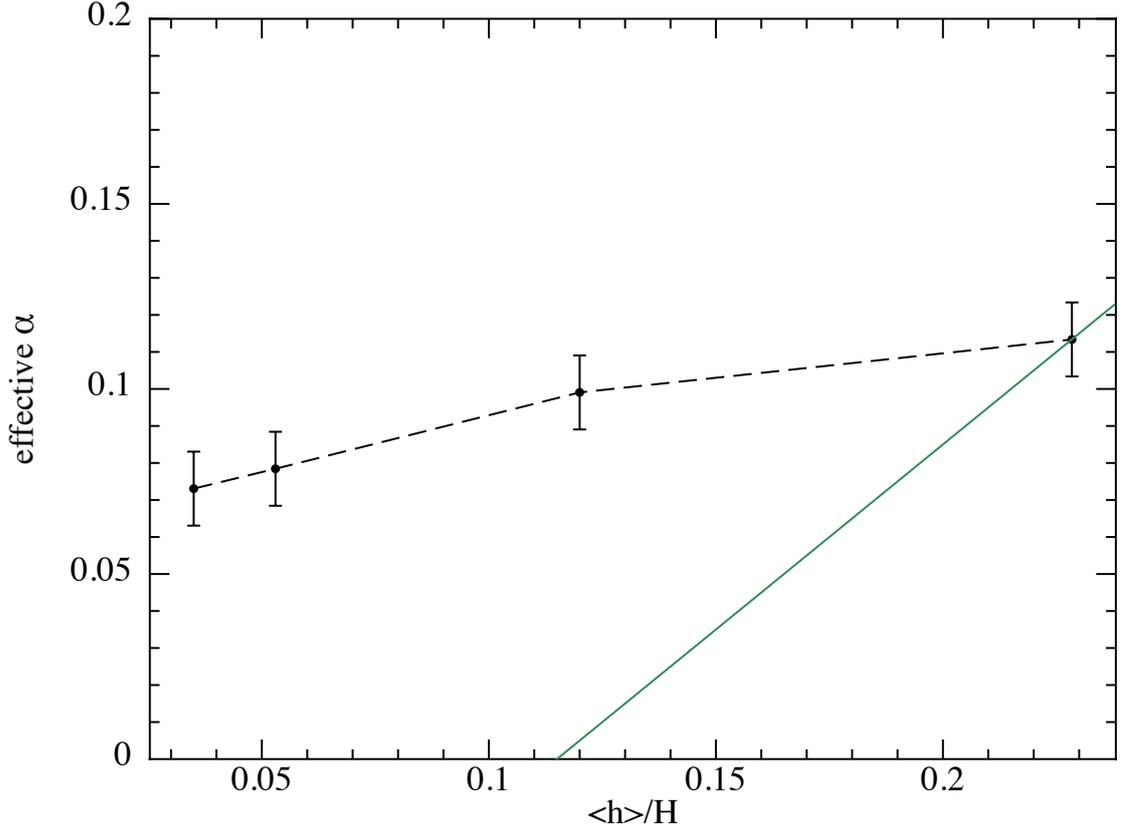


FIGURE 6.7: The ‘effective α ’ measured from the ring spreading in Figure 6.6 at our four highest numerical resolutions (increasing from right to left). We find a spreading which is independent of numerical viscosity, with a Shakura-Sunyaev $\alpha \approx 0.07$. Here the green line indicates the expected first-order scaling of the artificial viscosity terms.

z direction) varies between 0.04 – 0.045 during the first 4 orbits, so we adopt an average value of $H/R = 0.0425$. Figure 6.6 shows the comparison between the 1D code and the simulation after three orbits of our highest resolution simulation.

Unlike Lodato and Price (2010), we expect α to vary with time as the PPI develops. Hence we measure α_n at every tenth of an orbit, t_n , which represents the average viscosity prior to that time. As the PPI develops the effective viscosity remains roughly constant for 2 orbits, where it steeply increases and reaches a maximum by 2.5 orbits — when shocks are strongest in the simulation. The effective viscosity then decreases, returning to a constant value by 3 orbits. Figure 6.7 shows the effective α after the PPI has developed and saturated (i.e. at three orbits) as a function of resolution measured at R_0 . This effective α has contributions from both the resolution dependent numerical viscosity and the physical viscosity generated by the PPI. Figure 6.7 demonstrates that the viscosity in the torus is approximately independent of vertical resolution. As the viscosity derived from the higher resolution simulations has a weaker resolution dependence than expected from artificial viscosity only (green line in Figure 6.7), we conclude there is a physical origin for this viscosity that is related to the PPI with an effective $\alpha \approx 0.07$. These simulations also indicate that we require the ratio of the shell-averaged smoothing length to scale height

$\langle h \rangle / H \lesssim 0.1$ in the majority of the torus to guarantee that the rate of spreading in the torus does not change significantly with resolution.

6.4 Redistribution of the specific angular momentum

Figure 6.8 demonstrates the redistribution of specific angular momentum throughout the course of our thin torus simulation. We follow the method presented in Zurek and Benz (1986), where the specific angular momentum is parameterised by $l \propto r^q$. The specific angular momentum in the torus is measured as a function of the radius by Equation 2.60 for each time-step. We then fit the $l(r)$ profile using a least squares fit to find the q value at each time-step, with uncertainty from the fitting procedure. We find that the q value saturates at $q \approx 0.25$, shown by the purple dashed line in Figure 6.8. This is in agreement with Zurek and Benz (1986) (presented in Figure 1.22), despite their wider torus and our modern implementation of SPH.

6.5 Does the Papaloizou-Pringle instability lead to ballistic accretion?

We consider whether the PPI is able to drive angular momentum transport through ballistic accretion. This would occur if the motions of the gas become so eccentric that the gas is forced onto an orbit where it accretes straight into the black hole, so we begin by measuring the eccentricity of the gas \mathbf{e} , given by

$$\mathbf{e} = \frac{\mathbf{v} \times (\mathbf{r} \times \mathbf{v})}{GM} - \frac{\mathbf{r}}{|\mathbf{r}|}. \quad (6.5)$$

The magnitude of the eccentricity can be equivalently written as

$$|\mathbf{e}| = \sqrt{1 + \frac{2E|L|^2}{G^2M^2}}, \quad (6.6)$$

where E is the sum of the kinetic and potential energy of the orbit and L is the angular momentum. Figure 6.9 shows the eccentricity measured with either method for our high resolution simulation for the first three orbits. In the case that orbits are eccentric we expect $|\mathbf{e}| = 0$, however this is not the case even for our initial condition (black line). This is because the orbits in a torus, while being circular, are only Keplerian at R_0 (at $R = 1$ here). The deviation from Keplerian is larger the further from R_0 , explaining our initial condition where $|\mathbf{e}| \approx 0$ at R_0 but $|\mathbf{e}| \gtrsim 0.1$ at the inner and outer edges. This application of Equation 6.6 is thus misleading, because when applied to a torus it cannot directly show where orbits are circular. However, we know that the orbits in our initial condition are circular (we set them as such), so we take the profile at $t = 0$ to be at least indicative

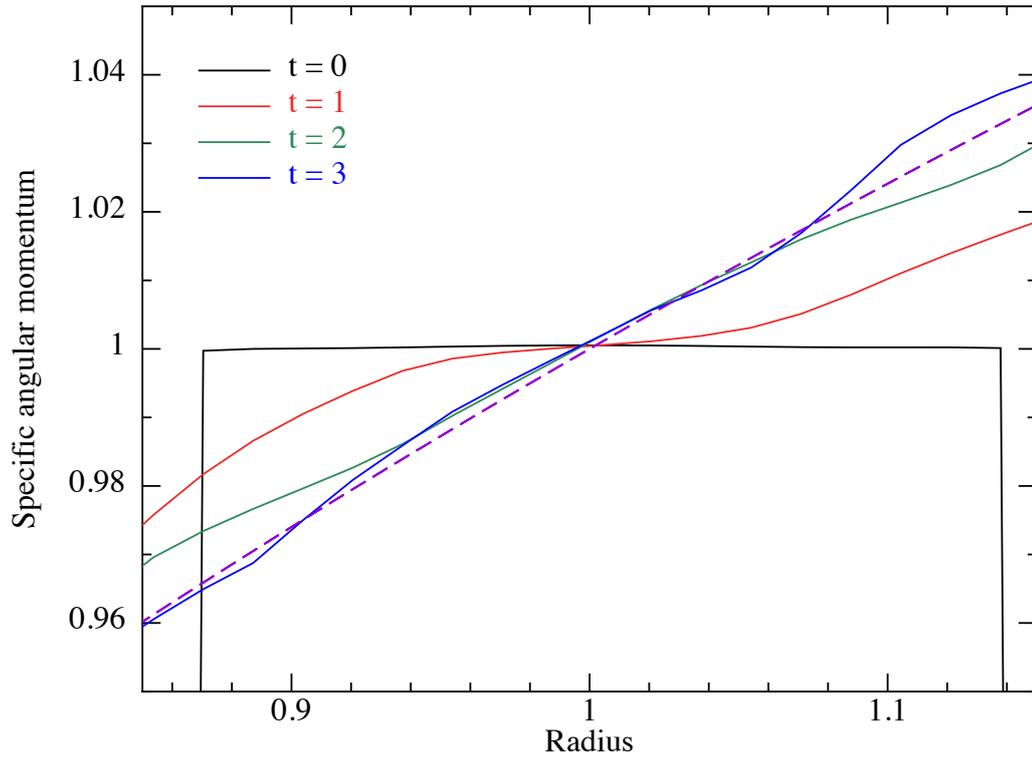


FIGURE 6.8: Evolution of the specific angular momentum profile (where $l \propto r^q$) for the highest resolution radially narrow tori simulation. Here the purple dashed line represents a saturated profile with $q = 0.25$. Following the method of Zurek and Benz (1986) we confirm that saturation of the PPI is equivalently achieved by redistribution of the specific angular momentum in the torus.

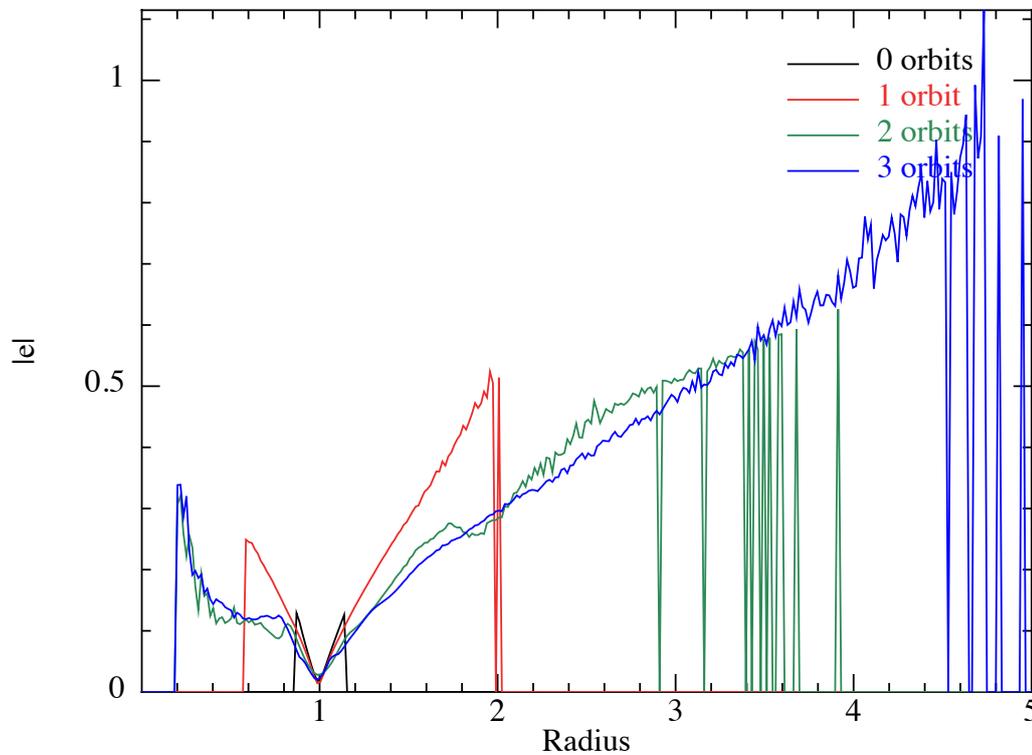


FIGURE 6.9: $|e|$ measured from the high resolution radially narrow torus simulation. This method does not work for non-Keplerian orbits, as the circular orbits in the initial condition are not represented by $|e| = 0$.

of circular orbits. For $R > R_0$, at later times the torus spreads to much larger radii but effectively maintains the initial profile of increasing $|e|$ to the outer torus edge. For $R < R_0$ however, the profile deviates between $0.4 < R < 0.8$, suggesting that material in this region may no longer be on circular orbits. Consideration of the final panel in Figure 6.3 confirms this; at the inner edge of the torus there are shock structures reaching inwards from the torus edge. While this suggests that there is eccentricity in the innermost part of the torus, this does not confirm ballistic accretion is occurring.

In an attempt to confirm that the torus is accreting ballistically we investigate the particles that have been accreted. We identify the ≈ 1500 particles that have accreted by three orbits in our high resolution simulation and track them backwards to find their initial positions. In the initial condition these particles all appear to hug the inner edge of the torus, consistent with accretion driven by viscosity (i.e. the particles at smaller radii in are accreted first). In the context of the interpretation from Figure 6.9, this suggests that material at the inner edge is being caught up in shocks, made slightly eccentric and accreted.

We finally consider the orbits of those individual particles that are accreted. Tracing them throughout the full evolution of the torus shows no eccentric motion visible in their orbits. This suggests that if the gas is driven to be eccentric by the PPI, it is not a large effect when compared to the shock structures that are generated. Interestingly, tracing the evolution of some randomly chosen particles demonstrates that particles are able to move radially inwards but then change their motion (still in the same orbital direction) and begin moving outwards. Such motion and the shocks found demonstrates that the PPI is driving complex mixing of the gas, which is ultimately what redistributes the angular momenta. We thus conclude that although the PPI does drive some eccentric motion at the inner edge, this is a small effect and ballistic accretion is not the main accretion mechanism.

6.6 Are TDE remnants unstable to the PPI?

Finally, we consider the $1M_\odot$ remnant of a TDE around a non-rotating $10^6 M_\odot$ black hole, calculated by Bonnerot et al. (2016). This TDE has a penetration factor of $\beta = 5$ where β is the ratio of the tidal to pericentre radii, an eccentricity of $e = 0.8$ and the final evolution after eight orbital periods is shown in their figure 8. We measure the specific angular momentum profile, $l(r) \propto r^q$ inside $R \sim 0.5$ AU using the method outlined in Zurek and Benz (Figure 4, 1986) and find that $0 \lesssim q \lesssim 0.1$. Although the torus formed from the TDE remnant is not yet axisymmetric, the specific angular momentum distribution confirms that if the innermost material circularised it would be susceptible to the PPI.

We construct a circularised version of the remnant by assuming a torus with a similar cross section and specific angular momentum profile. Similar to our thin torus, we set the particles initially on concentric shells around $R_0 = 0.5$ AU, using $d = 1.15$ and with the number of shells and particles per shell dependent on the resolution (we used 1×10^5 , 1×10^6

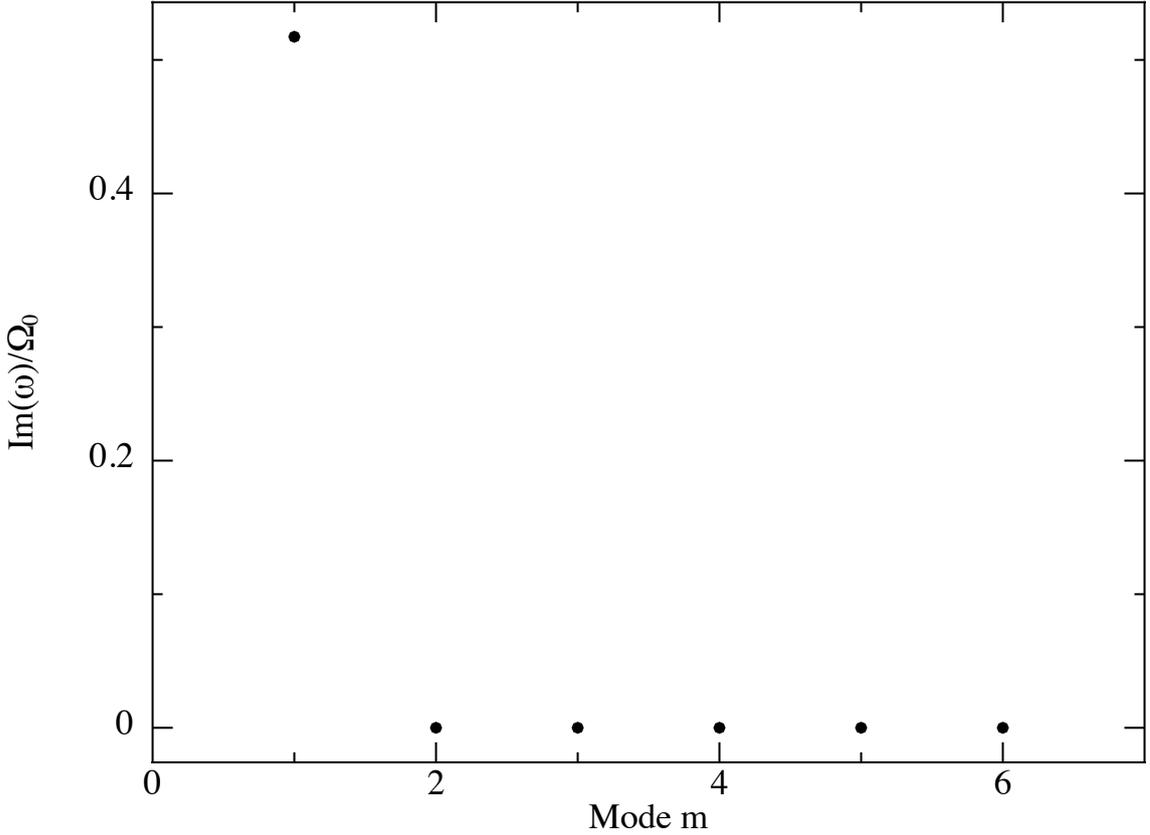


FIGURE 6.10: Growth rate calculated from Equation 6.1 as a function of azimuthal wavenumbers m for the tidal remnant torus simulation. In this case $m = 1$ is the fastest and only unstable mode, which is seen in Figure 6.11.

and 1×10^7 particles, respectively). This torus was relaxed into hydrostatic equilibrium using the relaxing potential as previously, until the potential energy ceased oscillating.

We then set the torus in orbit in a Keplerian potential with a $10^6 M_\odot$ central object and added an $m = 1$ density perturbation using the method described in § 6.2.1 (with amplitude $B = 0.05$). While the $m = 1$ mode is the fastest (and only) growing mode for a radially wide torus (see Figure 6.10), a non-axisymmetric density structure with this mode is already present in the remnant produced by Bonnerot et al. (2016) (their Figure 5). Additionally, if simulated with no added perturbation we find that the $m = 1$ mode develops. Again we use a Keplerian potential, but here with an accretion boundary at $R_{\text{acc}} = 0.1$ AU.

In comparison with the torus formed by Bonnerot et al. (2016) our torus is more compact (top panels, Figure 6.11). Although the original simulation by Bonnerot et al. (2016) has already spread close to the innermost stable orbit of the black hole, the rate of spreading is enhanced by the artificial viscosity included in their simulation. Taking this into account, a more compact torus than they have produced is expected when simulated with a higher resolution, justifying our more compact version. The cross-sectional extent of our torus is also constrained by the torus being in hydrostatic equilibrium. Rather than precisely matching the specific angular momentum measured from the remnant, we adopt constant

specific angular momentum (i.e. $q = 0$) — both are unstable to the PPI (Papaloizou and Pringle, 1984; Zurek and Benz, 1986).

Figure 6.11 shows the time evolution of the density cross section of this torus. Within the first three orbits (measured at R_0) an asymmetry develops but is not yet accompanied by appreciable spreading. The instability continues to grow, and by five orbits a shock has developed that extends from the inner to outer edges. By this stage the outer radius has spread by $\approx 10\%$ at the location of the over-density. The PPI continues to grow and the torus reaches its maximum size by six orbits, when the PPI visually saturates. By seven orbits both the strength of the shock and the asymmetry of the outer edge decrease, although the over-density remains (the partial remnants may be seen in the lower panels of Figure 6.11). The x - z plane cross section (right hand side) also shows the effects of the shocks at the outer edge of the torus. While the torus has not spread much in the z direction, it has doubled its radial extent (although the inner edge is constrained by our choice of accretion radius). For a TDE remnant that does circularise, this simulation suggests that the PPI develops within a few orbits, creating over-densities that remain even after the instability has saturated.

Figure 6.12 shows the mass that falls within the accretion boundary of the high resolution TDE simulation over ~ 20 orbits scaled with the black hole and star parameters specified in Bonnerot et al. (2016). Within the first five orbits this leads to a mass accretion rate of $\sim 100M_\odot/\text{yr}$ — this is well before the MRI is expected to be able to drive significant mass accretion (however, see § 6.7).

With the average mass accretion rate while the PPI is active but not yet saturated (e.g. around 5 orbits), the luminosity of the accreting material can be estimated using

$$L = \epsilon \dot{M} c^2. \quad (6.7)$$

When the PPI is constraining the accretion rate the luminosity is found to be $L \approx 1.1 \times 10^{48}$ erg/s, assuming $\epsilon = 0.1$. This is $\sim 10^4$ times larger than the Eddington luminosity for such a black hole ($L_{\text{Edd}} = 1.28 \times 10^{44}$ erg/s).

6.7 Which instability drives initial accretion?

The PPI growth rate for the only unstable mode in the tidal disruption remnant torus is $\sim 0.5\Omega_0$. As this is slower than the growth rate of the MRI ($0.75\Omega_0$), it would be expected that the faster growth of the MRI would mean that accretion is initially driven through this mechanism. However, the initial magnetic field in the torus formed from a tidal disruption event is expected to be around the same magnitude as that for a star, ~ 1 G. This weak initial field means that the MRI will have to grow for many orbits before accretion can be established. During this time the PPI may be able to saturate and hence drive accretion.

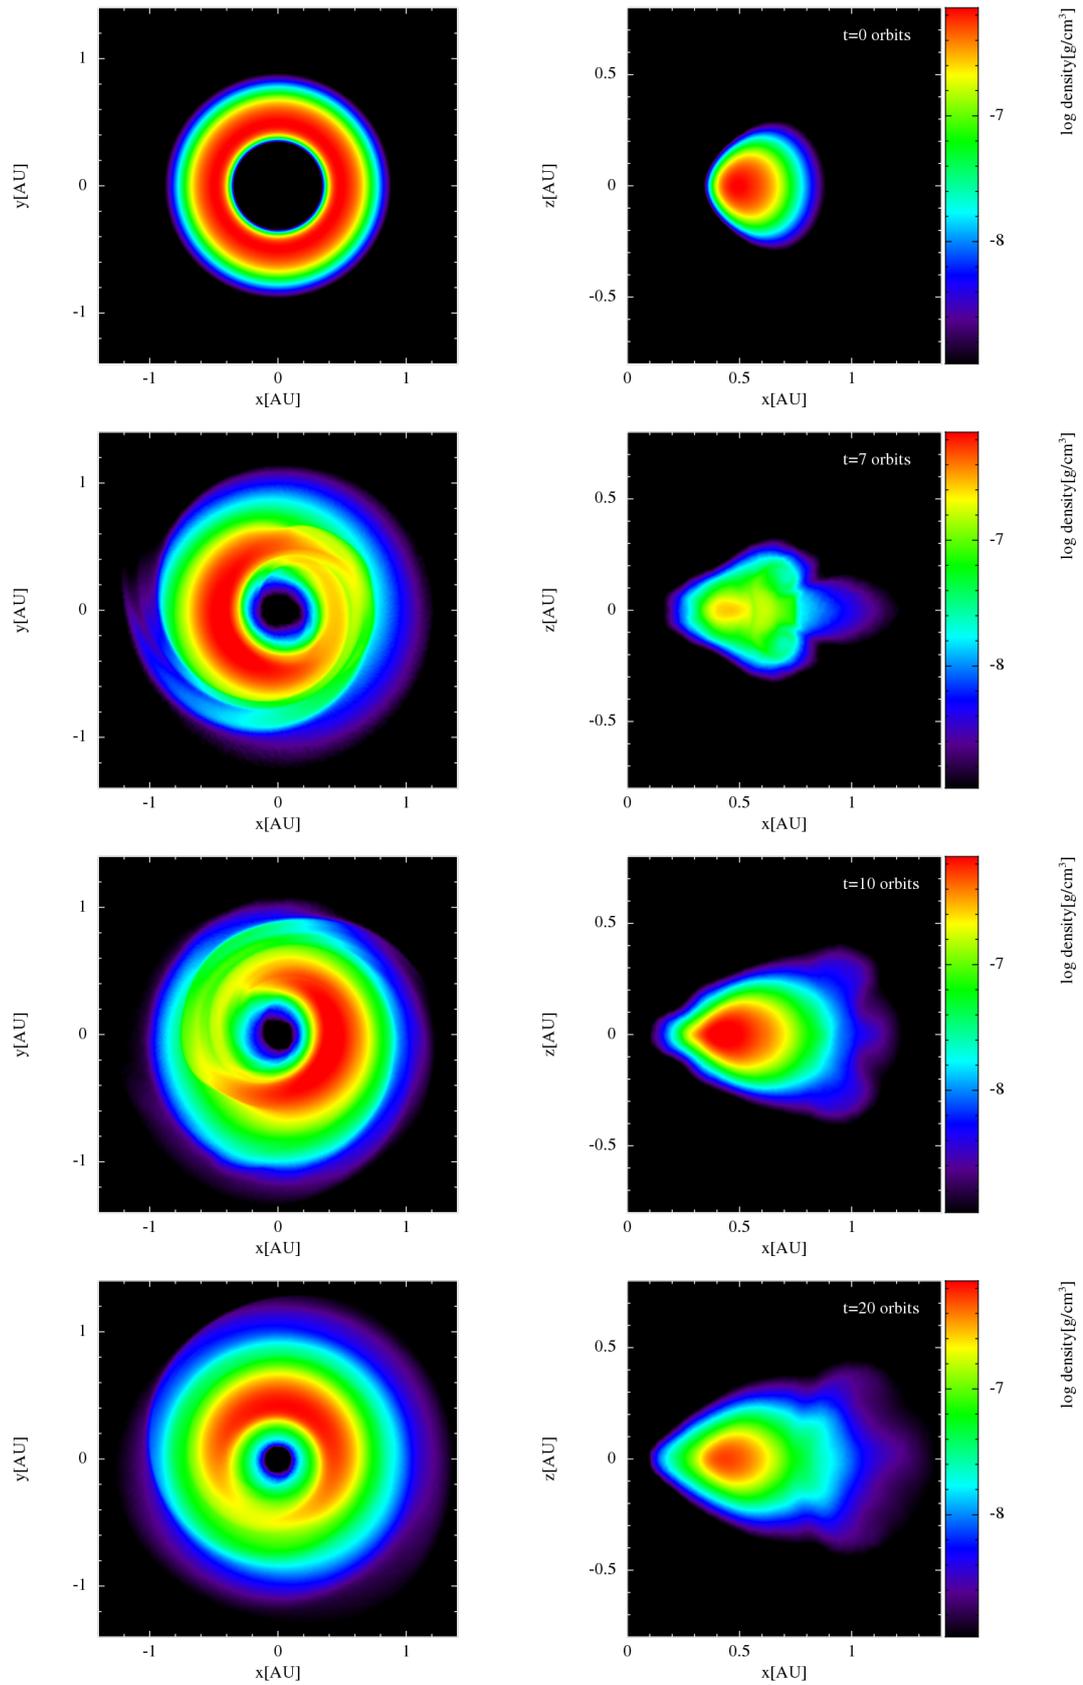


FIGURE 6.11: Development of the PPI in a torus that is initially similar to a tidal disruption remnant of a $1M_{\odot}$ star around a 10^6M_{\odot} black hole. The cross-sectional density evolution from our initial condition until ten orbits (at $R_0 = 0.5$ AU) is shown on a similar scale to Figure 8 of Bonnerot et al. (2016). As in our previous simulation there is significant spreading in the radial direction due to the PPI.

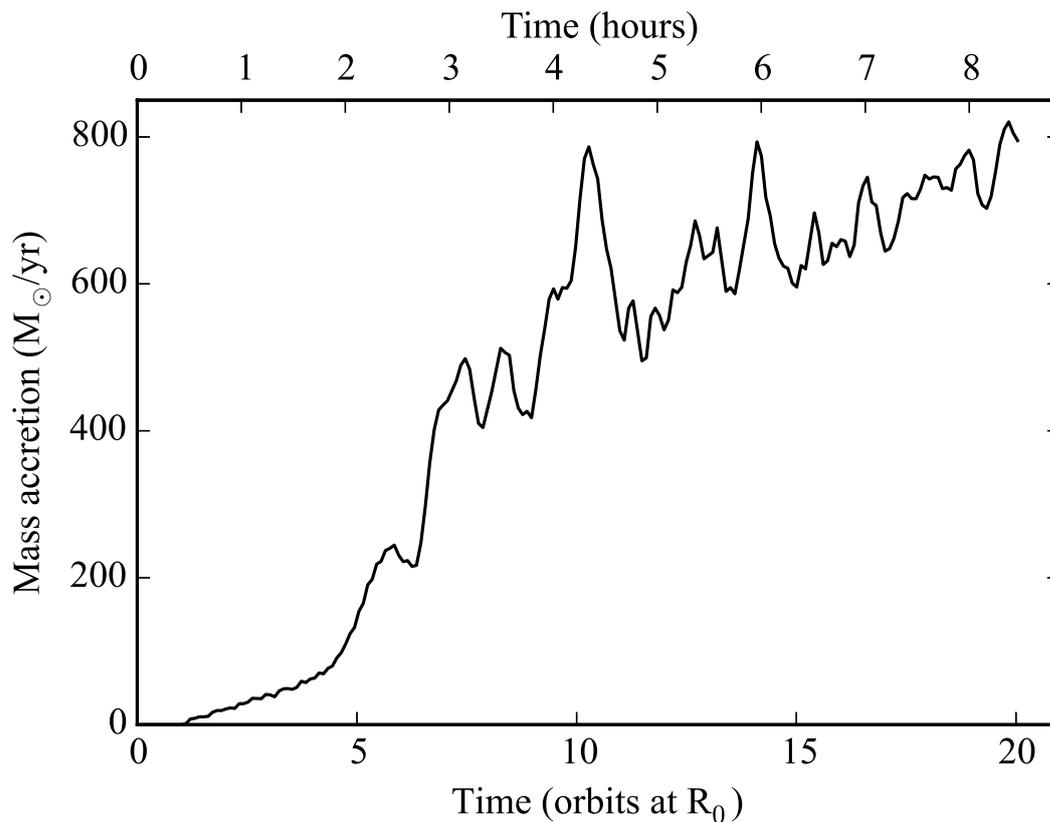


FIGURE 6.12: Mass accretion rate as a function of time in the tidal disruption remnant simulation shown in Figure 6.11. Accretion is primarily driven by the growth and saturation of the Papaloizou-Pringle instability in the remnant where the $m = 1$ mode dominates. The mass accretion rate is super Eddington ($\dot{M}_{\text{Edd}} = (\epsilon/0.1) \times 2.1 \times 10^{-2} M_{\odot}/\text{yr}$).

We thus consider the saturation timescales of both the PPI and the MRI to be indicative of the timescale for significant angular momentum transport. Figure 6.13 shows these timescales for a range of initial magnetic field strengths. We measure the saturation of the PPI in the tidal disruption remnant as the time taken for strong shocks to develop. We consider the equivalent of Figure 6.5 for this torus in our high resolution simulation (with 1×10^7 particles) and identify when the density variation stops increasing (the equivalent of ≈ 1.8 orbits in Figure 6.5). The uncertainty in this measurement (represented by the blue shaded region in Figure 6.13) is estimated by repeating this process for the simulation with 1×10^6 particles, as our lowest resolution simulation does not display any clear saturation. The maximum field strength required to establish the MRI will occur when the Alfvén speed and sound speed are comparable. However the MRI can be established at lower fields strengths than this, shown recently by Bugli et al. (2017) and in cases where there is zero net flux (e.g. Hawley, 2001). As a result, we consider weaker magnetic fields where the ratio between the gas and magnetic pressure $\beta \sim 100$. The magnetic field required for

saturation in this case is

$$B_{\text{sat}} \approx \frac{1}{10} H \Omega \sqrt{4\pi\rho}, \quad (6.8)$$

$$\approx 7.1 \times 10^5 \text{ G} \left(\frac{H}{4.5 \times 10^{12} \text{ cm}} \right) \left(\frac{\Omega}{5.6 \times 10^{-4} \text{ s}} \right) \left(\frac{\rho}{6.31 \times 10^{-7} \text{ g cm}^3} \right)^{1/2}.$$

Using the maximum density and scale-height from the simulation in Figure 6.11 the magnetic field strength required is $B_{\text{sat}} \approx 7.1 \times 10^5 \text{ G}$. The saturation timescale for the MRI, t_{MRI} , in terms of the initial magnetic field B_{initial} is thus

$$t_{\text{MRI}} = \frac{1}{0.75\Omega_0} \ln \left(\frac{B_{\text{sat}}}{B_{\text{initial}}} \right). \quad (6.9)$$

The comparison in Figure 6.13 demonstrates that for reasonable initial magnetic fields, the PPI does not grow fast enough to cause accretion before the MRI saturates. For example, in the case of the tidal disruption remnant shown in Figure 6.11 with an initial magnetic field of 1 G, the growth rate above suggests that the MRI will saturate in less than 3 orbits.

In the case that the PPI has already generated a density perturbation before it is damped by the MRI, this structure remains (e.g. the lower panels of Figure 6.11). Thus as the MRI reaches its saturation value and begins to drive accretion, we anticipate that the remnant will accrete as normal but retaining any structure generated by the growing PPI.

6.8 Discussion

Informed by the resolution study of the thin ring simulation in Figure 6.7, we would consider the simulation of the tidal remnant with $\langle h \rangle / H \approx 0.05$ to be resolved. A resolution study (Figure 6.14) suggests that the converged mass accretion rate would be ~ 4 times lower than that shown in Figure 6.12, despite Figure 6.5 and 6.7 suggesting that the development of the PPI is fully resolved.

Assuming that a converged mass accretion rate is less than one order of magnitude lower than measured from our highest resolution simulation (a conservative estimate based on Figure 6.12), the mass accretion rate due to the PPI is still calculated to be 10^3 times Eddington. The initial simulation by Bonnerot et al. (2016) also assumed an elliptic rather than parabolic orbit which increases the mass accreted while decreasing the time over which it is accreted, compounding our over-estimate of the mass accretion rate. However, even taking these into account we still predict a super-Eddington mass accretion rate due to the PPI.

For the measured super-Eddington luminosity we expect that strong outflows would develop. These may be either in the form of an expanding Eddington-limited spherical bubble (Loeb and Ulmer, 1997) or in the form of strong winds (Lodato and Rossi, 2011). Although

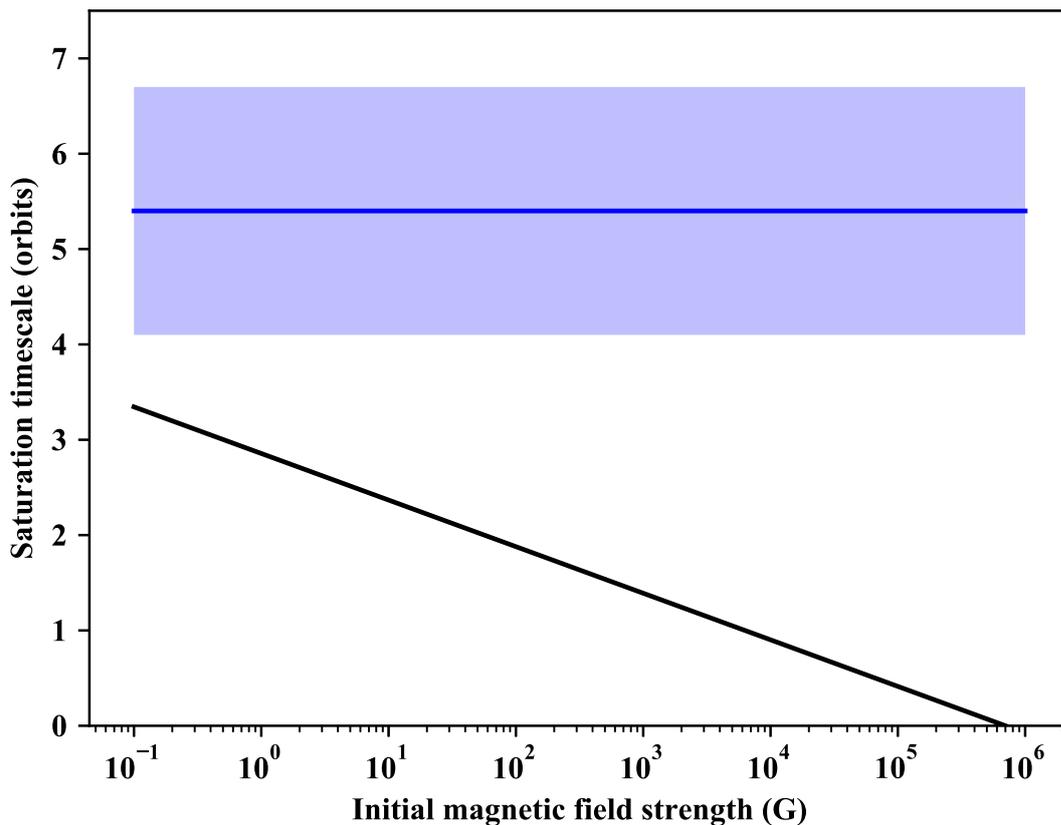


FIGURE 6.13: Timescale for saturation as a function of initial magnetic field strength of the Papaloizou-Pringle (blue) and magnetorotational instabilities (black) expected in the tidal disruption remnant shown in Figure 6.11. Although the PPI grows at a competitive rate to the MRI, it is unable to drive the initial accretion in the torus.

these outflows may drive radiative feedback, the current observations of TDEs (with the $t^{-5/3}$ fallback rate) do not suggest this is the case.

The original simulation by Bonnerot et al. (2016) assumed an elliptical rather than parabolic orbit such that our mass accretion is likely to be an overestimate. This choice means that the entire star forms into the torus rather than the $\lesssim 50\%$ expected from theoretical predictions, making our mass accretion rate roughly an order of magnitude larger than it should be. The accretion timescale is also shorted by the elliptical orbit (by a factor of ≈ 0.03 for $e = 0.8$), further compounding our overestimate. However, even while taking these differences into account (and the lower mass accretion rate discussed above) we still predict a super-Eddington mass accretion rate — suggesting a high mass accretion rate remains likely for parabolic encounters.

An interesting corollary of our work is the behaviour of the mass accretion rate after the PPI has saturated. In the early evolution of the torus when accretion is driven by the PPI, does the rate of accretion lead to self regulation or self-damping of the instability? As accretion due to numerical viscosity becomes important in our simulations around this

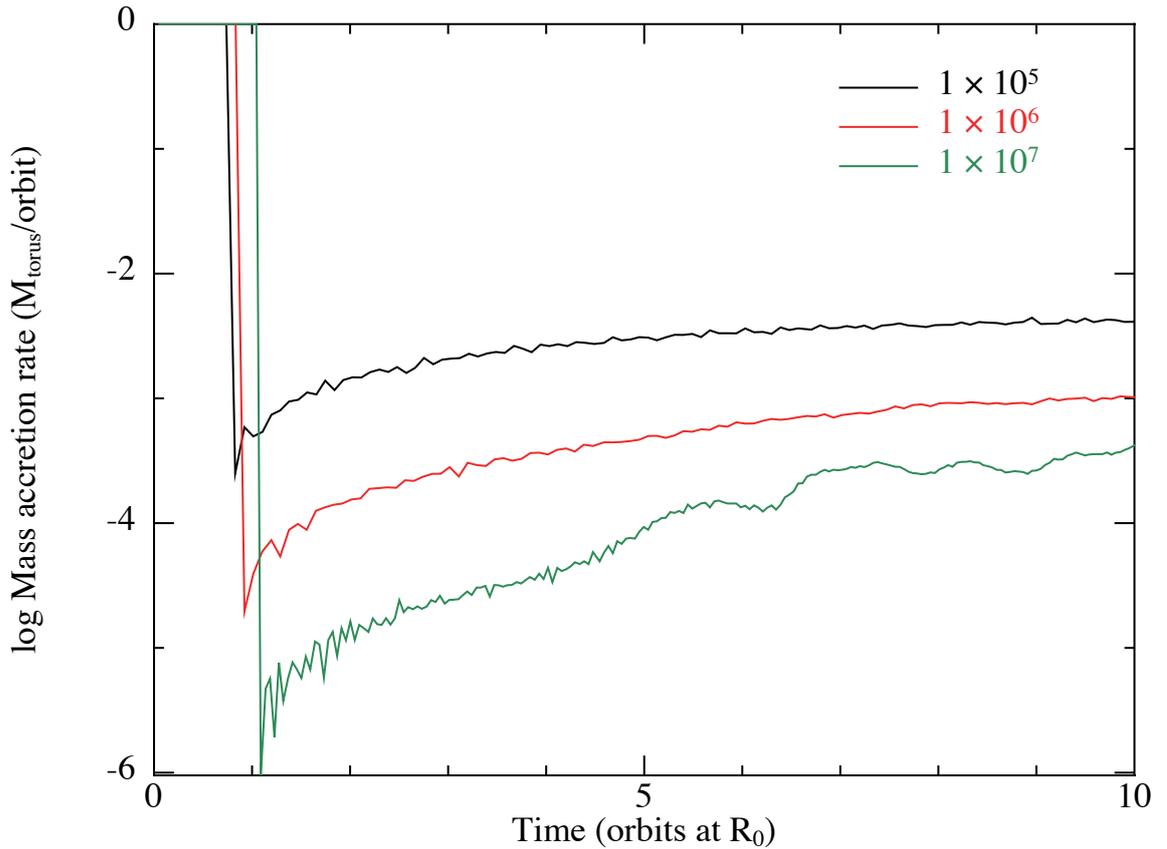


FIGURE 6.14: Mass accretion rate of the torus simulated in Figure 6.11 at three different resolutions. Even though the behaviour of the PPI is resolved in our simulations, the mass accretion rate generated from it is not yet and may be lower by a factor of ~ 4 (but not low enough to suggest a sub-Eddington mass accretion rate).

point, this is difficult to consider with our current simulations, but may be worth exploring in future work.

6.9 Conclusions

We have investigated the evolution of circularised tori with well defined inner and outer boundaries, similar to the compact discs formed by the tidal disruption of stars around a supermassive black hole. We demonstrate unmagnetised remnants of tidal disruption events are unstable to the $m = 1$ mode of the Papaloizou-Pringle instability. We find that this instability drives ring spreading and angular momentum transport that may be parameterised in terms of a Shakura-Sunyaev α viscosity. As the accretion in a tidal disruption event will be driven by the MRI after a few orbits, these accretion rates are not expected to be sustained for a long time.

Chapter 7

Conclusions

We have considered the evolution of accretion flows around a rotating compact object where the angular momentum vector of the flow is misaligned to the compact object spin vector. While misaligned accretion discs are expected and observed in different types of compact objects, we restrict ourselves to accretion discs around rotating black holes and in particular, those in low mass X-ray binaries. We have shown that relaxing the assumption of aligned accretion and taking into account general relativistic effects leads to a greater diversity of accretion flows — rather than just being flat, continuous and accreting on the viscous timescale, geometrically thin discs may have tilt oscillations or be torn into discrete, precessing rings. In the geometrically thick case, we considered an aligned torus that is formed after a tidal disruption event including general relativistic effects and demonstrated the rapid development of non-axisymmetric perturbations. After modelling these features using simulations we considered their observational consequences.

In Chapter 3 we considered the generic evolution of a geometrically thin, wave-like inclined accretion disc around a rotating black hole, complementing previous studies in the diffusive regime by Nixon et al. (2012). For the first time using three-dimensional simulations, at small inclinations we demonstrated the tilt profile of the disc settles into an oscillatory steady state with a misaligned inner edge — in agreement with analytical predictions. Importantly, the misaligned inner edge suggests that if a jet is launched from such a disc it will not necessarily be aligned with the black hole spin. At larger inclinations between the black hole spin and accretion disc we confirmed the phenomenon of disc breaking and tearing was possible in the wave-like regime. The differential precession of these precessing rings causes initially adjacent rings to develop opposing angular momenta, leading to rapid angular momentum cancellation and an enhanced accretion rate. In comparison to a fully aligned disc, we demonstrated that the mass accretion rate could be increased by up to an order of magnitude. In hand with the previous work in the diffusive regime this suggested that it is inevitable that strongly misaligned discs will tear and have an enhanced mass accretion rate.

In Chapter 4 we continued small inclination discs, conducting a comparison to a similar grid based simulation of a warped accretion disc by Krolik and Hawley (2015). For the disc parameters chosen we confirmed that purely hydrodynamic simulations using the isotropic α viscosity model represented an equivalent evolution to a full MHD treatment.

We then (again) demonstrated that GR effects like apsidal precession are critical to the disc evolution, recovering tilt oscillations and a misaligned inner edge when this effect was included. Finally, we clarified the applicability of the analytic disc tearing criterion by demonstrating that a disc with the parameters chosen by Krolik and Hawley (2015) can tear only at higher inclination.

Although we have confirmed that it is inevitable for strongly misaligned discs to tear, computational restrictions mean that the long term behaviour of tearing remains poorly understood. Is disc tearing a steady behaviour or isolated to a particular disc state (or transition between states)? In strongly misaligned discs that experience tearing, how does this affect observational estimates of the mass accretion rate? How is tearing affected by relaxing the assumptions made in our simulations (e.g. a locally isothermal equation of state)? At low inclinations the oscillatory tilt profile recovered is a steady state, so how much does this affect observations that measure the black hole spin while assuming the disc is aligned? What range of physically motivated parameters result in tilt oscillations? Do we expect tilt oscillations to be an observable phenomenon?

In Chapter 5 we attempted to demonstrate a relationship between disc tearing and quasi-periodic oscillations in low mass X-ray binaries. We argue that the features from observations including the radius in the disc that QPOs are understood to be generated from, the strong geometric dependence of low-frequency QPOs and the frequencies with which they occur are suggestive of the tearing rings that are generated in our simulations. The nature and the frequency range of low-frequency QPOs suggests that the accretion disc must be misaligned and be inhomogeneous. Tearing discs provides a natural mechanism to generate structures that meet both of these criteria. We focussed on simulations with parameters that were physically motivated from the LMXB XTE J1550-564, confirming that disc tearing could i) generate the structures that are inferred by current phenomenological models of LFQPOs, ii) that this process was repeatable and not a function of our initial conditions and iii) that these structures precessed at approximately the rate expected by observations. We developed a crude light curve implementation that included the effects of obscuration to measure the luminosity observed from this system and hence a power density spectrum. We demonstrated that this method was capable of representing features that occurred from the innermost precessing structures rather than the initial tearing phase. Despite these developments we were unable to produce a feature in a power density spectrum that was both resolution independent and significant. We discuss in depth the numerical (rather than physical) constraints that we attribute this to.

Simulations to demonstrate the connection between tearing discs and quasi-periodic oscillations are particularly computationally demanding, requiring long duration high resolution simulations with a large radial extent. Additionally, in order to measure any kind of luminosity from such a simulation is complex and requires careful consideration. This investigation in particular would benefit from significant code development, including using a physically motivated cooling model (previous simulations have shown that this is likely to lead to the formation of a corona, see Figure 3 of Nixon and Salvesen, 2014). Correctly

modelling radiation transfer is also expected to alter the disc evolution, so is also suitable to be included. To measure a power density spectrum accurately from our simulations, ideally an implementation that models both obscuration and reflection should also be used. Finally, irrespective of whether the disc is in the diffusive or wave-like regime, inclining the disc has strong observational consequences. LMXBs are expected to be in the wave-like regime while they are in the hard state (Ingram, Done, and Fragile, 2009) where tilt oscillations are found in our simulations. These oscillations manifest in the spatial region where high-frequency quasi-periodic oscillations are understood to be sourced, but these are observed in a different disc state. Do tilt oscillations represent a potential avenue for these quasi-periodic oscillations?

In Chapter 6 we considered the evolution of a torus formed after a tidal disruption event. Bonnerot et al. (2016) simulated the evolution of a tidal disruption where apsidal precession was taken into account and the gas cooled inefficiently, finding rapid circularisation and the formation of a thick torus. We demonstrated that when this torus was not already accreting it was unstable to the Papaloizou-Pringle instability. We considered the subsequent evolution of a circularised, more compact version of the remnant. The instability developed quickly in the remnant, preventing the torus from staying circularised. Using a radially narrow torus we showed that this instability was capable of driving angular momentum transport but that particularly high resolution was required to isolate this from the effects of numerical viscosity. Applied to our simulation of the tidal disruption remnant, this suggested that the mass accretion rate when the torus first forms depends on this instability. As the magnetorotational instability in such a torus would take ~ 3 orbits to develop from the initial weak magnetic field, we concluded that the Papaloizou-Pringle instability may drive some mass accretion rate in the early life of the torus. Initial accretion driven from the PPI rather than the MRI may be possible in tori that have a particularly weak initial magnetic field, such as after a tidal disruption event.

All of the simulations in this work were purely hydrodynamical. We demonstrated for the evolution of warped, geometrically thin discs with low inclinations warps that this is equivalent to a full MHD treatment. However, in the high inclination case where tearing occurs no such comparison is available. In simulations that exhibit tearing it is difficult to be as confident that our approaches are equivalent because of the inherent difficulty in resolving the regions between precessing rings, the rings themselves and their interactions (even when the tearing region in the disc may be resolved well). It is not intuitively clear how the presence of magnetic fields would affect the evolution of the tearing rings — would their tearing be delayed? Would magnetic fields alter the communication between the rings (and hence the mass accretion rate)? Thus including magnetic fields in tearing simulations is ultimately important, however their implementation in a simulation that uses particle methods remains particularly challenging. Additionally, simulations of tori would benefit from the inclusion of MHD as we could then directly quantify the competition between the development of the PPI and the MRI, and hence their impact on the accretion rate.

Finally, the applicability of the post-Newtonian approximation to model general relativistic effects used in this work should be examined carefully. At larger radii where tearing occurs, the approximation agrees well with GR predictions such that disc tearing is not expected to be affected by including full GR. However, the approximation is not particularly good inside $\approx 10R_g$, where tilt oscillations are found. The inclusion of a full GR description is thus expected to alter the shape of the oscillatory tilt profile, but it is not immediately clear by how much. The evolution of tidal disruption events would be reproduced more accurately in full GR, as those that have been simulated so far with PHANTOM can only include the effects of apsidal precession (e.g. Bonnerot et al., 2016). With GR, these could be extended to investigate the effects of nodal precession on the circularisation process and in the event that a torus or disc did form, whether it exhibits global precession. A full GR implementation is currently being developed for PHANTOM, so the validity of this approximation could be tested in the near future.

Appendix A

Measuring disc precession

Here we outline how the precession in the disc was measured from the simulations and show some example results. In order to analyse the properties of the discs from the simulations, we discretise the disc into a set of thin spherical annuli and average the properties of interest across the particles in each of these bins. This process is described in detail in Section 3.2.6 of Lodato and Price (2010). The twist, $\gamma(R)$, in our disc at a given radius is found by considering the unit angular momentum vectors at each radius bin in the disc. With $l_x(R)$, $l_y(R)$ and $l_z(R)$ being the unit vectors in the Cartesian coordinate system, we assign

$$\gamma(R) = \tan^{-1} \left(\frac{l_y(R)}{l_x(R)} \right), \quad (\text{A.1})$$

for each radial annulus. This is repeated at every time step, so that we have a description of the twist as a function of time at each radial bin in our simulations. An example of the twist in this format is in Figure A.1. Here the twist is increasing in the disc when the gradient is positive. As the disc twists through a full 2π radians, the twist then jumps back to zero because Equation A.1 does not take into account the cumulative twist angle.

The precession time can be measured from Figure A.1 directly by recording how long it takes for the disc to twist all the way around, equivalent to finding when the twist drops back to zero. This can be approximated by working out the gradient of the twist as a function of time and then using it to calculate the precession time in the disc. This is equivalent to calculating

$$t_p = 2\pi \left(\frac{d\gamma}{dt} \right)^{-1}. \quad (\text{A.2})$$

Because this calculation has been done at each radial annulus, we now have the precession time as a function of the radius in the disc, averaged over the length of the simulation. An example of this was shown in Figure 2.5.

The above analysis does not take the inclination of the disc into account. This disc was inclined at 30° , but this angle did not come into our expression for t_p or explicitly in our analysis from the twist. As outlined in the derivation by Larwood et al. (1996), the Lense-Thirring precession is independent of the inclination between the disc and the black hole spin. Repeating the above analysis with discs at other angles confirms this.

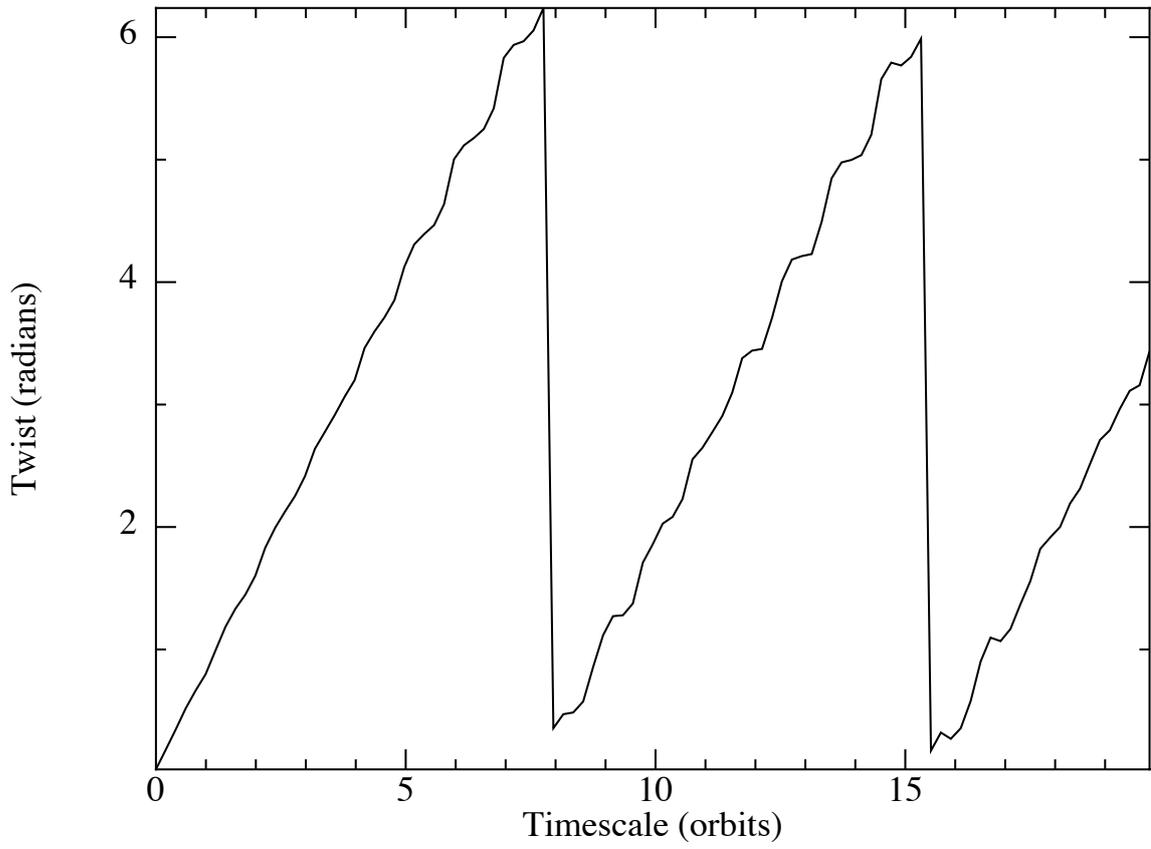


FIGURE A.1: An example of the description of twist in the disc as a function of time at a given radius. Here the time is plotted in orbits measured at the inner edge.

Appendix B

Simulated power density spectra

B.1 PDS Test

To test the program written to calculate the power density spectra (PDS) is correct, here the code is tested against some fake data. The data, $g(t)$, is constructed using

$$g(t) = \sin(t) + \sin(2t) + \cos(4t) + \cos(5t), \quad (\text{B.1})$$

where t is randomly sampled. This should result in frequencies detected at 1, 2, 4 and 5 Hz. When scaled by 2π , this suggests frequencies of 0.159, 0.318, 0.637 and 0.795. The PDS generated in Figure B.2 identifies frequencies at 0.159, 0.318, 0.637 and 0.795. Here the black line represents the full PDS and the red a coarser, binned PDS. The small feature around 2 Hz is noise, rather than a harmonic.

B.2 Power density spectra

We present additional power density spectra produced from the parameter sweeps S1 - S3 (see Table 5.1). Each parameter sweep covers the 36 combinations of viscosity parameter $\alpha = 0.03, 0.1, 0.3$, black hole spin parameter $a = 0.1, 0.5, 0.9$, disc inclination $\beta = 30^\circ, 60^\circ$ and aspect ratio $H/R = 0.02, 0.05$. Figures B.3-B.5 show results for $\beta = 30^\circ$, $H/R = 0.02$ for all values of viscosity and spin, but these results are representative of the full parameter sweep. In each figure, the PDS calculated from the luminosity approximation (in the optically thin limit, without the effects of obscuration) is presented in black and the PDS calculated from the mass accretion rate is presented in red. The lowest frequency in the PDS is set by the length of the simulation and the highest by the orbital frequency at the innermost stable orbit, not including general relativistic corrections. Figure B.6 shows PDS for simulation S6 (see Figure 5.11) including the effects of obscuration comparing with and without poorly resolved particles, demonstrating the influence that poorly resolved particles have on the PDS features.

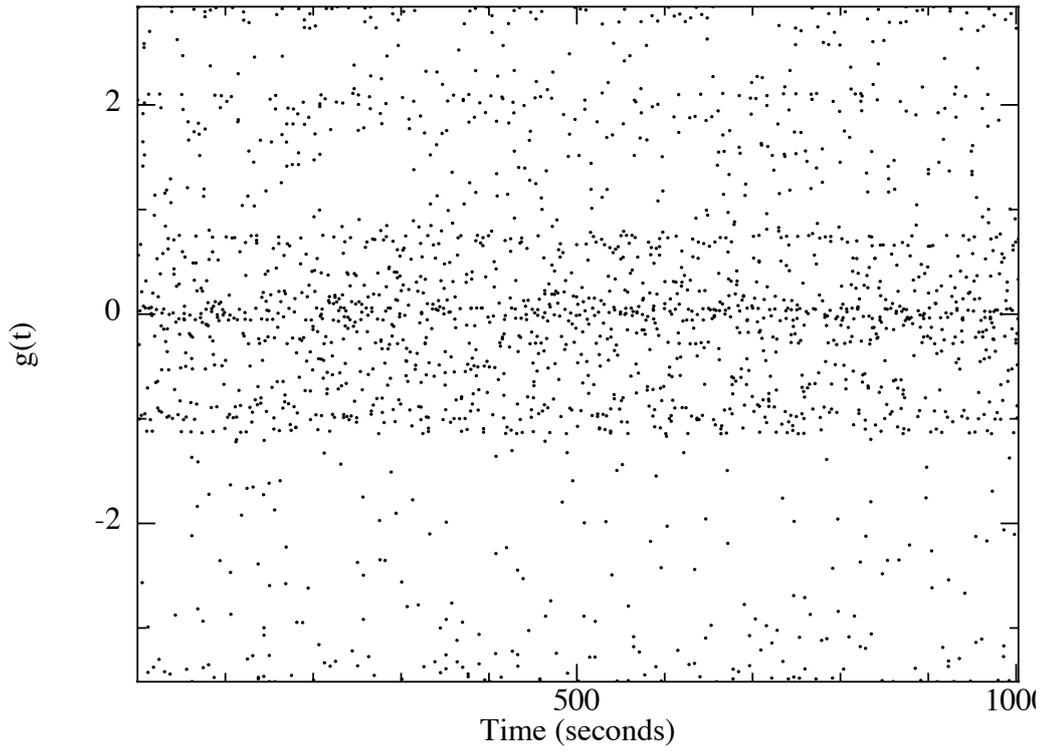


FIGURE B.1: Mock data used to test that the PDS technique works. The data is constructed from Equation B.1, using random values for t .

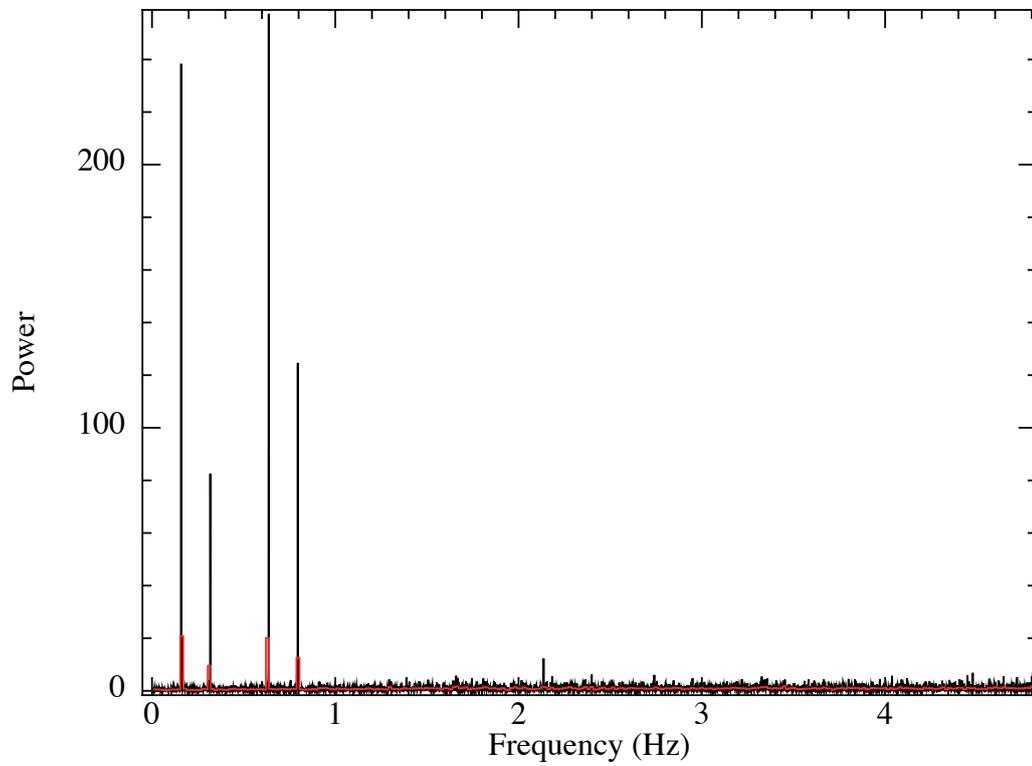


FIGURE B.2: The PDS derived from Figure B.1 using our implementation of the Lomb-Scargle method.

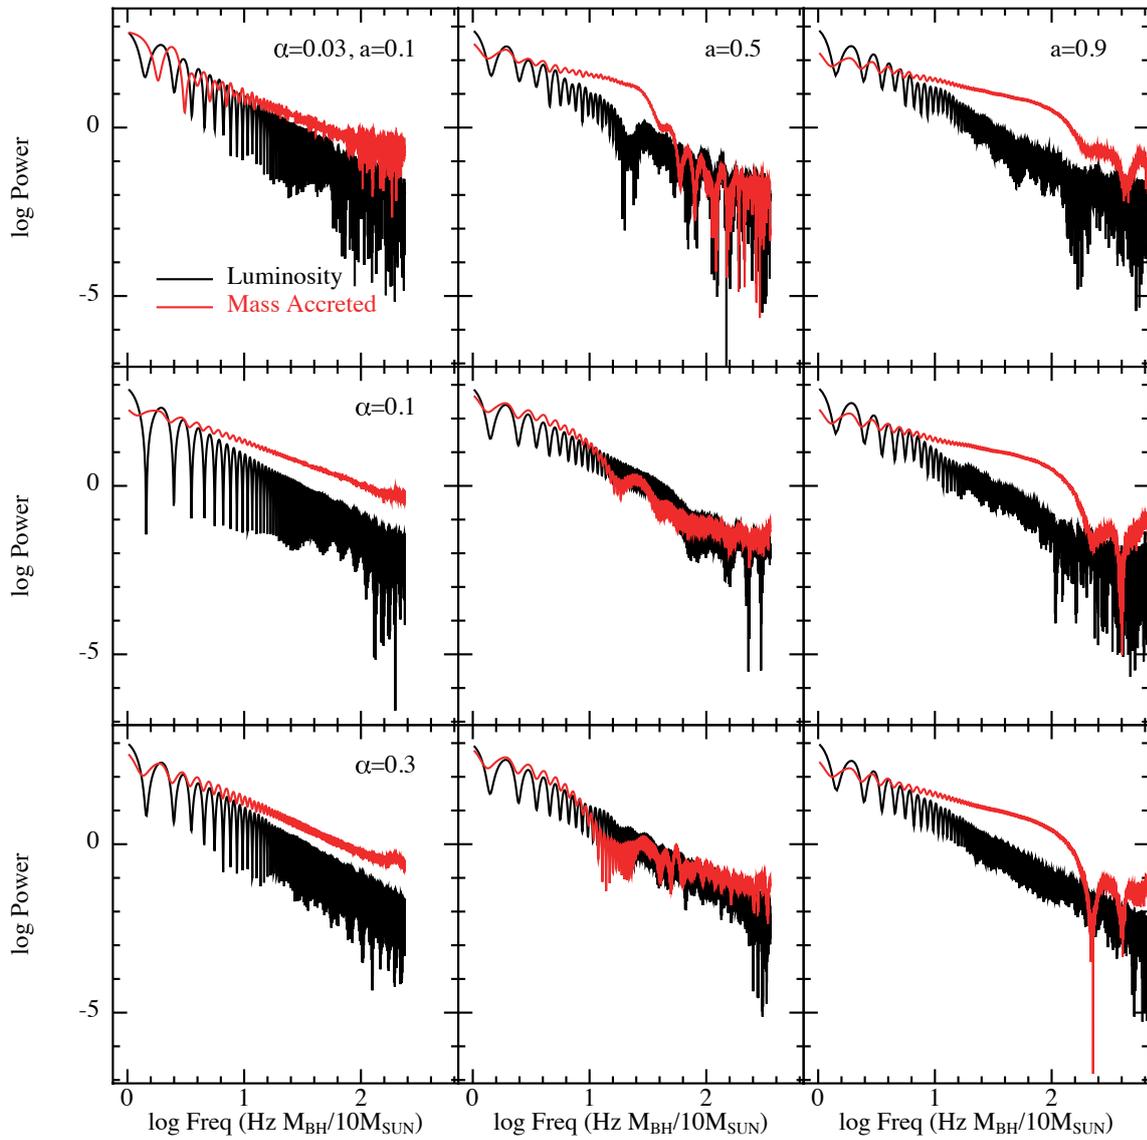


FIGURE B.3: Power spectra from both mass accretion rate and luminosity from the inner disc simulations (S1). Despite these simulations showing visual evidence of precession, nothing significant is present in the powerspectra. Here $H/R = 0.02$ and $\beta = 30^\circ$, but these results are typical for the other combinations of parameters.

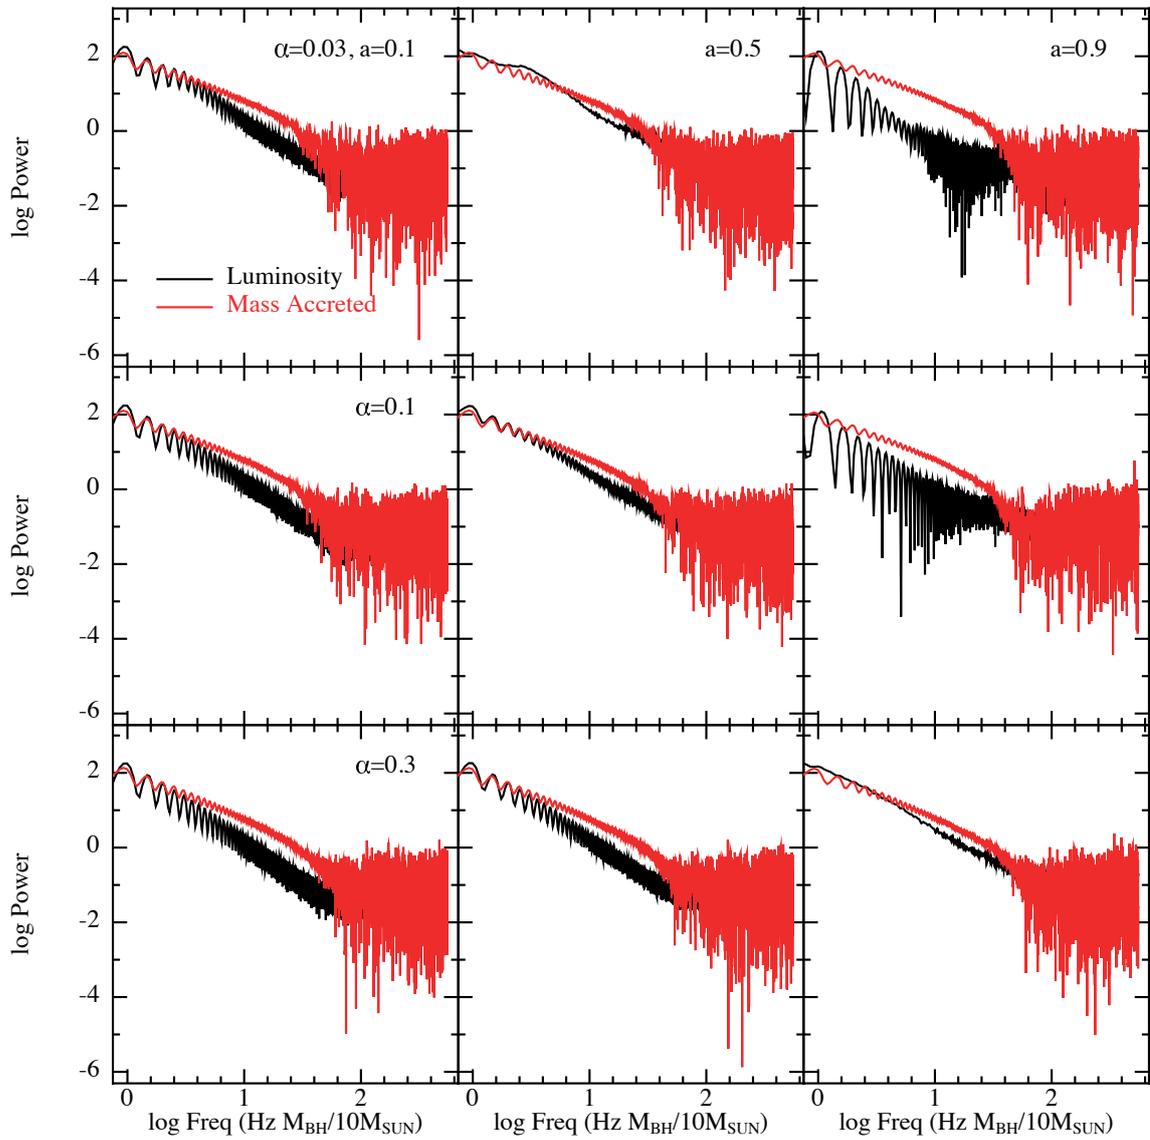


FIGURE B.4: Power spectra from both mass accretion rate and luminosity from the outer disc simulations (S2). These simulations show no signs of tearing, so it is not surprising that nothing significant is present in the powerspectra. Here $H/R = 0.02$ and $\beta = 30^\circ$, but these results are typical for the other combinations of parameters.

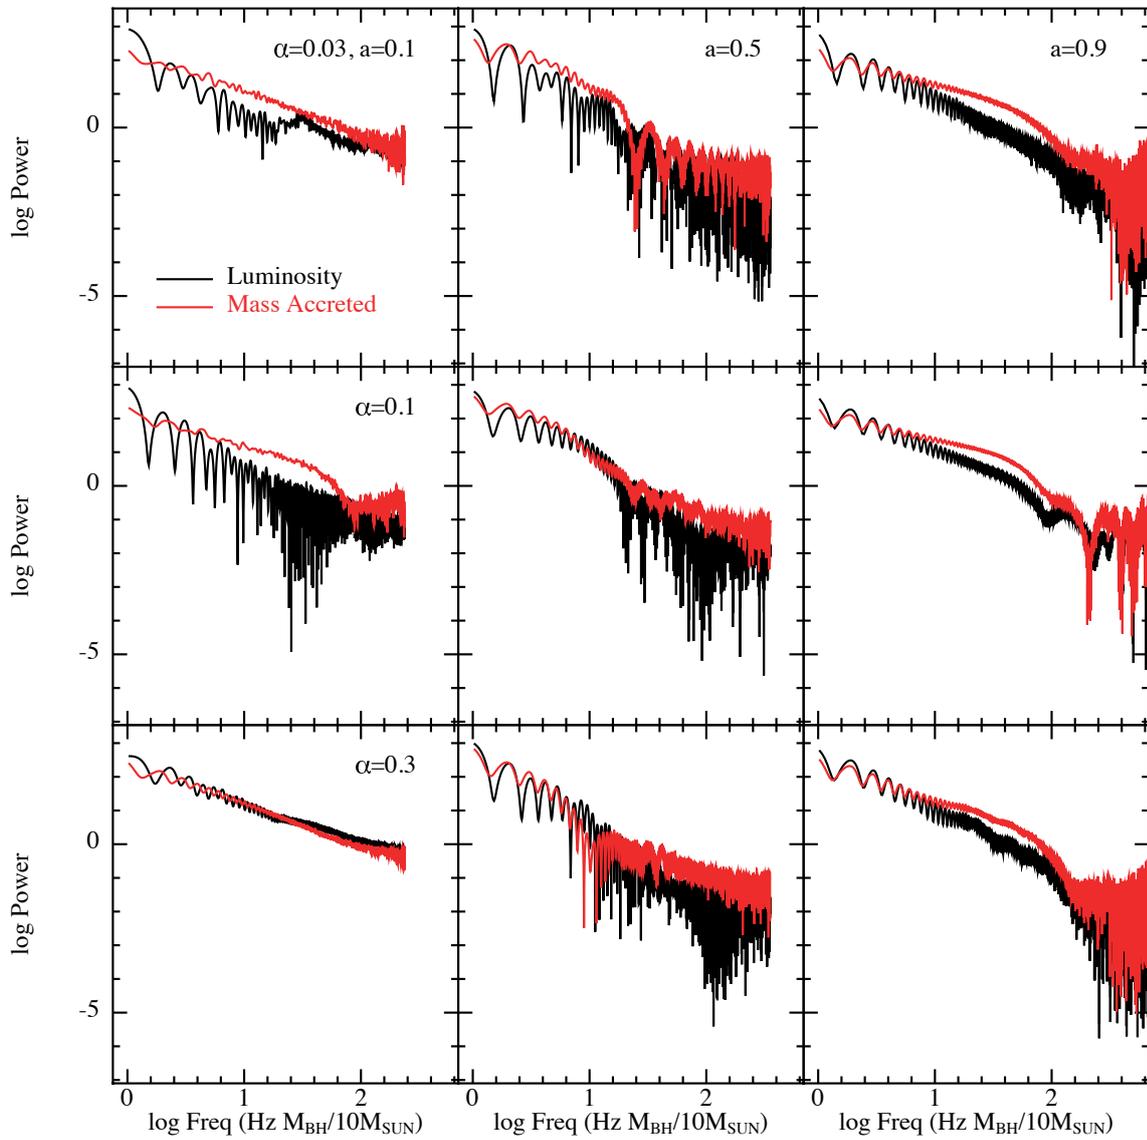


FIGURE B.5: Power spectra from both mass accretion rate and luminosity from the inner disc simulations (S3). Despite these simulations showing visual evidence of precession, nothing significant is present in the powerspectra. Here $H/R = 0.02$ and $\beta = 30^\circ$, but these results are typical for the other combinations of parameters.

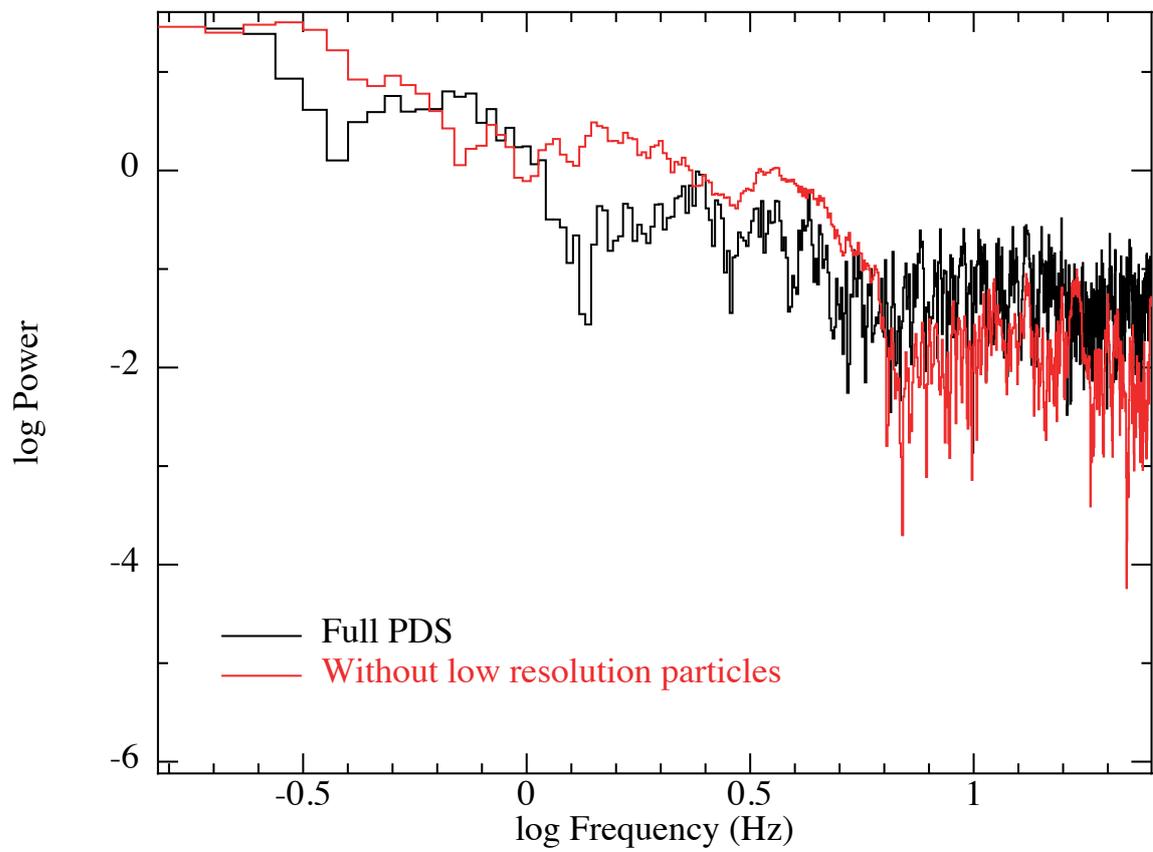


FIGURE B.6: Power spectra from simulation (S6), including the effects of obscuration, identifying the contribution of the low resolution particles.

Bibliography

- Artymowicz, P. and S. H. Lubow (1994). “Dynamics of binary-disk interaction. 1: Resonances and disk gap sizes”. In: *ApJ* 421, pp. 651–667.
- Axelsson, M., L. Hjalmarsson, and C. Done (2013). “Fast variability as a probe of the smallest regions around accreting black holes”. In: *MNRAS* 431, pp. 1987–1994.
- Ayal, S., M. Livio, and T. Piran (2000). “Tidal Disruption of a Solar-Type Star by a Supermassive Black Hole”. In: *ApJ* 545, pp. 772–780.
- Babkovskaia, N. et al. (2006). “A warped $m = 2$ water maser disc in V778 Cyg?”. In: *MNRAS* 370, pp. 1921–1927.
- Balbus, S. A. (2003). “Enhanced Angular Momentum Transport in Accretion Disks”. In: *ARA&A* 41, pp. 555–597.
- Balbus, S. A. and J. F. Hawley (1991). “A powerful local shear instability in weakly magnetized disks. I - Linear analysis. II - Nonlinear evolution”. In: *ApJ* 376, pp. 214–233.
- (1998). “Instability, turbulence, and enhanced transport in accretion disks”. In: *Reviews of Modern Physics* 70, pp. 1–53.
- Bardeen, J. M. and J. A. Petterson (1975). “The Lense-Thirring Effect and Accretion Disks around Kerr Black Holes”. In: *A&A* 195, p. L65.
- Belloni, T., M. Méndez, and J. Homan (2005). “The distribution of kHz QPO frequencies in bright low mass X-ray binaries”. In: *A&A* 437, pp. 209–216.
- Belloni, T., D. Psaltis, and M. van der Klis (2002). “A Unified Description of the Timing Features of Accreting X-Ray Binaries”. In: *ApJ* 572, pp. 392–406.
- Belloni, T. M. (2010). “States and Transitions in Black Hole Binaries”. In: *Lecture Notes in Physics, Berlin Springer Verlag*. Ed. by T. Belloni. Vol. 794. Lecture Notes in Physics, Berlin Springer Verlag, p. 53.
- Belloni, T. M. and S. E. Motta (2016). “Transient Black Hole Binaries”. In: *Astrophysics of Black Holes: From Fundamental Aspects to Latest Developments*. Ed. by C. Bambi. Vol. 440. Astrophysics and Space Science Library, p. 61.
- Belloni, T. M., A. Sanna, and M. Méndez (2012). “High-frequency quasi-periodic oscillations in black hole binaries”. In: *MNRAS* 426, pp. 1701–1709.

- Belloni, T. M. and L. Stella (2014). “Fast Variability from Black-Hole Binaries”. In: *Space Science Rev.* 183, pp. 43–60.
- Benz, W. (1990). “Smooth Particle Hydrodynamics: A Review”. In: *The Numerical Modelling of Nonlinear Stellar Pulsations: Problems and Prospects*. Ed. by J. Robert Buchler. Dordrecht: Springer Netherlands, pp. 269–288.
- Blaes, O. M. (1987). “Stabilization of non-axisymmetric instabilities in a rotating flow by accretion on to a central black hole”. In: *MNRAS* 227, pp. 975–992.
- Blaes, O. M. and W. Glatzel (1986). “On the stability of incompressible constant angular momentum cylinders”. In: *MNRAS* 220, pp. 253–258.
- Blaes, O. M. and J. F. Hawley (1988). “Nonaxisymmetric disk instabilities - A linear and nonlinear synthesis”. In: *ApJ* 326, pp. 277–291.
- Bonnerot, C., E. M. Rossi, and G. Lodato (2017). “Long-term stream evolution in tidal disruption events”. In: *MNRAS* 464, pp. 2816–2830.
- Bonnerot, C. et al. (2016). “Disc formation from tidal disruptions of stars on eccentric orbits by Schwarzschild black holes”. In: *MNRAS* 455, pp. 2253–2266.
- Brenneman, L. W. and C. S. Reynolds (2006). “Constraining Black Hole Spin via X-Ray Spectroscopy”. In: *ApJ* 652, pp. 1028–1043.
- Brookshaw, L. (1985). “A method of calculating radiative heat diffusion in particle simulations”. In: *Proceedings of the Astronomical Society of Australia* 6, pp. 207–210.
- Bugli, M. et al. (2017). “Papaloizou-Pringle instability suppression by the magnetorotational instability in relativistic accretion discs”. In: *ArXiv e-prints*.
- Cabanac, C. et al. (2010). “Variability of X-ray binaries from an oscillating hot corona”. In: *MNRAS* 404, pp. 738–748.
- Cannizzo, J. K. and N. Gehrels (2009). “A New Paradigm for Gamma-ray Bursts: Long-term Accretion Rate Modulation by an External Accretion Disk”. In: *ApJ* 700, pp. 1047–1058.
- Casassus, S. et al. (2015). “Accretion Kinematics through the Warped Transition Disk in HD142527 from Resolved CO(6-5) Observations”. In: *ApJ* 811, 92, p. 92.
- Casella, P., T. Belloni, and L. Stella (2005). “The ABC of Low-Frequency Quasi-periodic Oscillations in Black Hole Candidates: Analogies with Z Sources”. In: *ApJ* 629, pp. 403–407.
- Coughlin, E. R. and M. C. Begelman (2014). “The General Relativistic Equations of Radiation Hydrodynamics in the Viscous Limit”. In: *ApJ* 797, 103, p. 103.

- Coughlin, E. R. et al. (2016). “Post-periapsis pancakes: sustenance for self-gravity in tidal disruption events”. In: *MNRAS* 455, pp. 3612–3627.
- Coughlin, E. R. et al. (2017). “Tidal disruption events from supermassive black hole binaries”. In: *MNRAS* 465, pp. 3840–3864.
- Cui, W. et al. (1999). “Strong Aperiodic X-Ray Variability and Quasi-Periodic Oscillation in X-Ray Nova XTE J1550-564”. In: *ApJL* 512, pp. L43–L46.
- Cullen, L. and W. Dehnen (2010). “Inviscid smoothed particle hydrodynamics”. In: *MNRAS* 408, pp. 669–683.
- Debes, J. H. et al. (2017). “Chasing Shadows: Rotation of the Azimuthal Asymmetry in the TW Hya Disk”. In: *ApJ* 835, 205, p. 205.
- Dehnen, W. and H. Aly (2012). “Improving convergence in smoothed particle hydrodynamics simulations without pairing instability”. In: *MNRAS* 425, pp. 1068–1082.
- Dipierro, G. et al. (2015). “On planet formation in HL Tau”. In: *MNRAS* 453, pp. L73–L77.
- Dipierro, G. et al. (2016). “Two mechanisms for dust gap opening in protoplanetary discs”. In: *MNRAS* 459, pp. L1–L5.
- Done, C., M. Gierliński, and A. Kubota (2007). “Modelling the behaviour of accretion flows in X-ray binaries. Everything you always wanted to know about accretion but were afraid to ask”. In: *A&A Rev.* 15, pp. 1–66.
- Dönmez, O. (2014). “On the development of the Papaloizou-Pringle instability of the black hole-torus systems and quasi-periodic oscillations”. In: *MNRAS* 438, pp. 846–858.
- Doğan, S. et al. (2015). “Tearing up a misaligned accretion disc with a binary companion”. In: *MNRAS* 449, pp. 1251–1258.
- Esin, A. A., J. E. McClintock, and R. Narayan (1997). “Advection-Dominated Accretion and the Spectral States of Black Hole X-Ray Binaries: Application to Nova Muscae 1991”. In: *ApJ* 489, pp. 865–889.
- Español, P. and M. Revenga (2003). “Smoothed dissipative particle dynamics”. In: *Physical Review E* 67.2, p. 026705.
- Evans, C. R. and C. S. Kochanek (1989). “The tidal disruption of a star by a massive black hole”. In: *ApJL* 346, pp. L13–L16.
- Facchini, S., G. Lodato, and D. J. Price (2013). “Wave-like warp propagation in circumbinary discs - I. Analytic theory and numerical simulations”. In: *MNRAS* 433, pp. 2142–2156.
- Flebbe, O. et al. (1994). “Smoothed Particle Hydrodynamics: Physical viscosity and the simulation of accretion disks”. In: *ApJ* 431, pp. 754–760.

- Fragile, P. C. and P. Anninos (2005). “Hydrodynamic Simulations of Tilted Thick-Disk Accretion onto a Kerr Black Hole”. In: *ApJ* 623, pp. 347–361.
- Fragile, P. C., G. J. Mathews, and J. R. Wilson (2001). “Bardeen-Petterson effect and QPOs in low-mass X-ray binaries”. In: *20th Texas Symposium on relativistic astrophysics*. Ed. by J. C. Wheeler and H. Martel. Vol. 586. American Institute of Physics Conference Series, pp. 544–546.
- Fragile, P. C. et al. (2007). “Global General Relativistic Magnetohydrodynamic Simulation of a Tilted Black Hole Accretion Disk”. In: *ApJ* 668, pp. 417–429.
- Fragner, M. M. and R. P. Nelson (2010). “Evolution of warped and twisted accretion discs in close binary systems”. In: *A&A* 511, A77, A77.
- Fragos, T. et al. (2010). “Black Hole Spin-Orbit Misalignment in Galactic X-ray Binaries”. In: *ApJL* 719, pp. L79–L83.
- Franchini, A., G. Lodato, and S. Facchini (2016). “Lense-Thirring precession around supermassive black holes during tidal disruption events”. In: *MNRAS* 455, pp. 1946–1956.
- Franchini, A., S. Motta, and G. Lodato (2016). “Constraining black hole spins with low-frequency quasi-periodic oscillations in soft states”. In: *MNRAS*.
- Gezari, S. et al. (2012). “An ultraviolet-optical flare from the tidal disruption of a helium-rich stellar core”. In: *Nature* 485.7397, pp. 217–220.
- Gierliński, M. et al. (2008). “A periodicity of ~ 1 hour in X-ray emission from the active galaxy RE J1034+396”. In: *Nature* 455, pp. 369–371.
- Gingold, R. A. and J. J. Monaghan (1977). “Smoothed particle hydrodynamics - Theory and application to non-spherical stars”. In: *MNRAS* 181, pp. 375–389.
- Glatzel, W. (1987). “On the stability of compressible differentially rotating cylinders”. In: *MNRAS* 225, pp. 227–255.
- Goldreich, P., J. Goodman, and R. Narayan (1986). “The stability of accretion tori. I - Long-wavelength modes of slender tori”. In: *MNRAS* 221, pp. 339–364.
- Greene, J., C. D. Bailyn, and J. A. Orosz (2001). “Optical and Infrared Photometry of the Microquasar GRO J1655-40 in Quiescence”. In: *ApJ* 554, pp. 1290–1297.
- Gribbin, J. R., P. A. Feldman, and S. H. Plegemann (1970). “Physical Sciences: Short-term Periodic Variations following an Optical Flare of Sco X-1”. In: *Nature* 225.5238, pp. 1123–1124.
- Guillochon, J. and M. McCourt (2017). “Simulations of Magnetic Fields in Tidally Disrupted Stars”. In: *ApJL* 834, L19, p. L19.

- Guillochon, J. and E. Ramirez-Ruiz (2013). “Hydrodynamical Simulations to Determine the Feeding Rate of Black Holes by the Tidal Disruption of Stars: The Importance of the Impact Parameter and Stellar Structure”. In: *ApJ* 767, 25, p. 25.
- (2015). “A Dark Year for Tidal Disruption Events”. In: *ApJ* 809, 166, p. 166.
- Hasegawa, Y. et al. (2017). “Magnetically Induced Disk Winds and Transport in the HL Tau Disk”. In: *ApJ* 845.1, p. 31.
- Hawley, J. F. (1987). “Non-linear evolution of a non-axisymmetric disc instability”. In: *MNRAS* 225, pp. 677–694.
- (1991). “Three-dimensional simulations of black hole tori”. In: *ApJ* 381, pp. 496–507.
- (2001). “Global Magnetohydrodynamic Simulations of Cylindrical Keplerian Disks”. In: *ApJ* 554, pp. 534–547.
- Hayasaki, K., N. Stone, and A. Loeb (2013). “Finite, intense accretion bursts from tidal disruption of stars on bound orbits”. In: *MNRAS* 434, pp. 909–924.
- Heil, L. M., P. Uttley, and M. Klein-Wolt (2015). “Inclination-dependent spectral and timing properties in transient black hole X-ray binaries”. In: *MNRAS* 448, pp. 3348–3353.
- Hjellming, R. M. and M. P. Rupen (1995). “Episodic ejection of relativistic jets by the X-ray transient GRO J1655 - 40”. In: *Nature* 375, pp. 464–468.
- Homan, J. et al. (2001). “Correlated X-Ray Spectral and Timing Behavior of the Black Hole Candidate XTE J1550-564: A New Interpretation of Black Hole States”. In: *ApJS* 132, pp. 377–402.
- Ichimaru, S. (1977). “Bimodal behavior of accretion disks - Theory and application to Cygnus X-1 transitions”. In: *ApJ* 214, pp. 840–855.
- Ingram, A. and C. Done (2011). “A physical model for the continuum variability and quasi-periodic oscillation in accreting black holes”. In: *MNRAS* 415, pp. 2323–2335.
- Ingram, A., C. Done, and P. C. Fragile (2009). “Low-frequency quasi-periodic oscillations spectra and Lense-Thirring precession”. In: *MNRAS* 397, pp. L101–L105.
- Ingram, A. and S. Motta (2014). “Solutions to the relativistic precession model”. In: *MNRAS* 444, pp. 2065–2070.
- Ingram, A. and M. van der Klis (2013). “An exact analytic treatment of propagating mass accretion rate fluctuations in X-ray binaries”. In: *MNRAS* 434, pp. 1476–1485.
- (2015). “Phase-resolved spectroscopy of low-frequency quasi-periodic oscillations in GRS 1915+105”. In: *MNRAS* 446, pp. 3516–3525.

- Ingram, A. et al. (2016). “A quasi-periodic modulation of the iron line centroid energy in the black hole binary H1743-322”. In: *MNRAS* 461, pp. 1967–1980.
- Ivanov, P. B. and A. F. Illarionov (1997). “The oscillatory shape of the stationary twisted disc around a Kerr black hole”. In: *MNRAS* 285, pp. 394–402.
- Jonker, P. G. and G. Nelemans (2004). “The distances to Galactic low-mass X-ray binaries: consequences for black hole luminosities and kicks”. In: *MNRAS* 354, pp. 355–366.
- Joss, P. C., Y. Avni, and S. Rappaport (1978). “Accreting neutron stars in highly compact binary systems and the nature of 3U 1626-67”. In: *ApJ* 221, pp. 645–651.
- Jubelgas, M., V. Springel, and K. Dolag (2004). “Thermal conduction in cosmological SPH simulations”. In: *MNRAS* 351, pp. 423–435.
- Kalemci, E. et al. (2001). “The X-ray variability and quasi-periodic oscillations in XTE J1550-564 during the decay of the 2000 outburst”. In: *American Astronomical Society Meeting Abstracts #198*. Vol. 33. Bulletin of the American Astronomical Society, p. 802.
- Kato, S. (1990). “Trapped One-Armed Corrugation Waves and QPOS”. In: *Publications of the Astronomical Society Of Japan* 42.1, pp. 99–113.
- King, A. R., J. E. Pringle, and M. Livio (2007). “Accretion disc viscosity: how big is alpha?” In: *MNRAS* 376, pp. 1740–1746.
- King, A. R. et al. (2013). “Accretion disc viscosity: what do warped discs tell us?” In: *MNRAS* 431, pp. 2655–2660.
- Kinney, A. L. et al. (2000). “Jet Directions in Seyfert Galaxies”. In: *ApJ* 537, pp. 152–177.
- Kiuchi, K. et al. (2011). “Gravitational Waves from the Papaloizou-Pringle Instability in Black-Hole-Torus Systems”. In: *Physical Review Letters* 106.25, 251102, p. 251102.
- Kluźniak, W. et al. (2005). “QPOs in cataclysmic variables and in X-ray binaries”. In: *A&A* 440.2, pp. L25–L28.
- Komossa, S. (2015). “Tidal disruption of stars by supermassive black holes: Status of observations”. In: *Journal of High Energy Astrophysics* 7, pp. 148–157.
- Krolik, J. H. and J. F. Hawley (2015). “A Steady-state Alignment Front in an Accretion Disk Subjected to Lense-thirring Torques”. In: *ApJ* 806, 141, p. 141.
- Kumar, S. and J. E. Pringle (1985). “Twisted accretion discs - The Bardeen-Petterson effect”. In: *MNRAS* 213, pp. 435–442.
- Lacy, J. H., C. H. Townes, and D. J. Hollenbach (1982). “The nature of the central parsec of the Galaxy”. In: *ApJ* 262, pp. 120–134.
- Lai, D. (1999). “Magnetically Driven Warping, Precession, and Resonances in Accretion Disks”. In: *ApJ* 524, pp. 1030–1047.

- Laibe, G. and D. J. Price (2014). “Dusty gas with one fluid in smoothed particle hydrodynamics”. In: *MNRAS* 440, pp. 2147–2163.
- Larwood, J. (1998). “On the precession of accretion discs in X-ray binaries”. In: *MNRAS* 299, pp. L32–L36.
- Larwood, J. D. et al. (1996). “The tidally induced warping, precession and truncation of accretion discs in binary systems: three-dimensional simulations”. In: *MNRAS* 282, pp. 597–613.
- Lense, J. and H. Thirring (1918). “Über den Einfluß der Eigenrotation der Zentralkörper auf die Bewegung der Planeten und Monde nach der Einsteinschen Gravitationstheorie”. In: *Physikalische Zeitschrift* 19, p. 156.
- Li, Z. B. et al. (2013). “The energy dependence of the centroid frequency of the low-frequency quasi-periodic oscillations in XTE J1550-564”. In: *MNRAS* 428, pp. 1704–1714.
- Liptai, D. et al. (2017). “Does turbulence determine the initial mass function?” In: *MNRAS* 465, pp. 105–110.
- Liska, M. et al. (2017). “Formation of Precessing Jets by Tilted Black-hole Discs in 3D General Relativistic MHD Simulations”. In: *ArXiv e-prints*.
- Livio, M. and H. C. Spruit (1991). “On the mechanism of angular momentum transport in accretion disks”. In: *A&A* 252, pp. 189–192.
- Lodato, G. (2008). “Classical disc physics”. In: *New Astro. Rev.* 52, pp. 21–41.
- Lodato, G., A. R. King, and J. E. Pringle (2009). “Stellar disruption by a supermassive black hole: is the light curve really proportional to $t^{-5/3}$?” In: *MNRAS* 392, pp. 332–340.
- Lodato, G. and D. J. Price (2010). “On the diffusive propagation of warps in thin accretion discs”. In: *MNRAS* 405, pp. 1212–1226.
- Lodato, G. and J. E. Pringle (2006). “The evolution of misaligned accretion discs and spinning black holes”. In: *MNRAS* 368, pp. 1196–1208.
- (2007). “Warp diffusion in accretion discs: a numerical investigation”. In: *MNRAS* 381, pp. 1287–1300.
- Lodato, G. and E. M. Rossi (2011). “Multiband light curves of tidal disruption events”. In: *MNRAS* 410, pp. 359–367.
- Loeb, A. and A. Ulmer (1997). “Optical Appearance of the Debris of a Star Disrupted by a Massive Black Hole”. In: *ApJ* 489, pp. 573–578.
- Lomb, N. R. (1976). “Least-squares frequency analysis of unequally spaced data”. In: *Astrophysics and Space Science* 39, pp. 447–462.

- Lubow, S. H. and G. I. Ogilvie (2000). “On the Tilting of Protostellar Disks by Resonant Tidal Effects”. In: *ApJ* 538, pp. 326–340.
- Lubow, S. H., G. I. Ogilvie, and J. E. Pringle (2002). “The evolution of a warped disc around a Kerr black hole”. In: *MNRAS* 337.2, pp. 706–712.
- Lucas, W. E. et al. (2013). “Misaligned streamers around a Galactic Centre black hole from a single cloud’s infall”. In: *MNRAS* 433, pp. 353–365.
- Lucy, L. B. (1977). “A numerical approach to the testing of the fission hypothesis”. In: *AJ* 82, pp. 1013–1024.
- Lynden-Bell, D. and J. E. Pringle (1974). “The evolution of viscous discs and the origin of the nebular variables.” In: *MNRAS* 168, pp. 603–637.
- Maccarone, T. J. (2002). “On the misalignment of jets in microquasars”. In: *MNRAS* 336, pp. 1371–1376.
- Martin, R. G., C. A. Tout, and J. E. Pringle (2010). “Supernova kicks and misaligned microquasars”. In: *MNRAS* 401, pp. 1514–1520.
- Martin, R. G. et al. (2014a). “Giant Outbursts in Be/X-Ray Binaries”. In: *ApJL* 790, L34, p. L34.
- Martin, R. G. et al. (2014b). “The Kozai-Lidov Mechanism in Hydrodynamical Disks”. In: *ApJL* 792, L33, p. L33.
- Martin, R. G. et al. (2016). “Planet-disc evolution and the formation of Kozai-Lidov planets”. In: *MNRAS* 458, pp. 4345–4353.
- McCabe, C. et al. (2011). “Spatially Resolving the HK Tau B Edge-on Disk from 1.2 to 4.7 μm : A Unique Scattered Light Disk”. In: *ApJ* 727, 90, p. 90.
- Meru, F. and M. R. Bate (2011). “Non-convergence of the critical cooling time-scale for fragmentation of self-gravitating discs”. In: *MNRAS* 411, pp. L1–L5.
- (2012). “On the convergence of the critical cooling time-scale for the fragmentation of self-gravitating discs”. In: *MNRAS* 427, pp. 2022–2046.
- Mewes, V. et al. (2016). “Numerical relativity simulations of thick accretion disks around tilted Kerr black holes”. In: *Phys. Rev. D* 93.6, 064055, p. 064055.
- Meyer, F. and E. Meyer-Hofmeister (1984). “Outbursts in dwarf novae accretion disks”. In: *A&A* 132, pp. 143–150.
- Miller, J. M. et al. (2001). “High-Frequency Quasi-Periodic Oscillations in the 2000 Outburst of the Galactic Microquasar XTE J1550-564”. In: *ApJ* 563, pp. 928–933.
- Monaghan, J. J. (1989). “On the problem of penetration in particle methods”. In: *J. Comput. Phys.* 82.1, pp. 1–15.

- Monaghan, J. J. (1992). “Smoothed Particle Hydrodynamics”. In: *Annual Review of Astronomy and Astrophysics* 30.1, pp. 543–574.
- (1997). “SPH and Riemann Solvers”. In: *Journal of Computational Physics* 136, pp. 298–307.
- (2005). “Smoothed particle hydrodynamics”. In: *Reports on Progress in Physics* 68, pp. 1703–1759.
- (2012). “Smoothed particle hydrodynamics and its diverse applications”. In: *Annual Review of Fluid Mechanics* 44, pp. 323–346.
- Monaghan, J. J. and J. C. Lattanzio (1985). “A refined particle method for astrophysical problems”. In: *A&A* 149, pp. 135–143.
- Morales Teixeira, D. et al. (2014). “Conservative GRMHD Simulations of Moderately Thin, Tilted Accretion Disks”. In: *ApJ* 796, 103, p. 103.
- Morgan, E. H., R. A. Remillard, and J. Greiner (1997). “RXTE Observations of QPOs in the Black Hole Candidate GRS 1915+105”. In: *ApJ* 482, pp. 993–1010.
- Morris, J. P. (1996). “A study of the stability properties of smooth particle hydrodynamics”. In: *PASA* 13, pp. 97–102.
- Morris, J. P. and J. J. Monaghan (1997). “A switch to reduce SPH viscosity”. In: *Journal of Computational Physics* 136.1, pp. 41–50.
- Motta, S. E. (2016). “Quasi periodic oscillations in black hole binaries”. In: *Astronomische Nachrichten* 337, p. 398.
- Motta, S. E. et al. (2014a). “Black hole spin measurements through the relativistic precession model: XTE J1550-564”. In: *MNRAS* 439, pp. L65–L69.
- Motta, S. E. et al. (2014b). “Precise mass and spin measurements for a stellar-mass black hole through X-ray timing: the case of GRO J1655-40”. In: *MNRAS* 437, pp. 2554–2565.
- Motta, S. E. et al. (2015). “Geometrical constraints on the origin of timing signals from black holes”. In: *MNRAS* 447, pp. 2059–2072.
- Motta, S. E. et al. (2017). “Links between quasi-periodic oscillations and accretion states in neutron star low-mass X-ray binaries”. In: *MNRAS* 468, pp. 2311–2324.
- Muno, M. P. et al. (2001). “Radio Emission and the Timing Properties of the Hard X-Ray State of GRS 1915+105”. In: *ApJ* 556, pp. 515–532.
- Murray, J. R. (1996). “SPH simulations of tidally unstable accretion discs in cataclysmic variables”. In: *MNRAS* 279, pp. 402–414.
- Narayan, R. and E. Quataert (2005). “Black Hole Accretion”. In: *Science* 307, pp. 77–80.

- Nealon, R., D. J. Price, and C. J. Nixon (2015). “On the Bardeen-Petterson effect in black hole accretion discs”. In: *MNRAS* 448, pp. 1526–1540.
- Nealon, R. et al. (2016). “Apsidal precession, disc breaking and viscosity in warped discs”. In: *MNRAS* 455, pp. L62–L66.
- Nelson, R. P. and J. C. B. Papaloizou (2000). “Hydrodynamic simulations of the Bardeen-Petterson effect”. In: *MNRAS* 315, pp. 570–586.
- Nespoli, E. et al. (2003). “A transient variable 6 Hz QPO from GX 339-4”. In: *A&A* 412, pp. 235–240.
- Nixon, C. (2015). “Accretion disc viscosity: a limit on the anisotropy”. In: *MNRAS* 450, pp. 2459–2465.
- Nixon, C. and A. King (2015). “Warp propagation in astrophysical discs”. In: *ArXiv e-prints*.
- Nixon, C., A. King, and D. Price (2013). “Tearing up the disc: misaligned accretion on to a binary”. In: *MNRAS* 434, pp. 1946–1954.
- Nixon, C. and G. Salvesen (2014). “A physical model for state transitions in black hole X-ray binaries”. In: *MNRAS* 437, pp. 3994–3999.
- Nixon, C. et al. (2012). “Tearing up the Disk: How Black Holes Accrete”. In: *ApJL* 757, L24, p. L24.
- Nixon, C. J. and A. R. King (2012). “Broken discs: warp propagation in accretion discs”. In: *MNRAS* 421, pp. 1201–1208.
- Nixon, C. J., A. R. King, and D. J. Price (2012). “Rapid AGN accretion from counter-rotating discs”. In: *MNRAS* 422, pp. 2547–2552.
- Ogilvie, G. I. (1999). “The non-linear fluid dynamics of a warped accretion disc”. In: *MNRAS* 304, pp. 557–578.
- (2003). “On the dynamics of magnetorotational turbulent stresses”. In: *MNRAS* 340, pp. 969–982.
- (2006). “Non-linear bending waves in Keplerian accretion discs”. In: *MNRAS* 365, pp. 977–990.
- Ogilvie, G. I. and G. Dubus (2001). “Precessing warped accretion discs in X-ray binaries”. In: *MNRAS* 320, pp. 485–503.
- Orosz, J. A. et al. (2001). “Dynamical Evidence for a Black Hole in the Microquasar XTE J1550-564”. In: *The Astronomer’s Telegram* 70.
- Papaloizou, J. C. B. and D. N. C. Lin (1995). “On the dynamics of warped accretion disks”. In: *ApJ* 438, pp. 841–851.

- Papaloizou, J. C. B. and J. E. Pringle (1983). “The time-dependence of non-planar accretion discs”. In: *MNRAS* 202, pp. 1181–1194.
- (1984). “The dynamical stability of differentially rotating discs with constant specific angular momentum”. In: *MNRAS* 208, pp. 721–750.
- (1985). “The dynamical stability of differentially rotating discs. II”. In: *MNRAS* 213, pp. 799–820.
- Papaloizou, J. C. B. and C. Terquem (1995). “On the dynamics of tilted discs around young stars”. In: *MNRAS* 274, pp. 987–1001.
- Patterson, J., E. L. Robinson, and R. E. Nather (1977). “The 71 Second Oscillation of DQ Herculis”. In: *Bulletin of the American Astronomical Society*. Vol. 9. Bulletin of the American Astronomical Society, p. 557.
- Phinney, E. S. (1989). “Manifestations of a Massive Black Hole in the Galactic Center”. In: *The Center of the Galaxy*. Ed. by M. Morris. Vol. 136. IAU Symposium, p. 543.
- Piran, T. et al. (2015). “Disk Formation Versus Disk Accretion — What Powers Tidal Disruption Events?” In: *ApJ* 806, 164, p. 164.
- Poutanen, J., J. H. Krolik, and F. Ryde (1997). “The nature of spectral transitions in accreting black holes - The case of CYG X-1”. In: *MNRAS* 292, pp. L21–L25.
- Poutanen, J. et al. (2007). “Supercritically accreting stellar mass black holes as ultraluminous X-ray sources”. In: *MNRAS* 377, pp. 1187–1194.
- Press, W. H. et al. (1993). *Numerical Recipes in FORTRAN; The Art of Scientific Computing*. 2nd. New York, NY, USA: Cambridge University Press.
- Price, D. J. (2007). “splash: An Interactive Visualisation Tool for Smoothed Particle Hydrodynamics Simulations”. In: *PASA* 24, pp. 159–173.
- (2012). “Smoothed particle hydrodynamics and magnetohydrodynamics”. In: *J. Comput. Phys.* 231, pp. 759–794.
- Price, D. J. and M. R. Bate (2007). “The impact of magnetic fields on single and binary star formation”. In: *MNRAS* 377, pp. 77–90.
- Price, D. J. and C. Federrath (2010). “A comparison between grid and particle methods on the statistics of driven, supersonic, isothermal turbulence”. In: *MNRAS* 406, pp. 1659–1674.
- Price, D. J. and J. J. Monaghan (2004). “Smoothed Particle Magnetohydrodynamics - I. Algorithm and tests in one dimension”. In: *MNRAS* 348, pp. 123–138.
- (2005). “Smoothed Particle Magnetohydrodynamics - III. Multidimensional tests and the $\nabla \cdot B = 0$ constraint”. In: *MNRAS* 364, pp. 384–406.

- Price, D. J. et al. (2017). “Phantom: A smoothed particle hydrodynamics and magnetohydrodynamics code for astrophysics”. In: *ArXiv e-prints*.
- Pringle, J. E. (1981). “Accretion discs in astrophysics”. In: *ARA&A* 19, pp. 137–162.
- (1992). “A simple approach to the evolution of twisted accretion discs”. In: *MNRAS* 258, pp. 811–818.
- (1996). “Self-induced warping of accretion discs”. In: *MNRAS* 281, pp. 357–361.
- Pringle, J. E. and A. King (2014). *Astrophysical Flows*.
- Rees, M. J. (1988). “Tidal disruption of stars by black holes of 10 to the 6th–10 to the 8th solar masses in nearby galaxies”. In: *Nature* 333, pp. 523–528.
- Reis, Ruben C et al. (2012). “A 200-second quasi-periodicity after the tidal disruption of a star by a dormant black hole”. In: *Science* 337.6097, pp. 949–951.
- Remillard, R. A. and J. E. McClintock (2006). “X-Ray Properties of Black-Hole Binaries”. In: *ARA&A* 44, pp. 49–92.
- Remillard, R. A. et al. (1999). “X-Ray Nova XTE J1550-564: Discovery of a Quasi-periodic Oscillation near 185 HZ”. In: *ApJL* 517, pp. L127–L130.
- Remillard, R. A. et al. (2002a). “Characterizing the Quasi-periodic Oscillation Behavior of the X-Ray Nova XTE J1550-564”. In: *ApJ* 564, pp. 962–973.
- Remillard, R. A. et al. (2002b). “Evidence for Harmonic Relationships in the High-Frequency Quasi-periodic Oscillations of XTE J1550-564 and GRO J1655-40”. In: *ApJ* 580, pp. 1030–1042.
- Robertson, B. E. et al. (2010). “Computational Eulerian hydrodynamics and Galilean invariance”. In: *MNRAS* 401, pp. 2463–2476.
- Rodriguez, J. et al. (2004). “An X-Ray Timing Study of XTE J1550-564: Evolution of the Low-Frequency Quasi-periodic Oscillations for the Complete 2000 Outburst”. In: *ApJ* 612, pp. 1018–1025.
- Saitoh, T. R. et al. (2012). “Flaring up of the Compact Cloud G2 during the Close Encounter with Sgr A* in Summer 2013”. In: *arXiv:1212.0349*.
- Scargle, J. D. (1982). “Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data”. In: *ApJ* 263, pp. 835–853.
- Schnittman, J. D., J. Homan, and J. M. Miller (2006). “A Precessing Ring Model for Low-Frequency Quasi-periodic Oscillations”. In: *ApJ* 642, pp. 420–426.
- Shafee, R. et al. (2006). “Estimating the Spin of Stellar-Mass Black Holes by Spectral Fitting of the X-Ray Continuum”. In: *ApJL* 636, pp. L113–L116.

- Shakura, N.I. and R.A. Sunyaev (1973). “Black Holes in Binary-Systems - Observational Appearance”. In: *A&A* 24.3, pp. 337–355.
- Shen, R.-F. and C. D. Matzner (2014). “Evolution of Accretion Disks in Tidal Disruption Events”. In: *ApJ* 784, 87, p. 87.
- Shiokawa, H. et al. (2015). “General Relativistic Hydrodynamic Simulation of Accretion Flow from a Stellar Tidal Disruption”. In: *ApJ* 804, 85, p. 85.
- Smak, J. (1984). “Accretion in cataclysmic binaries. IV - Accretion disks in dwarf novae”. In: *Acta Astronomica* 34, pp. 161–189.
- Sobczak, G. J. et al. (2000). “Correlations between Low-Frequency Quasi-periodic Oscillations and Spectral Parameters in XTE J1550-564 and GRO J1655-40”. In: *ApJ* 531, pp. 537–545.
- Sorathia, K. A., J. H. Krolik, and J. F. Hawley (2013). “Magnetohydrodynamic Simulation of a Disk Subjected to Lense-Thirring Precession”. In: *ApJ* 777, 21, p. 21.
- Steiner, J. F. et al. (2011). “The spin of the black hole microquasar XTE J1550-564 via the continuum-fitting and Fe-line methods”. In: *MNRAS* 416, pp. 941–958.
- Stella, L. and M. Vietri (1998). “Lense-Thirring Precession and Quasi-periodic Oscillations in Low-Mass X-Ray Binaries”. In: *ApJL* 492, pp. L59–L62.
- (1999). “kHz Quasiperiodic Oscillations in Low-Mass X-Ray Binaries as Probes of General Relativity in the Strong-Field Regime”. In: *Physical Review Letters* 82, pp. 17–20.
- Strohmayer, T. E. (2001). “Discovery of a Second High-Frequency Quasi-periodic Oscillation from the Microquasar GRS 1915+105”. In: *ApJL* 554, pp. L169–L172.
- Strohmayer, T. E. et al. (2003). “X-Ray Bursts from the Accreting Millisecond Pulsar XTE J1814-338”. In: *ApJL* 596, pp. L67–L70.
- Strubbe, L. E. and E. Quataert (2009). “Optical flares from the tidal disruption of stars by massive black holes”. In: *MNRAS* 400, pp. 2070–2084.
- Tagger, M. and R. Pellat (1999). “An accretion-ejection instability in magnetized disks”. In: *A&A* 349, pp. 1003–1016.
- Titarchuk, L. and R. Fiorito (2004). “Spectral Index and Quasi-Periodic Oscillation Frequency Correlation in Black Hole Sources: Observational Evidence of Two Phases and Phase Transition in Black Holes”. In: *ApJ* 612, pp. 988–999.
- Ulmer, A. (1999). “Flares from the Tidal Disruption of Stars by Massive Black Holes”. In: *ApJ* 514, pp. 180–187.
- van der Klis, M. (2006). “Overview of QPOs in neutron-star low-mass X-ray binaries”. In: *Advances in Space Research* 38, pp. 2675–2679.

- van der Klis, M. et al. (1985). “Intensity-dependent quasi-periodic oscillations in the X-ray flux of GX5 - 1”. In: *Nature* 316, pp. 225–230.
- Varnière, P., M. Tagger, and J. Rodriguez (2012). “A possible interpretation for the apparent differences in LFQPO types in microquasars”. In: *A&A* 545, A40, A40.
- Verlet, L. (1967). “Computer “experiments” on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules”. In: *Physical review* 159.1, p. 98.
- Watkins, S. J. et al. (1996). “A new prescription for viscosity in smoothed particle hydrodynamics.” In: *A&AS* 119, pp. 177–187.
- Wen, F. and L. Dong (2011). “Papaloizou–Pringle instability of magnetized accretion tori”. In: *MNRAS* 410.3, pp. 1617–1630.
- Wendland, H. (1995). “Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree”. In: *Advances in computational Mathematics* 4.1, pp. 389–396.
- Wijnands, R., J. Homan, and M. van der Klis (1999). “The Complex Phase-Lag Behavior of the 3-12 HZ Quasi-Periodic Oscillations during the Very High State of XTE J1550-564”. In: *ApJL* 526, pp. L33–L36.
- Wurster, J., D. Price, and B. Ayliffe (2014). “Ambipolar diffusion in smoothed particle magnetohydrodynamics”. In: *MNRAS* 444, pp. 1104–1112.
- Wurster, J., D. J. Price, and M. R. Bate (2017). “The impact of non-ideal magnetohydrodynamics on binary star formation”. In: *MNRAS* 466, pp. 1788–1804.
- Zhuravlev, V. V. et al. (2014). “No Evidence for Bardeen-Petterson Alignment in GRMHD Simulations and Semi-analytic Models of Moderately Thin, Prograde, Tilted Accretion Disks”. In: *ApJ* 796, 104, p. 104.
- Zurek, W. H. and W. Benz (1986). “Redistribution of angular momentum by nonaxisymmetric instabilities in a thick accretion disk”. In: *ApJ* 308, pp. 123–133.