

Reflections upon black hole accretion discs: Advanced X-ray modelling and novel timing methods

Edward James Richard Nathan

St Hugh's College
University of Oxford

*A thesis submitted for the degree of
Doctor of Philosophy*

Trinity 2022

Abstract

The regions close to black holes (BHs) have the most extreme gravitational fields; therefore we are able to probe Einstein's theory of gravity by studying how matter accretes onto these objects. The best data we have of black hole accretion comes from X-ray binaries (XRBs) within the Milky Way, which shine brightly in X-rays. The accretion discs are very small and therefore are unresolved in observations, so we have to study the properties of the X-ray point source in great detail to understand the structure of, and processes within, the accretion system. A key observational feature we observe is the reflection of high energy X-ray photons incident onto the accretion disc, which imprint characteristic features onto the observed X-ray spectrum such as the Fe $K\alpha$ line. Relativistic distortions to the energies and arrival times of the reflected photons enable the mapping of the accretion region.

In this thesis I employ advanced techniques to model the reflection spectrum in observations of two BH XRBs, GRS 1915+105 and H1743-322. First, I study quasi-periodic oscillations (QPOs) in the observed X-ray flux from these systems. There are many theories to the origin of these signals, and these theories make different predictions of how the energy spectrum of the source is modulated over the course of a QPO cycle. I use novel Fourier domain timing techniques in order to phase-resolve the energy spectrum of the source over a QPO cycle, to look for these telltale modulation. I constrain modulations in the illumination of the accretion disc by the high energy X-ray photons, adding growing body of evidence that the QPOs have a geometrical origin.

I also present pioneering work testing a model able to constrain the mass-distance ratio of a BH, using the spectral shape and observed flux of the reflected X-rays to constrain the size and ionization state of the accretion disc. I performed a proof-of-principle test using observations of H1743-322, which I combined with a prior distance measurement to constrain the mass of the BH in this system for the first time. This demonstration shows how X-ray spectroscopy can be used to make mass measurements of the many BHs within the Milky Way which are too obscured for current optical techniques.

Reflections upon black hole accretion
discs: Advanced X-ray modelling and
novel timing methods



Edward James Richard Nathan
St Hugh's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2022

Acknowledgements

Personal

I wish to express my gratitude to my supervisor, Dr. Adam Ingram, whose enthusiasm and support are unmatched. He has encouraged me to explore every rabbit hole I have encountered, yet has been able to provide the much needed voice of reason when the time came to move on. I must thank him for every conversation we have had, the insightful, the fascinating, the fun. Much like a gold watch, he has passed to me a passion for X-ray astronomy which I know I shall carry for a very long time.

I also would like to thank Professor Rob Fender, along with every member of his research group at Oxford (no matter how transient), for keeping my mind open to the broad spectrum of ways in which to look at compact objects.

I interrupt this program with a shout-out Lauren, Andy, Sean, Laurence, Jamie, and David; a.k.a. Tower FM: a great group of friends to complete a DPhil with (or, indeed, to spend time talking and doing anything other than completing it). The last four years would never have been the same without you.

Likewise, to everyone at St. Hugh's and its MCR. Over the years I have been there, there have been far too many people to name whom I have had the privilege to call my friends. However, to the Joshs ('JC', and 'JT'), Amartya, Florence, Anwar, Alice, and Caroline, thank you for helping me finding myself at Oxford, albeit often at the pub. Also to Sophie, Imogen, Vedika, Jared, Daphné, Rebecca, and Sandra, who at every moment got my final year rocking (just like a rolling stone).

My time at Oxford has been smooth sailing, not least due to the committee and members of Oxford Sailing Club. With particular note to the 'Dream Team', thank you for being so welcoming and friendly as I rediscovered an old hobby.

I also cannot miss thanking all the teachers I have had. In particular, to my physics teachers Mr. Illingworth and Dr. Hollis. Without your support, guidance, and encouragement in my formative years, I may never have realised my passion.

Finally, to all my family, who have not only supported me at every moment, but also took a far more active role in my DPhil than they could have imagined when I returned home during the pandemic. This is for my Mum, who taught me everything she knows about physics. For my Dad, who has not only been my role-model, but also

a physical model for my future. For Poppy, my oldest friend. And for Mutley, who provides the best cuddles (and a government-approved reason to leave the house).

To everyone, thank you.

Institutional

I thank the Royal Society for providing the financial support for me to pursue the research within this thesis.

Abstract

The regions close to black holes (BHs) have the most extreme gravitational fields; therefore we are able to probe Einstein's theory of gravity by studying how matter accretes onto these objects. The best data we have of black hole accretion comes from X-ray binaries (XRBs) within the Milky Way, which shine brightly in X-rays. The accretion discs are very small and therefore are unresolved in observations, so we have to study the properties of the X-ray point source in great detail to understand the structure of, and processes within, the accretion system. A key observational feature we observe is the reflection of high energy X-ray photons incident onto the accretion disc, which imprint characteristic features onto the observed X-ray spectrum such as the Fe $K\alpha$ line. Relativistic distortions to the energies and arrival times of the reflected photons enable the mapping of the accretion region.

In this thesis I employ advanced techniques to model the reflection spectrum in observations of two BH XRBs, GRS 1915+105 and H1743-322. First, I study quasi-periodic oscillations (QPOs) in the observed X-ray flux from these systems. There are many theories to the origin of these signals, and these theories make different predictions of how the energy spectrum of the source is modulated over the course of a QPO cycle. I use novel Fourier domain timing techniques in order to phase-resolve the energy spectrum of the source over a QPO cycle, to look for these telltale modulation. I constrain modulations in the illumination of the accretion disc by the high energy X-ray photons, adding growing body of evidence that the QPOs have a geometrical origin.

I also present pioneering work testing a model able to constrain the mass-distance ratio of a BH, using the spectral shape and observed flux of the reflected X-rays to constrain the size and ionization state of the accretion disc. I performed a proof-of-principle test using observations of H1743-322, which I combined with a prior distance measurement to constrain the mass of the BH in this system for the first time. This demonstration shows how X-ray spectroscopy can be used to make mass measurements of the many BHs within the Milky Way which are too obscured for current optical techniques.

Contents

List of Abbreviations	xi
1 Introduction	1
1.1 Black Holes	2
1.2 X-ray binaries	4
1.3 Outbursts	6
1.4 X-ray Spectrum	8
1.4.1 Thermal disc emission	8
1.4.2 Compton scattering	11
1.4.3 Reflection	13
1.5 Spectral-timing states	20
1.6 QPO phenomenology	21
1.7 Models for the corona	24
1.7.1 Sandwich	25
1.7.2 Lamppost	26
1.7.3 Truncated disc	26
1.8 QPO models	27
1.8.1 The relativistic precession model	27
1.8.2 The precessing inner flow model	28
1.8.3 Corrugation modes	29
1.8.4 Accretion ejection instability	30
1.8.5 Propagating oscillatory shock	30
1.8.6 Pressure/accretion rate modes	31
1.9 Observational tests of QPO models	32
1.9.1 Inclination Dependence	32
1.9.2 Frequency-resolved spectroscopy	34
1.9.3 QPO Phase-resolved spectroscopy	35
1.10 Fourier methods	40
1.10.1 Power spectrum	41
1.10.2 Cross spectrum	43
1.10.3 Bi-spectrum	45

2	Phase-resolved spectroscopy of a quasi-periodic oscillation in the black hole X-ray binary GRS 1915+105 with <i>NICER</i> and <i>NuSTAR</i>	47
2.1	Introduction	48
2.2	Observations	49
2.2.1	<i>NuSTAR</i> Data Reduction	49
2.2.2	<i>NICER</i> Data Reduction	50
2.2.3	Energy spectrum	51
2.2.4	Power spectrum	52
2.2.5	Spectral Timing State	53
2.3	Phase Resolved Spectroscopy	53
2.3.1	QPO frequency tracking	56
2.3.2	Phase lag spectrum	58
2.3.3	Fractional RMS Spectrum	60
2.3.4	Phase difference between harmonics	61
2.3.5	Reconstructed Fourier Transformed Spectra	67
2.4	Theoretical Model	68
2.4.1	Corona	68
2.4.2	Disc	69
2.4.3	The complete model	74
2.5	Model Fits	74
2.5.1	Fitting procedure	74
2.5.2	Results	75
2.6	Discussion	82
2.6.1	Comparison with H1743-322	82
2.6.2	Asymmetric illumination profile	84
2.6.3	Inner Radius	86
2.6.4	Misalignment	88
2.6.5	Biases in the phase-averaged spectrum	90
2.6.6	Thermalisation lags	91
2.7	Conclusions	93
3	Proof-of-Principle mass measurement of the black hole in X-ray binary H1743-322	95
3.1	Introduction	96
3.2	Observations and Data Reduction	97
3.2.1	Observation Log	97
3.2.2	<i>NuSTAR</i> Data Reduction	97
3.2.3	<i>XMM-Newton</i> Data Reduction	97
3.2.4	Cross-Calibration	99

3.2.5	Timing Analysis	101
3.3	Spectral Fits	103
3.3.1	The model	103
3.3.2	Model fitting procedure	105
3.3.3	Results	107
3.4	Discussion	111
3.4.1	Proof-of-Principle	112
3.4.2	Mass Measurement	114
3.4.3	Higher Spin	114
3.4.4	Inner radius and QPOs	115
3.4.5	Density and inclination	116
3.5	Conclusions	117
4	Phase-resolved spectroscopy of quasi-periodic oscillations in H1743-	
	322 across three outbursts	119
4.1	Introduction	119
4.2	Data Analysis	120
4.2.1	Data selection and reduction	120
4.2.2	Phase-resolving method	121
4.2.3	Results	123
4.3	Phase-resolved fits	130
4.3.1	Fitting procedure	130
4.4	Discussion and Conclusions	137
5	Conclusions	141
Appendices		
A	Appendix of Chapter 2	147
A.1	Phase difference from Bispectrum	147
A.1.1	Treatment of Poisson noise in the Bispectrum	147
A.1.2	Jellyfish Bootstrapping	149
A.2	Model Details	149
A.2.1	Ray tracing	149
A.2.2	Model Normalisation	150
A.3	Phase offsets with different reference bands	150
A.4	MCMC corner plots of modulated parameters	152
B	Appendix of Chapter 4	157
B.1	Auto-bispectrum for epochs 14, 16b, and 18b	157
	Bibliography	161

List of Abbreviations

AEI	Accretion ejection instability
AGN	Active galactic nuclei
BH	Black hole
CENBOL	Centrifugal pressure supported boundary layer
DoF	Degree of Freedom
EPIC	European Photon Imaging Camera
FFT	Fast Fourier transform
FPMA / FPMB		Focal plane module A/B
FT	Fourier transform
FWHM	Full width at half maximum
HID	Hardness-intensity diagram
GRMHD	General relativistic magnetohydrodynamic
GTI	Good time interval
HRD	Hardness-RMS diagram
HSS	High soft state
ISCO	Innermost stable circular orbit
LHD	Low hard state
MCMC	Markov Chain Monte Carlo
MHD	Magnetohydrodynamic
<i>NICER</i>	The Neutron star Interior Composition Explorer
<i>NuSTAR</i>	The Nuclear Spectroscopic Telescope Array
NS	Neutron star
POS	Propagating oscillatory shock
QPO	Quasi-periodic oscillation
RGS	Reflection grating spectrometer

RMS	Root mean square
RPM	Relativistic precession model
<i>RXTE</i>	The Rossi X-ray Timing Explorer
SAS	Science analysis software
TCAF	Two component accretion flow
ULX	Ultraluminous X-ray source
XRB	X-ray binary

There is a theory which states that if ever anyone discovers exactly what the Universe is for and why it is here, it will instantly disappear and be replaced by something even more bizarre and inexplicable.

There is another theory which states that this has already happened.

— Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

1

Introduction

Contents

1.1	Black Holes	2
1.2	X-ray binaries	4
1.3	Outbursts	6
1.4	X-ray Spectrum	8
1.4.1	Thermal disc emission	8
1.4.2	Compton scattering	11
1.4.3	Reflection	13
1.5	Spectral-timing states	20
1.6	QPO phenomenology	21
1.7	Models for the corona	24
1.7.1	Sandwich	25
1.7.2	Lamppost	26
1.7.3	Truncated disc	26
1.8	QPO models	27
1.8.1	The relativistic precession model	27
1.8.2	The precessing inner flow model	28
1.8.3	Corrugation modes	29
1.8.4	Accretion ejection instability	30
1.8.5	Propagating oscillatory shock	30
1.8.6	Pressure/accretion rate modes	31
1.9	Observational tests of QPO models	32
1.9.1	Inclination Dependence	32
1.9.2	Frequency-resolved spectroscopy	34
1.9.3	QPO Phase-resolved spectroscopy	35
1.10	Fourier methods	40
1.10.1	Power spectrum	41
1.10.2	Cross spectrum	43
1.10.3	Bi-spectrum	45

1.1 Black Holes

A key prediction of general relativity is that objects of sufficient density collapse, forming a Black Hole (BH). While much theoretical work had occurred over the 20th century, most famously but not exclusively by Schwarzschild (1916) and Kerr (1963), it wasn't until Webster & Murdin (1972) when the first observational evidence of a potential (now confirmed) BH was reported in Cygnus X-1. As the defining feature of a BH is the *event horizon*, the boundary at which the Newtonian escape velocity is the speed of light, we can only infer the presence of the BH from its effect on its surroundings.

The first metric introduced for BHs was the Schwarzschild metric, which is spherically symmetric, treating BHs as a singular point mass (Schwarzschild 1916), with mass M . In this metric, the event horizon is a sphere with radius $R_H = 2GM/c^2$, or twice the *gravitational radius* $R_g = GM/c^2$. A further critical radius is the *innermost stable circular orbit* (ISCO), the smallest radius that a test-mass can circularly orbit the BH in a stable manner. While circular orbits are possible at smaller radii, they are unstable and any small deviation would destroy the orbit, likely sending it towards the event horizon. For a Schwarzschild BH, the ISCO lies at $R_{\text{ISCO}} = 6R_g$.

Kerr (1963) relaxed the spherical symmetry of the Schwarzschild metric, asserting only azimuthal symmetry. This led to the Kerr metric which accounts for a BH with angular momentum J_{BH} , with spin parameter $a = J_{\text{BH}}/(cM)$. The spin parameter is often given in a dimensionless form $a_* = cJ_{\text{BH}}/(GM^2)$. In this metric, the event horizon within the equatorial plane is at $R_H = (1 + \sqrt{1 - a_*^2})R_g$. When considering the ISCO, it is necessary to account for whether the orbit is in the same direction as the BH spin (prograde orbits), or in the opposite direction to the spin (retrograde orbits). Again, staying within the equatorial plane of the BH,

the spin parameter is defined as $0 < a < 1$ for prograde orbits, but $-1 < a < 0$ for retrograde. The ISCO then lies at (Bardeen et al. 1972)

$$\begin{aligned} R_{\text{ISCO}} &= \left(3 + Z_2 - \frac{a_*}{|a_*|} \sqrt{(3 - Z_1)(3 + Z_1 + 2Z_2)} \right) R_g, \text{ where} \\ Z_1 &= 1 + \sqrt[3]{1 - a_*^2} \left(\sqrt[3]{1 + a_*} + \sqrt[3]{1 - a_*} \right), \text{ and} \\ Z_2 &= \sqrt{3a_*^2 + Z_1^2}. \end{aligned} \quad (1.1)$$

For simple cases, this gives

$$R_{\text{ISCO}} = \begin{cases} 1 R_g, & a = 1 \quad (\text{maximal prograde spin}); \\ 6 R_g, & a = 0 \quad (\text{Schwarzschild}); \\ 9 R_g, & a = -1 \quad (\text{maximal retrograde spin}). \end{cases} \quad (1.2)$$

While the Kerr metric has been expanded upon to include BHs with non-zero electrical charge (Newman et al. 1965), it is not expected that astrophysical BHs would have a significant charge as it would attract neutralising material¹.

As the mass and spin are enough to describe any astrophysical BH, the question follows: what distribution of BH mass and spins exist in the Universe? BHs have been observed which fall into two categories of mass ranges: super-massive BHs which are expected to exist at the centre of every galaxy, with masses in the range $\sim 10^5 - 10^{11} M_\odot$ (e.g. Peterson et al. 2004); and stellar-mass BHs, the remnants left behind after the death of the most massive stars (possibly after some merger processes), with masses $\sim 3 - 100 M_\odot$ (Corral-Santana et al. 2016; Wang et al. 2021). Populations of BHs with mass in other ranges have been theorised, such as primordial BHs ($\ll 1M_\odot$ Villanueva-Domingo et al. 2021)) and intermediate BHs ($\sim 10^2 - 10^5 M_\odot$ Greene et al. 2020), however as yet there is no significant evidence of their existence. Conversely, the spin distribution of BHs is still an active area research, and is an important indicator of their formation and history (Middleton 2016; Reynolds 2021).

Both super-massive BHs and stellar mass BHs have been seen to accrete. The former are seen in galaxies across the Universe as active galactic nuclei (AGN), and

¹While, in theory, a BH could exist with a non-zero charge away from neutralising material, we would not be able to observe it as we rely on the material around a BH to infer its presence. Besides, this thesis is about accreting BHs!

accrete gas and dust from their immediate surroundings. The latter are seen in X-ray binary systems (XRBs), which are the focus of this thesis.

1.2 X-ray binaries

BH XRBs are systems with a BH in a binary orbit with a companion star, which supplies the gas that the BH accretes. The accreting material is heated to such high temperatures close to the BH that it glows brightly in X-rays, giving rise to their name. XRBs also exist where the accretor is instead a neutron star (NS), an extremely dense stellar remnant held up against gravitational collapse by degenerate neutron pressure. These systems are complicated by the solid surface and magnetic field of the NS. In this thesis I only focus on BH XRBs, in order to study simple accretion and general relativity. Moreover, I focus on BH XRBs within the Milky Way, which provide the highest quality data necessary for the advanced techniques I will be applying. XRBs are also observed in other galaxies, most notably ultraluminous X-ray sources (ULXs Middleton et al. 2012), some of which potentially contain a BH, with others having been shown to be powered by a NS accretor (Bachetti et al. 2014).

The effective ‘Roche’ potential of a binary system, which accounts for both gravitational attraction and angular momentum, is shown in Fig. 1.1. The ‘figure-of-eight’ equipotential depicted by the bold line divides the potential into two *Roche lobes*. Within each Roche lobe, the potential is dominated by the gravitational field of one of the bodies. At the L_1 Lagrangian point where the two Roche-Lobes meet, the gravitational force from the two bodies cancel. If the companion star becomes large enough (e.g., after undergoing expansion towards the end of its life) for it to fill its own Roche-Lobe, matter can ‘spill over’ through the L_1 point into the Roche-Lobe of the BH, a process known as *Roche lobe overflow*. Some systems with a high-mass companion star can have a powerful enough stellar wind to allow significant mass transfer without overflow; however the systems studied within this thesis are powered by Roche lobe overflow.




Figure 1.1: The Roche potential of a binary system, shown in terms of both lines of equipotential, and a 3D representation of the potential. Three Lagrange points (L_1 , L_2 , and L_3) are labelled, which are Saddle points in the potential. The system shown has a 2:1 mass ratio, with the larger mass located between L_1 and L_3 . Image from van der Sluys (2006).

Despite falling into the gravitational well of the BH, the gas is unable to directly accrete due to the orbital angular momentum it carries. Therefore it forms an accretion disc around the BH, in which the majority of the material slowly spirals inwards, converting the free energy of shear to radiation with high efficiency. The gravitational potential energy liberated by this slow inward spiral, and thus the luminosity available to radiate away, is proportional to the mass accretion rate \dot{m} . If the disc is very dense, and therefore very optically thick, it will be in local thermal equilibrium such that the blackbody temperature at a given disc radius is $T \propto \dot{m}^{1/4}$.

There are currently 68 known transients in the Milky Way thought to be BH XRBs, 19 of which have a dynamical mass measurement consistent with the existence of the BH (Corral-Santana 2021; Corral-Santana et al. 2016), with the rest being BH candidates. Fig. 1.2 is a cumulative histogram showing the number




Figure 1.2: A cumulative histogram showing the number of discovered BH transients (with those which have been dynamically confirmed in blue). This figure is taken from Corral-Santana (2021), which is an updated version of the work presented in Corral-Santana et al. (2016). The horizontal black bars indicate the years of operation of the X-ray missions which have been responsible for their discovery.

of discovered BH transients over time, including the years of operation of the X-ray missions responsible for their discovery.

1.3 Outbursts

Most XRBs are transient sources, which spend most of their time in quiescence during which they are either not or barely detectable. However, they occasionally undergo outbursts, during which their accretion rate, and thus X-ray luminosity, increases by at least an order of magnitude for a period of a few weeks to years. Fig. 1.3 shows the X-ray light curve from *RXTE* (the Rossi X-ray Timing Explorer) monitoring of the BH XRB GX 339-4 as it goes through three outbursts between 2002 and 2008. The dramatic increase in accretion rate that gives rise to an outburst




Figure 1.3: *RXTE* monitoring of the BH XRB GX 339-4 which shows three outbursts in 2002, 2004, and 2007 (in red, green, and blue respectively). Taken from Fig. 1 of Plant et al. (2014).

is thought to result from the *hydrogen ionisation instability*.

When the mass accretion rate \dot{m} through the disc is low during quiescence, the disc's temperature $T \propto \dot{m}^{1/4}$ is low and thus hydrogen within the disc is almost entirely unionised (Cannizzo & Reiff 1992). However, a small fluctuation in the mass accretion rate can push a small fraction of the photons in the Wien tail of the Planck function to energies over 13.6 eV: the hydrogen ionisation energy. As these photons are absorbed within the disc, this increases the energy within and therefore temperature of the local disc annulus. The increase in T causes even more photons to have enough energy to be able to ionise the hydrogen, and therefore be absorbed themselves, causing further heating. This instability causes the \dot{m} of the original annulus to be greater than the input mass accretion (of either the next annulus out, or from the companion). As the mass flowing out of one annulus propagates through the disc into the next inner annulus, this slightly inner annulus also begins to heat and triggers the instability there. This causes a heating wave to flow through the disc, as the instability is triggered in subsequently inner radii (Frank et al. 2002; Kato 1998; Cannizzo 1993; Lasota 2001).

Eventually, the balance of the incoming and outgoing mass accretion rates cause the original annulus to be drained of matter. Thus the density of the gas within the annulus decreases, and so is the gas pressure since $P \propto T$. As the temperature of the gas at the annulus decreases, fewer photons are able to ionise the hydrogen. In turn this leads to runaway cooling as the hydrogen recombines, and lowers the mass

flow rate. In the converse way as above, this recombination causes a cooling wave to propagate down the disc, bringing the outburst to an end (Cannizzo 1993) with a mass flow rate below that of the input level. This allows the density to build up again until the cycle is able to repeat (King & Ritter 1998; Cannizzo 1993; Lasota 2001).

1.4 X-ray Spectrum

1.4.1 Thermal disc emission

The most straightforward form of structured accretion is through an *accretion disc*, of which the simplest model is a geometrically thin, optically thick disc, where the orbital motion of the accreting material is much greater than the slower inwards drift towards the BH (Shakura & Sunyaev 1973; Novikov & Thorne 1973).

Consider a small amount of material of mass m , orbiting a BH of mass M , at a radius R . The energy associated with this material is

$$E - J\Omega, \quad (1.3)$$

where $E = E_{\text{kinetic}} + E_{\text{grav}}$ is the sum of the kinetic and gravitational potential energies, and Ω and $J = mr^2\Omega$ are the orbital velocity and angular momentum of the material. Here, term $J\Omega$ results from the work done on the material through shear forces within the disc. Within Newtonian gravity, the gravitational binding energy of the material which has originated from a far distance is straightforwardly

$$E_{\text{grav}} = -\frac{GMm}{R}. \quad (1.4)$$

As material in a Keplerian orbit has the orbital velocity

$$\Omega = \sqrt{\frac{GM}{r^3}}, \quad (1.5)$$

therefore the material has the kinetic energy

$$E_{\text{kinetic}} = \frac{GMm}{2R}. \quad (1.6)$$




Figure 1.4: A schematic diagram of an accretion disc getting hotter at smaller radii. This sketches how many blackbody spectra with a black body spectrum with peak temperature $T \propto r^{-4/3}$ can combine to give an overall ‘disc’ spectra.

Thus, it follows that by the time the material has reached the radius R the material has total energy

$$E = E_{\text{kinetic}} + E_{\text{grav}} - mr^2\Omega^2 = -\frac{3GMm}{2R}, \quad (1.7)$$

with matter losing energy as it moves a small radius dR

$$dE = \frac{3GMm}{2R^2}dR. \quad (1.8)$$

Consequently, the total energy lost by the material accreting from very far way down to a disc inner radius R_{in} is

$$E_{\text{lost}} = \frac{3R_{\text{g}}}{2R_{\text{in}}}c^2. \quad (1.9)$$

Assuming that this energy loss by the matter is radiated away, for an accretion rate \dot{m} the luminosity of a disc extending down to the BH ISCO is

$$L = E_{\text{lost}} = \eta\dot{m}c^2, \quad (1.10)$$

where the efficiency $\eta = \frac{R_{\text{g}}}{2R_{\text{in}}}$ is $\eta = \frac{1}{12}$ for a Schwartzchild BH, and up to $\eta = \frac{1}{2}$ for a maximally spinning BH with a prograde disc.

As we are considering an optically thick disc, on average the emitted photons would undergo many scatterings within the disc, therefore producing a thermal,

blackbody spectrum. However, as more energy is liberated from the disc at smaller radii, the disc is hotter closer to the BH. Thinking back again to mass accreting as it moves over a small dR , the energy released is radiated with a blackbody distribution and luminosity $dL(R) = \sigma_{\text{SB}}T^4(R)dA$, where $dA = 4\pi R dR$ is the area of the annulus from which it is radiating², and σ_{SB} is the Stefan-Boltzmann constant. Again, on the assumption that all the released energy goes into the thermal heating of the disc $dL = dE$, this can be combined with Eq. 1.8

$$\frac{3GM\dot{m}}{2R^2}dR = \sigma_{\text{SB}}T^4(R)4\pi R dR \quad (1.11)$$

and rearranged, showing that the temperature structure of the disc is

$$\sigma_{\text{SB}}T^4(R) = \frac{3GM\dot{m}}{8\pi R^3}. \quad (1.12)$$

The overall observed spectrum from the disc is therefore a ‘multi-coloured’ blackbody spectrum. This is illustrated in Fig. 1.4, with a sketch of a disc getting hotter at inner radii on the left, and with an example sum of a number of increasingly hotter blackbody distributions on the right, to illustrate the approximate shape of a thermal disc spectrum. BH XRBs are sometimes observed to have an X-ray spectrum dominated by such a multi-coloured blackbody disc component, with some examples shown in the left subfigure of Fig. 1.5.

Eq. 1.12 is slightly modified when considering the conservation of angular momentum, alongside the presence of an inner radius R_{in} of the disc within which there is no viscosity: the stress free inner boundary condition. This makes the disc temperature profile (e.g. Frank et al. 2002)

$$T^4(R) = \frac{3GM\dot{m}}{8\pi R^3} \left(1 - \sqrt{\frac{R_{\text{in}}}{R}}\right). \quad (1.13)$$

While this was derived in the Newtonian regime, the relativistic corrections are relatively small (Novikov & Thorne 1973). Equations 1.12 and 1.13 come from the two extreme assumptions about the stress at the inner boundary of the disc.

²This is twice the traditional area of an annulus, since the disc radiates from both the top and the bottom sides.




Figure 1.5: Left: An example of a number of disc dominated spectra of the BH XRB GRO 1655-40, using *RXTE* PCA and *HEXTE* data. From Fig. 15 of Done et al. (2007) Right: Hard state spectrum of the BH XRB GX 339-4 using *RXTE* PCA and *HEXTE* data. Adapted from Fig. 8 of Plant et al. (2014).

As expected, intermediate assumptions of the inner boundary condition provide intermediate distributions (Mummery & Balbus 2019a,b).

The disc spectrum is dependent on the exact position of the inner radius. At times, when the disc is thought to extend all the way down to R_{ISCO} , the disc spectrum can therefore be used to constrain the spin of the BH, provided the black hole mass M and the mass accretion rate \dot{m} are known. While \dot{m} can be inferred from the total disc flux, this also depends on knowing the inclination angle from which we view the disc and the distance to the system from Earth (for an example, see Steiner et al. 2012).

1.4.2 Compton scattering

Even the hottest disc spectra observed from XRBs peak at energies $\lesssim 1$ keV. However, often an accompanying spectral component of much higher energy photons is observed alongside the disc component, often dominating over it (examples of which are shown on the right in Fig. 1.5). This component is thought to result from Compton scattering of comparatively cool photons, perhaps originally from the disc, with a population of hot electrons (Sunyaev & Truemper 1979) in a region commonly referred to as the *corona*. Provided the corona is reasonably optically thin (with optical depth $\tau \sim 1 - 3$), on average photons will scatter several times before

escaping, but not enough times to reach thermal equilibrium with the electrons. Since the electrons have greater energy than the photons, this process (Compton up-scattering) transfers energy to the photons, producing high energy photons while simultaneously slightly cooling the corona.

When photons of frequency ν (and energy ϵ) scatter off electrons of velocity βc (and Lorentz factor $\gamma = (1 - \beta^2)^{-1/2}$), the average fractional energy transfer when averaged over all scattering angles is (Longair 2011; Rybicki & Lightman 1979)

$$\left\langle \frac{\Delta\epsilon}{\epsilon} \right\rangle = \frac{4}{3}\beta^2\gamma^2 - \frac{h\nu}{m_e c^2}. \quad (1.14)$$

The first term in this equation is classical Thomson scattering, and the second term is quantum recoil.

In the optically thin corona, if the electrons are in thermal equilibrium with a temperature T_e they have a random velocity from thermal motion given by a Maxwellian distribution with an average velocity of $v^2 \sim 3k_B T_e/m_e$. By defining the dimensionless constant $\Theta = \frac{k_B T_e}{m_e c^2}$, we get a typical $\beta^2 = 3\Theta$. Inserting this into Eq. 1.14 (for temperatures where the electrons are non-relativistic, therefore $\gamma \approx 1$), we get

$$\left\langle \frac{\Delta\epsilon}{\epsilon} \right\rangle \approx 4\Theta - \frac{h\nu}{m_e c^2}. \quad (1.15)$$

Therefore, any scattering where $4k_B T \gg h\nu$ on average imparts a fractional energy gain of 4Θ . The probability of a photon scattering is $1 - e^{-\tau}$, for a corona with an optical depth τ . While corona have a typical optical depth $\tau \sim 1 - 3$, as a simpler illustration we can consider an extremely optically thin corona with $\tau \ll 1$, which reduces the scattering probability to $\sim \tau$. Starting off with a seed population of soft X-ray photons with energy ϵ_{in} , a fraction τ of them will undergo this scattering to an energy $(1 + 4\Theta)\epsilon_{\text{in}}$. However, some ($\sim \tau$) of these up-scattered photons will scatter again, and this can continue until the photons have reached energies $h\nu \sim 4k_B T$ and then can no longer gain energy. As, before the photon energy reaches this maximum, each subsequent scattering affects the same proportion

of photons, and also gives the same fractional energy increase, the resulting flux spectrum is a power law $F(\epsilon) \propto \epsilon^{-\alpha}$ with the index (e.g. Done 2010)

$$\alpha \approx \frac{\log \tau}{\log(1 + 4\Theta)}, \quad (1.16)$$

thus the differential photon spectrum is $N(\epsilon) = \epsilon^{-1}F(\epsilon) \propto \epsilon^{-\Gamma}$, where $\Gamma = \alpha + 1$ is the *photon index*. This spectrum has a low-energy cut-off at the seed photon energy $\sim \epsilon_{\text{in}}$, and a high-energy cut-off related to the electron temperature at $\sim 3\Theta$ where the recoil effect becomes important for these photons which have had many scatterings.

1.4.3 Reflection

Most of the hard X-ray photons upscattered by the corona will escape the system, potentially making it to an observer. However, some of these photons will be incident back onto the accretion disc. These photons undergo reprocessing within the optically thick disc, before leaving (potentially having been scattered back into the observer's line of sight). These photons form what is known as the *reflection* spectrum (Guilbert & Rees 1988; Lightman & White 1988; Ross & Fabian 2005; García & Kallman 2010).

Photons incident on the disc interact with electrons within the disc, either scattering off free electrons or being absorbed by bound electrons. The relative probability of scattering or absorption depends on the relative abundance of the two populations, and the interaction cross-sections of the two processes.

1.4.3.1 Neutral Reflection

The simplest assumption we can make is that all of the bound atoms that photons interact with are neutral atoms, which is a good assumption for when the incident photon flux is low enough such that the produced ions always have time to recombine before the next photo-ionisation event. The cross-section for photoelectric absorption in this case is shown by the black line in Fig. 1.6 (left). We see that it decreases with photon energy as $\sim E^{-3}$, with edges at the ionisation potentials of each bound level of each element. In contrast, the cross-section for scattering off free electrons




Figure 1.6: Left: The photoelectric absorption cross-section of neutral material with solar abundance, along with the electron scattering cross-section. Right: The reflection spectrum seen from this same material. It shows how the reflection decreases with the increasing photoelectric absorption, including the absorption edges, along with the fluorescence lines. At the higher energies, the Compton hump dominates the reflection spectrum. Taken from Fig. 1.20 of Done (2010).

(grey line) is approximately constant for $E \lesssim 100$ keV before slowly dropping off for higher energies due to relativistic effects. This means that for photon energies $E \lesssim 10$ keV the photoelectric absorption dominates, so the resulting reflection spectrum increases with energy and includes absorption edges (blue line in Fig. 1.6, right). At higher energies photons are instead more likely to be scattered. However, as this scattering is not elastic, hard X-ray photons are down-scattered to produce a Compton hump peaking at $\sim 20 - 50$ keV (Fig. 1.6, right).

The photoelectric absorption of low-energy photons leads to emission lines being present in the ‘reflection’ spectrum, which can also be seen in Fig 1.6(right). Photons (with the required energy) are most likely to be absorbed by the most tightly bound electrons in an atom (say element X), those in the K ($n=1$) shell (Hussein 2007), liberating the electron and in the process leaving behind an excited X^+ ion. The vacancy left can then be filled by an electron in a higher orbital; for a vacancy in the K shell, an electron is most likely to come from the next shell (L, $n=2$) if it is occupied. When this higher energy electron de-excites to the lower energy K shell it does so by either emitting a photon, or another electron via the Auger effect. The photons emitted through this process form fluorescence lines, often




Figure 1.7: Left: The photoelectric absorption cross-section of ionised material with solar abundance, with increasing values of the ionisation parameter ξ (0 in black; 0.1 in pink; 1 in cyan; 10 in red; 100 in green; and 1000 in blue), along with the electron scattering cross-section. Right: The reflection spectra seen from this same material, for the same ionisations (in the same colours). It shows how the reflection decreases with the increasing photoelectric absorption, including the absorption edges, although the fluorescence lines are not included. At the higher energies, the Compton hump dominates the reflection spectrum, and is not affected by the ionisation state of the disc. Taken from Fig. 1.21 of Done (2010).

labelled in Siegbahn notation (where $K\alpha$ represents the $n=2$ to $n=1$ transition, $K\beta$ the $n=3$ to $n=1$ transition, and so on). As atoms with a lower atomic number are more likely to de-excite via the Auger effect (suppressing the fluorescence line emission), and as the energy of the Iron $K\alpha$ line at ~ 6.4 keV is in a region with a high reflection probability, this leads to this line being the dominant line in the observed reflection spectrum. As the accreting material in an XRB has a stellar origin, elements heavier than iron have very low abundance and therefore have an almost negligible effect on the observed reflection spectrum.

1.4.3.2 Ionised Reflection

In BH XRB systems, the inner regions of the disc receive high levels of irradiating flux. This means that not only is the observed reflection spectrum dominated by these regions, but also many ions will not have time to recombine before the next photo-ionization event. It is therefore necessary to account for the disc medium consisting of a superposition of different ionic species. The ionisation state of the

disc is often parameterised by the ionisation parameter (Tarter et al. 1969; Ross & Fabian 1993; Kallman & Bautista 2001; García et al. 2013a)

$$\xi = 4\pi \frac{F_x}{n_e}, \quad (1.17)$$

where n_e is the electron density within the disc, and F_x is the illuminating X-ray flux. A common choice energy band for the flux is the 1 Ry to 1000 Ry (≈ 13.6 eV to 13.6 keV) band (e.g. Kallman & Bautista 2001), with the lower bound corresponding to the hydrogen ionisation energy and the upper bound being a convenient choice where the photoelectric absorption cross-section is greatly diminished.

As the ionisation of the disc material increases, the relative proportion of ions to neutral atoms within the disc increases. As ions have a smaller absorption cross-section this in turn decreases the overall absorption cross-section of the material (and therefore increases the amount of reflection), particularly at lower energies. The higher ionisation also causes the resulting absorption edges to increase in energy, as it requires more energy to liberate an electron from a negative ion than it does a neutral atom. Fig. 1.7 (left) shows how the photoelectric absorption cross-section changes for $0 \leq \xi \leq 1000$. The resulting reflection spectrum for the same values of ξ , calculated using a simple model that only considers absorption and scattering (the PEXRIV model in XSPEC) is shown in Fig. 1.7 (right).

Beyond the changes in the photoelectric cross-section, the ionization state of the disc also influences the fluorescence lines, which increase in energy for higher ionic species. The fluorescence lines are also become broader, as the irradiation primarily effects the outer layer of the disc, causing these outer levels to be more ionised. As the more ionised material has a smaller absorption cross-section, more photons are able to penetrate this layer towards the less ionised material inside the disc, and it is these layers that produce the fluorescence photons which we see. However, in order to be observed these fluorescence photons have to pass through the outer layer and escape. This upper layer, being highly ionised, contains many free electrons so that the photons are subjected to Compton scattering, causing broadening of the fluorescence lines (Young et al. 2001). When the disc is entirely ionised, the




Figure 1.8: Three example spectra computed with XILLVER for (from bottom to top) $\log \xi = 0$ in blue, 2.0 in black, and 3.6 in red. The other model parameters used to calculate these were spectral index $\Gamma = 2.0$, coronal electron temperature $kT_e = 300$ keV, solar iron abundance, and a 30° observing inclination. Each spectrum was offset by a factor of 10 to separate it clearly from the one below.

reflection spectrum has almost no features, becoming close to the illuminating spectrum, up to the energy of the Compton hump.

The resulting reflection spectrum calculated taking all of the discussed effects into account is shown in Fig. 1.8, for ionisation parameters $\log \xi = 0, 2.0, 3.6$ to cover a range from completely unionised to very highly ionised. This uses the most advanced reflection model available at the moment, XILLVER (García & Kallman 2010; García et al. 2013a), which self-consistently calculates the temperature and ionisation state in the disc atmosphere then solves the radiative transfer equation for a slab of constant density. The overarching spectral shape changes as before, however in the XILLVER models the broadening and weakening of the spectral lines becomes obvious.

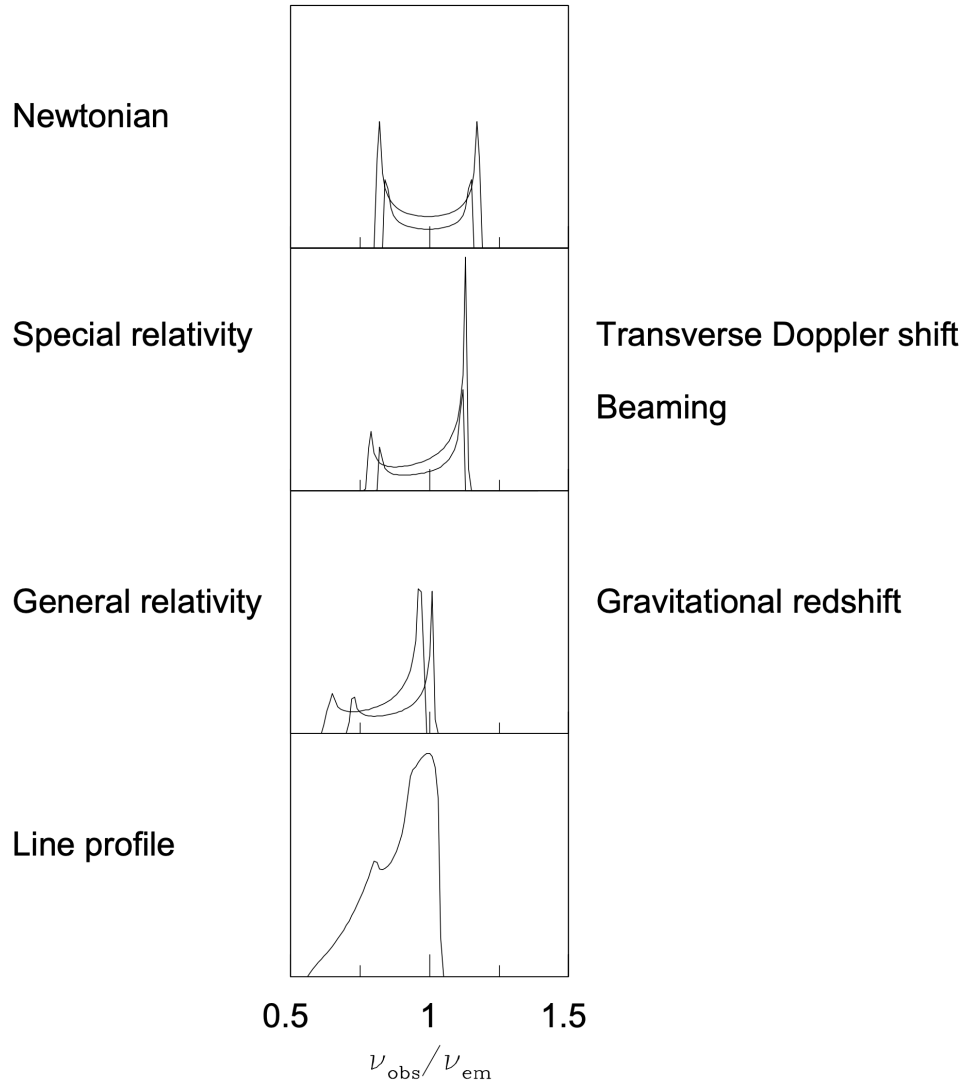


Figure 1.9: The effects on the light emitted from the disc from two narrow radii. The upper panel has no relativistic effects, and only considers Newtonian Doppler shifts. The second panel adds in special relativistic effects: transverse Doppler shifting, and relativistic beaming. The third panel shows the gravitational redshift from general relativity. The final, lowest panel shows the final broadened and shifted line profile from the emission as integrated over the disc. Image adapted from Fig. 3 of Fabian et al. (2000).

1.4.3.3 Relativistic smearing

The discussion of the reflection spectrum thus far has only considered the spectrum in the local restframe of the disc. From our, observers, point of view, we see the disc in a highly relativistic regime (Fabian et al. 1989), particularly the areas of the disc closest to the BH where it is expected the disc illumination (and

therefore reflection) will be the strongest. A number of effects are at play, and these are demonstrated in the upper three panels of Fig. 1.9 given a δ -function spectral line, considering two separate disc radii. The fourth panel shows the overall smearing observed from the disc.

1. The first effect is the ‘Newtonian’ Doppler shifting. We see the light from a patch of disc approaching us blue-shifted to higher energies, whereas light from a patch of disc receding from us is red-shifted to lower energies. Annuli of the disc closer to the BH have faster orbital velocities, and therefore the most significant broadening.
2. The next consideration is special relativistic effects due to the highly relativistic motion close to the BH, which come in two forms. There is an additional red-shift from transverse Doppler shifts, due to the time dilation of rapidly moving material. There is also beaming of the light from the approaching material, causing the the blue-wing of the line profile to be boosted, and anti-beaming of the receding material causing suppression of the red-wing. Again, this is strongest for the fastest material, closest to the BH.
3. Finally, the material close to the black hole is in a strong gravitational field, so we see the light from that region of the disc undergoing gravitational redshift as described by general relativity – an effect strongest for the material closest to the BH, producing a stronger redshift for the innermost annuli.

The reflection spectrum that we observe originates from all parts of the disc, and the strength of each of these effects depends on the radii *in units of* R_g of the patch of the disc that the X-rays are coming from. Therefore pure reflection spectroscopy is independent of the BH mass, however it does depend on the disc inner radius (again, in units of R_g), the inclination angle we observe the disc from; and the emissivity profile which weights the reflection spectrum to different radii.




Figure 1.10: An example hardness intensity diagram (HID; upper) and hardness RMS diagram (HRD; lower), sketching out a typical outburst which proceeds from the point marked ‘A’, moving around the ‘q’ shape anti-clockwise. The corresponding points are also labelled on the HRD. Imagine from Fig. 1 of Kylafis & Belloni (2015).

1.5 Spectral-timing states

During an outburst an XRB exhibits a number of different spectral and timing behaviours, which are classified into different states. These states are typically classified by the source’s location on the hardness-intensity diagram (HID) and hardness-RMS diagram (HRD) (Belloni 2010; Fender & Muñoz-Darias 2016; Ingram & Motta 2019; Belloni & Motta 2016), where RMS means root-mean-squared variability amplitude. Fig. 1.10 shows a schematic of the HID (upper) and HRD (lower). Hardness is not a uniquely defined property, but all ways of measurement ensure it is a comparison of the amount of hard, high energy X-ray flux ($\gtrsim 10$ keV) to soft, low-energy X-ray flux ($\lesssim 5$ keV). Intensity is simply the overall flux observed³, and RMS is the total variability present in the source light curve.

When an XRB comes out of quiescence, it begins to rise up the right-hand side

³Therefore, care must be taken when comparing sources at different distances.

of the HID. The X-ray spectrum is dominated by a strong coronal power law with a hard photon index of $1.6 - 1.7$, such as on the right subfigure of Fig. 1.5. In this low hard state (LHS) the light curve typically shows a large amount of intrinsic variability (often with $\gtrsim 10\%$ RMS). In this state the total RMS within the light curve is seen to decrease as the flux increases, and a compact radio jet is often seen with radio luminosity that correlates with the X-ray luminosity (Gallo et al. 2003).

Eventually the system reaches its maximal brightness, however it still has a hard spectrum (sitting around the top-right part of the ‘q’ of the HID). It then begins to transition through the intermediate state, with the source softening despite maintaining its brightness. Here, the disc spectrum becomes hotter and brighter, while the power law spectral index increases.

As the source moves through the intermediate state, the continuous jet emission stops, and is instead characterised by discrete ejecta. The discrete jet ejection appears to occur only for a very narrow range in X-ray hardness, which is referred to as the ‘jet line’ (Fender et al. 2004). Occasionally, a source never leaves the hard state (or only briefly enters the intermediate state) before returning to quiescence, in what is known as a *failed outburst* (Alabarta et al. 2021).

After the intermediate state, the source ends up in the high soft state (HSS) located to the left part of the HID, with the spectrum dominated by a strong, hot disc spectrum in the soft X-rays, although the power law does not vanish completely (see the black spectrum on the left subfigure of Fig. 1.5). The HSS is also seen to have low X-ray variability, which is consistent with most of the variability having come from the power law. In this state, jets are quenched.

After spending some time in the HSS, the source fades in brightness, before transitioning back through the intermediate state to the LHS, eventually returning to quiescence typically a few weeks or months after the outburst began.

1.6 QPO phenomenology

Quasi-periodic oscillations (QPOs) are often seen in the X-ray flux from XRBs. When studied in the power spectrum (see 1.10.1) they appear as narrow peaks,




Figure 1.11: Examples of the three types of low frequency QPOs seen in the power spectra from BH XRBs. The top panel shows a type-A QPO seen from XTE J1859+226, while the middle and lower panels show respectively a type-B and type-C QPO seen from GX 339-4. Figure adapted from Fig. 1 of Ingram & Motta (2019).

with a width $\lesssim 1/2$ of the centroid frequency. In BH XRB systems, QPOs are classified as either being low frequency with a centroid frequency $\lesssim 30$ Hz, or high frequency with a centroid frequency $\gtrsim 60$ Hz (Belloni 2010; Ingram & Motta 2019). Low frequency QPOs are then sub-divided into three categories: type-A, -B, and -C; Fig. 1.11 shows an example of each type.

QPOs are typically modelled in the power spectrum for frequency ν with the Lorentzian function (van Straaten et al. 2002)

$$L(\nu) = \frac{a_0^2 \Delta}{\left[\frac{\pi}{2} + \arctan\left(\frac{\nu_0}{\Delta}\right) \right] \left[\Delta^2 + (\nu - \nu_0)^2 \right]}, \quad (1.18)$$

where ν_0 is the QPO centroid frequency, and Δ is the half width at half maximum

(HWHM), such that the quality factor is $Q = \frac{\nu_0}{2\Delta}$. $L(\nu)$ is defined such that the overall RMS power of the QPO is a_0 (therefore $\int_0^\infty L(\nu)d\nu = a_0^2$).

Beginning at the end of the LHS and persisting a little way into the intermediate state, type-C low frequency QPOs are often observed. They are very strong (with up to $\sim 20\%$ RMS) and narrow ($Q \gtrsim 8$), and can range from a centroid frequency of $\sim 0.1 - 20$ Hz. These QPOs are seen accompanied by broadband noise (i.e. strong RMS on a wide range of timescales, characterised in the power spectrum high power between two break frequencies). As a source makes its way through the LHS into the intermediate state type-C QPOs are seen to increase in frequency, and properties such as the disc flux and the photon index are seen to correlate with the QPO frequency (Vignarca et al. 2003). As type-C QPOs are very strong, they are commonly observed with a 2nd harmonic with a frequency twice that of the fundamental, (and sometimes even higher harmonics are detectable).

During the intermediate state, sometimes a type-B QPO is observed, defined as a QPO with very low associated broad band noise. The presence of a type-B QPO is used to split the intermediate state into the *soft intermediate state* (when a type-B QPO is present), and the *hard intermediate state* (without a type-B QPO) (Belloni & Motta 2016). These QPOs are known to be associated with jet launching (Fender et al. 2004; Fender et al. 2009), with a switch from a type-C to a type-B QPO having been seen to correspond to the launching of apparent superluminal ejecta in recent outbursts of the BH XRBs MAXI J1535-571 (Russell et al. 2019) and MAXI J1820+070 (Homan et al. 2020). Type-B QPOs are typically less strong ($\lesssim 5\%$ RMS) and less sharply peaked ($Q \gtrsim 6$) than type-C and they have typical frequencies in the narrow range of $\sim 5 - 6$ Hz. Often, a weak second harmonic accompanies the fundamental.

The least common type of low frequency QPO observed in BH XRBs is type-A, which are seen shortly after the XRB has entered the soft state. They show as a broad ($Q \lesssim 3$) peak with a centroid frequency of 6-8 Hz; they are also quite weak with RMS around only a few percent, and occur alongside very weak red noise.




Figure 1.12: Four examples of Coronal geometries around a black hole, where the optically thick accretion disc is sketched in grey, and the population of hot electrons (the corona) is in yellow. The top left sketches the lamppost model (although it is assumed the structure also has a symmetric component below the disc too). The top right sketch shows the sandwich corona geometry. The lower two sketches are possible configurations of the truncated disc model, with the corona within a disc truncation radius. Figure taken from Fig. 8.1 of Bambi (2017).

As type-C QPOs are both the most commonly observed form of low frequency QPO, and are also show strongest variability, I will focus on studying these in this thesis.

1.7 Models for the corona

While the corona has long been known to be a population of hot electrons close to the BH, the exact nature and geometry of the corona is still a matter of hot academic debate despite decades of study. There are a number of geometries to describe the corona which are used in the literature; I will present a summary of the key models here.

XRBS are not resolved in observations and so only appear as point sources. Therefore, to understand the accretion geometry we have to use spectral and variability information, which encode the key processes and interactions of the disc and corona. The exact geometry of the corona has implications for the seed

photons from the disc which are up-scattered. Different geometries predict different fractions of the disc photons being intercepted and up-scattered by the corona, and likewise the illumination pattern on the disc by these photons. This illumination of the disc by the corona changes the overall reflection spectrum we see, and so is encapsulated by the *emissivity profile* $\epsilon(r)$. This emissivity profile is defined by the ‘reflected’ emission coming from the disc (Laor 1991); however, because the disc is optically thick, the reflection happens locally and so the emissivity profile should match the illumination profile.

1.7.1 Sandwich

The sandwich, or slab, coronal geometry has the disc embedded within the hot electrons (see upper-right of Fig. 1.12), which are themselves held close to the accretion disc via magnetic fields (Haardt & Maraschi 1991, 1993; Field & Rogers 1993; Stern et al. 1995). In this geometry, it is difficult to explain the hard spectrum observed from the Comptonised component, since a sandwich corona as described will intercept a large luminosity of disc photons, keeping the electron temperature low and the power-law index of the emitted spectrum above ~ 2 (Dove et al. 1997; Poutanen & Coppi 1998; Malzac et al. 2005).

The sandwich geometry can be modified to try to rectify this issue. For instance perhaps the corona is patchy rather than homogeneous (Haardt et al. 1994), possibly being comprised of individual magnetic flares. As the disc isn’t entirely covered, fewer photons enter the corona, so there is a smaller cooling effect, allowing the harder spectra that we observe. Likewise, the electrons in the corona could have a bulk velocity flow away from the disc, reducing the disc flux intercepted by the corona in its rest frame due to Doppler de-boosting. However, both of these effects this would also reduce the contribution of reflection to the observed spectrum to an extent that is difficult to reconcile with the observational data (Done & Nayakshin 2007).

1.7.2 Lamppost

The lamppost geometry (Matt et al. 1991; Martocchia & Matt 1996) considers the coronal emission coming from a spot at a height h directly above the BH with respect to the plane of the disc (with an assumed symmetrical counterpart below the disc too), as illustrated in the top-left of Fig. 1.12. This model is not entirely physical, as it is implausible that the emitting corona is localised to one small spot, however it is mathematically convenient, as the the radial emissivity profile can then be calculated through ray-tracing.

Lamppost models have usually allowed a ‘reflection fraction’ to vary to allow the proportion of reflected light compared to the directly observed coronal light to be a model parameter while otherwise maintaining isotropic emission. However, recent versions of lamppost models (such as in Ingram et al. (2022), which I use in chapter 3) include terms that allow the emission to be non-isotropic.

While this model is just a convenient approximation, it does possibly represent an over-density at the base of an outflowing jet (Miyamoto & Kitamoto 1991; Fender et al. 1999; Markoff et al. 2005; Kara et al. 2019), it is simply the first approximation of a compact corona that sits above the disc. Recent work has considered a ‘double lamppost’ model to represent a vertically extended corona (e.g. Chainakun & Young 2017), which is used to explain a common feature where timing methods imply a higher lamppost reflecting off larger disc radii, whereas spectral methods indicate strong reflection from very small radii.

1.7.3 Truncated disc

The truncated disc model (Eardley et al. 1975; Ichimaru 1977; Done et al. 2007) describes a situation where (in the hard and intermediate states) the disc does not extend all the way down to the ISCO of the BH, but is instead truncated at some larger radius (such as the lower two sketches of Fig. 1.12). Within this truncation radius the disc density is very low due to the rapid inflow velocity of the accreting matter, despite the same mass accretion rate. This means the disc has low opacity and also a low electron-proton collision rate. As the protons are

much more massive than the electrons, they gain most energy through accretion. However, it is the electrons that are dominant in Compton scattering as they have a much higher scattering cross section. Therefore the electron temperature is set by the balance of the weak coupling with the hotter protons and the radiative cooling, with the thermal energy of the protons advecting into the BH. Such an accretion flow is known as an advection dominated accretion flow (ADAF; Ichimaru 1977; Rees et al. 1982; Narayan & Yi 1995).

As the XRB enters the soft state, the accretion rate increases and the disc moves further inwards. The resulting greater covering fraction of the disc by the corona means more seed photons are up-scattered, which in turn cools the electrons in the corona.

1.8 QPO models

There are many models in the literature that try to explain the origin of low frequency QPOs, particularly type-C QPOs. Here, I will attempt to summarise key models that are favoured to explain the observed properties (see section 1.9).

Most models relate the QPO to three key frequencies of motion within the the disc. The first is the orbital frequency of the disc $\nu_\phi(r)$; the second is the vertical epicyclic frequency $\nu_\theta(r)$; and the third is the radial epicyclic frequency $\nu_r(r)$ (see e.g. Kato 1998; Nowak & Lehr 1998). While in Newtonian gravity these three frequencies are the same, within general relativity this is not true and they can differ significantly in extreme-gravity (for example, when close to a BH). It follows that an element on an elliptical orbit would have its semi-major axis precess at the periastron frequency $\nu_{\text{per}} = \nu_\phi - \nu_r$. Within the Kerr metric, an orbit tilted out of the plane of a spinning mass would undergo nodal precession with frequency $\nu_{\text{nod}} = \nu_\phi - \nu_\theta$, also known as Lense-Thirring precession $\nu_{\text{LT}} = \nu_{\text{nod}}$ (Lense & Thirring 1918).

1.8.1 The relativistic precession model

The relativistic precession model (RPM; Stella & Vietri 1998; Stella et al. 1999) suggests that type-C low frequencies QPOs are caused by Lense-Thirring precession

at some characteristic radius. This radius is plausibly the disc truncation radius r_{tr} if the disc does not extend down to the ISCO. If there are hot-spots (but see Psaltis & Norman (2000) for an alternative modulation mechanism) at this characteristic radius⁴ on orbits which are inclined with respect to the spin of the BH, these orbits would precess at the Lense-Thirring frequency. The origin of the inclination of the disc with respect to the BH spin axis could likely be a result of the misalignment of the plane of the binary orbit. Such a misalignment could initially result from an asymmetric supernova at the birth of the BH (Fragos et al. 2010). Although dissipative processes will eventually align the system, this will only be after the BH accretes $\gtrsim 10\%$ of its initial mass (King & Nixon 2016), meaning that a natal misalignment could be preserved in most BH XRBs.

The RPM was also designed to explain the presence of two high frequency QPOs, with frequencies at ν_{per} and ν_{ϕ} respectively. In the rare occurrences that all three QPOs are observed in the same observation (assuming they originate at the same radius, which is not necessarily the disc inner radius) then it is possible to use these three frequencies to solve for the radius of origin, along with the mass and spin of the BH (Motta et al. 2014a; Ingram & Motta 2014). Likewise, it is possible to combine the frequencies of two QPOs with prior knowledge of the BH mass to find the spin and radius, (Motta et al. 2014b; Ingram & Motta 2014). Schnittman et al. (2006) considered the effect of Lense-Thirring precession on a ring of material orbiting the BH, rather than simply a hot-spot.

1.8.2 The precessing inner flow model

Ingram et al. (2009) presented a model where the QPO is also caused by Lense-Thirring precession, however here it is a hot inner flow that is precessing within a truncated disc. If the truncated disc is misaligned with the spin axis of the BH, with a misalignment angle β between the BH spin axis and the normal vector to the disc, it would feed the inner flow such that it could undergo the nodal precession. The model assumes that the inner flow precesses as a solid body which has an outer

⁴These hot-spots could be generated by turbulence if the radius is the truncation radius.

radius set at the disc truncation radius, and either propagates down to the ISCO, or a radius where the density drops substantially (which could be greater for BHs with higher spin). The precession is such that it keeps the misalignment angle β between the inner flow and the BH spin axis constant, however this means that the misalignment between the inner flow and the truncated disc varies from 0 to 2β over the course of a precession cycle. A moderate misalignment of $\beta \sim 10 - 15^\circ$ would lead to QPOs of the typical observed RMS (see e.g. Veledina et al. 2013; Ingram & van der Klis 2015; Ingram et al. 2017). Recent high-resolution general relativistic magnetohydrodynamic simulations have seen thick inner flows, and even the associated jet, precessing due to the Lense-Thirring effect (Fragile et al. 2007; Liska et al. 2018, 2019).

The effect of the precession of a hot inner flow depends on the exact observational angle to the system. When the inner flow is seen most edge-on the Doppler boosting from it is maximal (with the Doppler blue shifts being dominant over the red shifts). Conversely, when the inner flow is most face-on, while the Doppler boosting is minimal the projected area is maximal (Veledina et al. 2013; Ingram & van der Klis 2015). Also, as the inner flow is not spherically symmetric, the Comptonised emission is not isotropic, and the observed intensity and temperature of the inner flow could change as the viewing angle changes (Ingram et al. 2017; You et al. 2020).

A key prediction of this model is changes in the reflection spectrum over the course of the QPO cycle. As the inner flow precesses, it changes the emissivity profile on the disc. At any given time different parts of the disc would be preferentially illuminated. As we observe different parts of the disc to have different red-shifts, this modulated illumination pattern also modulates the shape of the observed reflection spectrum, causing a rocking of the iron $K\alpha$ line between higher and lower energies as blue- and red- shifted parts of the disc are illuminated (Ingram & Done 2012).

1.8.3 Corrugation modes

Corrugation modes (c-modes) are discoseismic, transverse standing waves trapped in a region close to the inner radius of the disc (Kato & Fukue 1980; Tsang & Lai 2009),

as visualised in Fig 2. of Tsang & Butsky (2013). The standing wave is set up from a white noise of initial perturbations via reflection off the disc inner radius r_{in} and the *inner vertical resonance* $r_{\text{IVR}} > r_{\text{in}}$. The resonant frequency of the standing wave is

$$\nu_c = m\nu_\phi(r_{\text{IVR}}) - j^{1/2}\nu_\theta(r_{\text{IVR}}), \quad (1.19)$$

where m and j are integers corresponding to different modes. This means that the fundamental mode $m = j = 1$ is the Lense-Thirring precession frequency of the IVR $\nu_c = \nu_\phi(r_{\text{IVR}}) - \nu_\theta(r_{\text{IVR}}) \equiv \nu_{\text{LT}}(r_{\text{IVR}})$. This model also predicts a QPO phase-dependent rocking of the iron line, but the effect is much more subtle than for the precession model (Tsang & Butsky 2013).

1.8.4 Accretion ejection instability

The Accretion Ejection Instability (AEI; Tagger & Pellat 1999; Varnière & Tagger 2002; Tagger et al. 2004; Tagger 2007) also describes normal modes of standing waves in the inner accretion disc. It can occur when a thin accretion disc is threaded by a strong vertical magnetic field, strong enough that the magnetic and gas pressure within the disc are comparable. Perturbations in the strength of the magnetic field trigger spiral density waves that travel outwards. Outward travelling waves with frequency ν reflect back at the inner Lindblad resonance $r_{\text{ilr}} > r_{\text{in}}$, which is the smaller of the two radii that satisfy the condition $(\nu - m\nu_\phi)^2 - \nu_r^2 = 0$, where the integer m is the mode number. The resulting inward travelling wave is in turn reflected at r_{in} until a standing wave is set up that rotates about the disc surface at some resonant eigenfrequency ν_i . The m^{th} mode corresponds to a spiral structure with m turns between the disc inner radius and the inner Lindblad resonance. The frequency of the spiral wave depends on the BH mass and spin, along with disc parameters such sound speed, and the strength of the vertical magnetic field.

1.8.5 Propagating oscillatory shock

The Propagating Oscillatory Shock (POS) model assumes a two component accretion flow (TCAF) which has a Keplerian disc, sandwiched by a sub-Keplerian corona

(Chakrabarti 1997). The POS model predicts that a shock forms at some radius r_{shk} where the centrifugal pressure in the Keplerian disc slows the accretion flow. At smaller radii, the flow is in the form of a sub-Keplerian halo referred to as a *centrifugal pressure supported boundary layer* (CENBOL), in which the protons and electrons may have different temperatures due to inefficient coupling between them.

The POS model suggests that QPOs are a result of the shock radius oscillating about its mean radius, which can occur when the cooling timescale of the halo is comparable with the infalling timescale of the accreting material.

1.8.6 Pressure/accretion rate modes

A class of QPO models consider oscillations which have a dynamical origin in some resonance of an intrinsic property of the accretion flow, such as pressure, density or accretion rate. For example, Cabanac et al. (2010) suggested a model of a pressure wave with a resonant frequency that is trapped within the corona, when assuming a truncated disc / hot inner flow geometry. O’Neill et al. (2011) notes QPOs in the accretion rate of a magnetohydrodynamic (MHD) simulation that they attributed to dynamo cycles in the azimuthal disc magnetic field. However, the QPO frequency in their simulation depended on radius, such that it is not clear how these accretion rate oscillations could give rise to a coherent QPO in the X-ray flux. Other results from MHD simulations include Wang et al. (2012), who explored toroidal Alfvén wave oscillations and Machida & Matsumoto (2008), who found an inner torus that deforms itself quasi-periodically between being circular and being in crescent shaped, which modulated the mass accretion and outflow rates.

An oscillation of an intrinsic accretion flow property can modulate the X-ray flux in several ways, since the corona radiates via Compton scattering. Karpouzas et al. (2021) explored a model initially considered for kHz QPOs in NSs (Kumar & Misra 2014; Karpouzas et al. 2020), which perturbs the Kompaneets equation by introducing a sinusoidal oscillation into the external heating rate (i.e. the mass accretion rate, which is the mechanism that provides the external heating). This leads to the electron temperature to oscillate, and feedback onto the source of seed

photons additionally leads to a modulation of the seed photon temperature with a time lag resulting from the time it takes photons to undergo many scatterings in the corona. The oscillations in these two temperatures in turn introduce oscillations into the power-law index. The predicted spectral variability can then be tested against spectral-timing properties of the QPOs observed in XRBs. Such a process occurring in a fixed geometry predicts that there should be no systematic dependence of QPO properties on the angle from which we view the source.

1.9 Observational tests of QPO models

As is usual when there are multiple theories to explain a phenomenon, we need to examine the various observational properties of low frequency QPOs to begin to understand the mechanism behind them.

The models for type-C QPOs, both those described above and others in the literature, can be broadly classified into *geometrical* models, where the observed change in the flux is attributed to part of the system physically varying in size or shape quasi-periodically; or *intrinsic* models, where a property such as the pressure or accretion rate is modulated while the geometry remains otherwise stable. I will outline a number of tests used to study QPOs in the literature, which will build into work I have done on QPOs which will be presented in chapters 2 and 4.

1.9.1 Inclination Dependence

Any geometrical origin of a QPO is likely to result in an inclination dependence of the QPO properties, whereas an intrinsic origin of the QPO would not. For example, a precessing torus would appear more variable when viewed at high inclinations (i.e. edge-on) as it tilts towards or away from us; a precessing jet however would be more variable if viewed at low inclinations (face-on) as the Doppler boosting is much stronger for smaller inclination angles⁵.

⁵Except, of course, for the case of a precessing jet viewed directly down the precession axis, for which no variability at all would be seen.




Figure 1.13: Fig. 3 of van den Eijnden et al. (2017), which shows the measured phase lag ϕ for type-C QPOs of different frequencies seen in different sources. The sources were broadly classified into high inclination and low inclination, and the results were binned logarithmically in frequency space to produce the blue and red points respectively to encapsulate the differing high and low inclination trends.

Motta et al. (2015) used a sample of *RXTE* observations containing QPOs to show that the type-C QPOs are indeed stronger in higher inclination sources, however the trend is reversed for type-B QPOs, which are stronger in lower inclination sources. This is strong evidence that low frequency QPOs have a geometric origin, and also highlights that type-B and type-C QPOs likely have different mechanisms.

van den Eijnden et al. (2017) looked at the inclination dependence of the QPO phase-lags, that is whether the QPO within a soft energy band either lags or leads the QPO measured in a hard band. They considered the phase lags between hard ($\sim 7 - 13$ keV) and soft ($\sim 2 - 7$ keV) photons for different frequency QPOs, with sources split into ‘low’ and ‘high’ inclination. The main result (Fig. 3 of their paper) is reproduced in Fig. 1.13 here. It illustrates that for $\nu_{\text{QPO}} \lesssim 2$ Hz all the sources they considered showed a small hard lag (where modulations in the hard X-ray flux lag behind those in soft X-ray flux). However, for type-C QPOs with

$\nu_{\text{QPO}} \gtrsim 2$ Hz, low inclination sources showed much larger hard lags, whereas high inclination sources showed large soft lags (where modulations in the soft X-ray flux lag behind those in hard X-ray flux). van den Eijnden et al. (2017) also looked at type-B QPOs which too showed soft lags for higher inclination, however there were only a small number of sources so the results weren't statistically significant.

1.9.2 Frequency-resolved spectroscopy

Frequency-resolved spectroscopy involves finding the energy spectrum of the power in different Fourier frequencies. In particular, it can be used to see the strength of the QPO harmonics in given energy bands. This involves finding light curves in narrow energy bands, computing the power spectrum of each of these light curves, and then extracting the power at the frequencies of interest. A detailed discussion of an application of this method can be found in section 2.3.3 of this thesis.

Sobolewska & Życki (2006) used frequency-resolved spectroscopy to study the fundamental harmonic of type-C QPOs, with frequencies ranging between 1 – 10 Hz, in a number of BH XRBs as observed with *RXTE*. By using the absolute RMS normalisation for the power of the QPO (see section 1.10.1), it is possible to model the RMS spectra of the QPO with the same models as the phase-averaged⁶ energy spectrum. They found, when using a model including both the quasi-thermal disc emission and Comptonised emission, that the phase-averaged spectrum required both the disc and Comptonised components, however the RMS spectrum of the QPO fundamental did not include the disc emission. They also found some evidence of reflection features in the QPO spectrum. This key result highlights that the QPO fundamental is tied to the coronal emission, not the intrinsic disc emission. However, they also noticed that the QPO spectrum doesn't always match the exact Comptonised spectrum, with the QPO tending to have a harder RMS spectrum when

⁶In the literature, this is sometimes also referred to as the 'time-averaged' spectrum. However, as work of this type ends up averaging over the course of an observation (or at least over a period of an observation much longer than the time period associated with the frequencies concerned), I prefer to use the term 'phase-averaged spectrum' as it represents the spectrum averaged over the entire phase of all the oscillation frequencies.

the phase-averaged spectrum is harder, and conversely the QPO RMS spectrum is softer when the phase-averaged spectrum is softer.

Axelsson et al. (2014) studied the RMS spectrum of a type-C QPO, but went further by studying the RMS spectrum of not only fundamental but also second harmonic of the QPO. While they found similar results as Sobolewska & Życki (2006) for the fundamental harmonic, they found spectrum of the second harmonic was much softer than the phase-averaged spectrum. They interpreted this result within the context of the model of a truncated disc with a hot inner flow model undergoing Lense-Thirring: if the hot inner-flow is hotter or viewed with a smaller optical depth for smaller radii, it would produce a softer power law closer to the disc truncation radius, but a harder power law closer to the BH. The viewing angle dependence of the emission from a Comptonising slab has harmonics which are weighted depending on optical depth; higher harmonics are stronger from regions with lower optical depths, which also give a softer power law spectra. Similar results were also found in another source (Axelsson & Done 2016).

These observations of the QPO modulating the corona rather than the disc disfavour theorised QPO mechanisms in which the oscillation comes from within the disc (such as c-modes). Such models are not, however ruled out. For example an oscillation may originate in the inner part of the disc that is blocked from direct view by the corona, meaning that we only see the constant outer disc component in the phase-averaged spectrum whereas the variable seed photon luminosity seen by the corona gives rise to a modulation of the Comptonised flux (Ingram & Motta 2019).

1.9.3 QPO Phase-resolved spectroscopy

Frequency-resolved spectroscopy considers how the amplitude of the QPO harmonics depend on photon energy. As it has been repeatedly shown that the QPO RMS spectra is not flat, and that phase-lags between hard and soft bands are non-zero, it becomes obvious that the energy spectrum must be varying with QPO phase. Phase-resolved spectroscopy is a technique to unpick how the energy spectrum varies with QPO phase. This can also be considered as trying to find the waveform

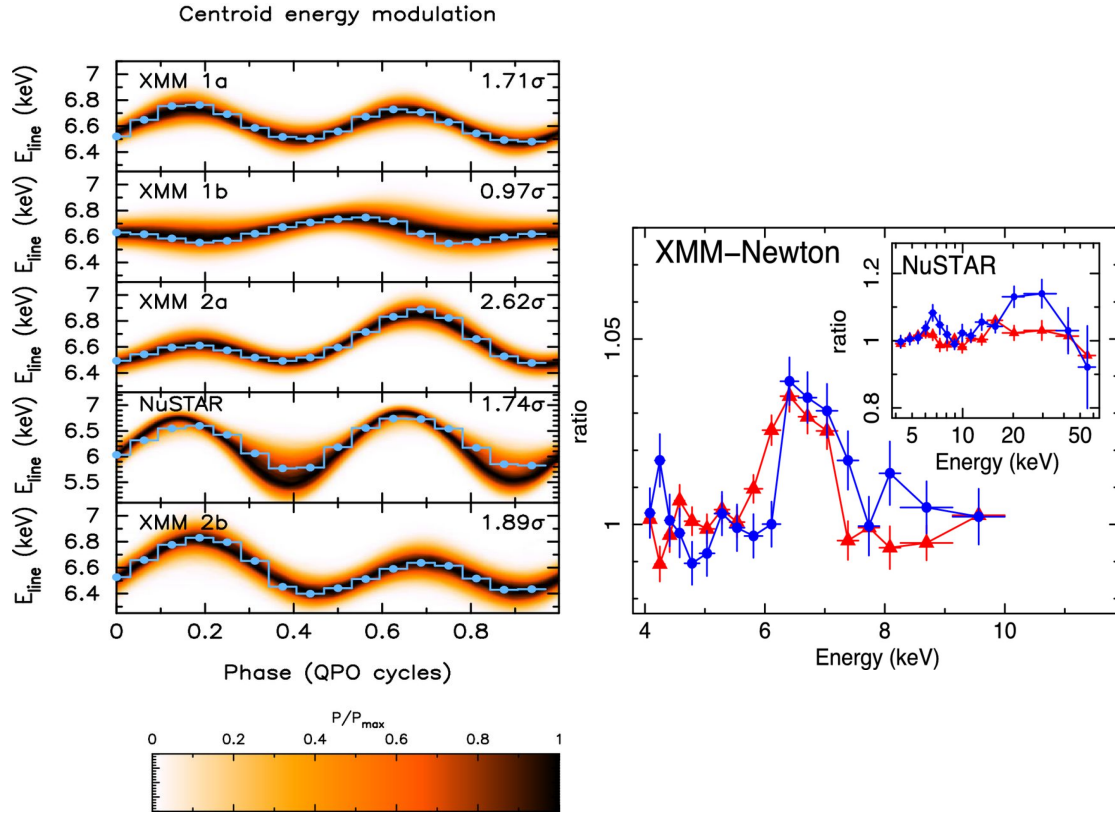


Figure 1.14: Figs. 7 and 3 from Ingram et al. (2016), which show their key result of detecting the modulation of the iron line centroid energy with QPO phase. The left subfigure shows how the centroid energy of a Gaussian iron line varies over the course of a QPO cycle, with each panel considering a separate observation. The blue line shows the centroid energy of the line which was included in spectral fits of spectra which have been reconstructed at 16 different phases of the QPO cycle. An illustrative MCMC was also run, which is shown as the background heat map as a visual guide to the significance of the modulation. The right subfigure shows two of these reconstructed spectra (at 7/16 cycles marked with red triangles, and 11/16 cycles marked with blue circles), for *XMM-Newton*, and *NuSTAR* in the inset. This demonstrates the differences at these two phases in both the iron line ($\sim 6 - 7$ keV), and also the Compton hump ($\sim 20 - 50$ keV).

of the QPO within different energy bands. In this thesis (in chapters 2 and 4) will be using the QPO phase-resolving method which was originally introduced by Ingram & van der Klis (2015).

Ingram et al. (2016) used the waveform reconstruction method of performing phase-resolved spectroscopy to analyse a type-C QPO seen in the BH XRB H1743-332. They fit first with a very simple model of an absorbed power law plus a Gaussian to fit the iron line. They allowed the power-law index, the centroid energy and width of the iron line, and the normalisations of the two components to be

independently modulated with QPO phase, assuming a two harmonic sinusoidal oscillation for each. Fig. 1.14 (left, Fig. 7 from their paper) shows how centroid energy of the Gaussian is modulated with QPO phase in each of the five observations they considered, found by fitting spectra corresponding to 16 QPO phases, along with the statistical significance of this modulation. When all observations were combined, their work showed that the iron line is modulated with QPO phase with 3.7σ significance overall. This is strong evidence of a geometric QPO mechanism, as a only a geometrical mechanism could change the shifting and broadening of the iron line that we see, and was an effect previously predicted for the Lense-Thirring precession model (Ingram & Done 2012). That said, this could be interpreted as a intrinsic QPO model with a modulated ionisation state, however the observed shift in energy would require a large change in illuminating flux which is not observed in the iron line normalisation. The large change is also disfavoured by the Compton hump which is strongest when the iron-line is blue-shifted as can be seen in Fig. 1.14 (right, Fig. 3 from their paper), which is opposite to the effect of varying the ionization parameter.

Ingram et al. (2017) went further, and fit the data with a full tomographic reflection model. They considered a model which could preferentially illuminate different patches of the disc at different QPO phase γ , such that the specific intensity (disc at disc radius r , azimuthal angle ϕ) of the reflected X-rays is

$$I(E; r, \phi, \gamma) \propto r^{-q} \left\{ 1 + A_1 \cos \left(\frac{1}{2} [\gamma - \phi + \phi_1] \right) + A_2 \cos (\gamma - \phi + \phi_2) \right\} I(E), \quad (1.20)$$

where $I(E)$ is the rest-frame emitted spectrum for photon energy E . In this model, the radial dependence is simply a power law, whereas there is an azimuthal dependence based upon the amplitude parameters $A_{\{1,2\}}$ and azimuth phase parameters $\phi_{\{1,2\}}$. When $A_1 = A_2 = 0$, there is simply axisymmetric, phase-invariant reflection. However, when one or both of $A_{\{1,2\}}$ is non-zero, then this represents one or two bright “patches” which move around the disc with QPO phase. Through ray-tracing this specific intensity from over the disc they were able to calculate an observed reflection spectrum at any QPO phase, taking into account all relativistic effects

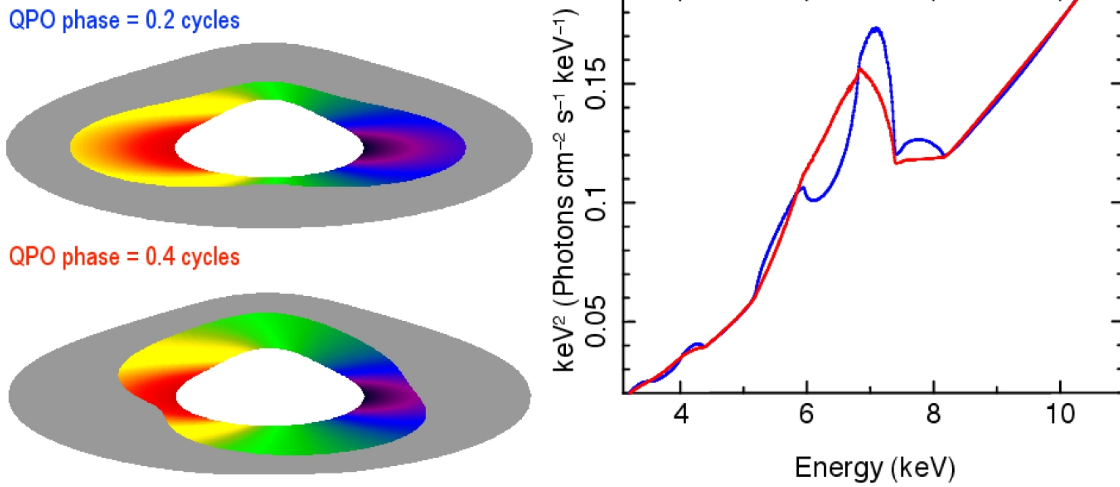


Figure 1.15: Fig. 2 from Ingram et al. (2017), showing their best fitting disc emissivity profile at two QPO phases, along with the iron line profile within the spectrum model at these two phases (0.2 cycles in blue, 0.4 cycles in red).

that broaden and smear the observed spectrum. This model was designed with the precessing inner-flow model in mind (with the bright spots caused by illumination from above/below the oblate flow), although it is also able to represent other geometries, such as a precessing jet, or indeed axisymmetric models which have $A_1 = A_2 = 0$. It is this technique, attempting to rule out the symmetrical case of $A_1 = A_2 = 0$, that I will employ in chapters 2 and 4 of this thesis.

Ingram et al. (2017) found that that an axi-asymmetric profile ($A_1, A_2 \neq 0$) was required at a 2.4σ confidence level. Fig. 1.15 (Fig. 2 from their paper) shows the best fitting axi-asymmetric illumination profile at two different QPO phases (0.2 and 0.4 cycles), along with the model spectrum at these phases. The change in structure and centroid energy of the line is very dramatic, particularly the prominent blue-wing at 0.2 cycles. Along with the illumination profile, they also found that the *reflection fraction* was modulated with QPO phase at 3.5σ significance (here, they defined this as the ratio of number of photons intercepted by the disc to the number of photons radiated by the corona, equivalent to the covering fraction of the corona⁷). This was a key result which proves that the QPO is caused by a geometric effect with high significance. They also examined whether the model was able to explain

⁷It should be noted that I used a slightly different definition of the reflection fraction in the work covered by this thesis: the *observer's reflection fraction*, as discussed in section 2.4.2.1.

the data if forced to be axisymmetric. They instead tried to use QPO phase modulations in the ionisation parameter, disc inner radius, and the radial emissivity index q to explain the spectral modulations. They found that all of these attempts produced worse fits than the axi-asymmetric model fit with $> 3.5\sigma$ significance.

Stevens & Uttley (2016) used a slightly different cross-correlation function method to phase resolve a type-B QPO in the BH XRB GX 339-4 observed by *RXTE*. Their best-fitting model included a typical multi-temperature black body which was constant with QPO phase, a single temperature black body component with a QPO-phase variable temperature that represents the inner-disc, a variable power-law, and a Gaussian iron line profile. While the signal-to-noise didn't allow them to constrain variability of the iron line, they did find significant variability of both the inner black body component and the power law. The variability of the black body flux came directly from variation in its temperature, which could result from accretion rate fluctuations, which the authors rule out due to the relatively small variability of the black body in comparison to the power law. They also considered if the black body variability could result from variable irradiation (from the variable power law), but they ruled this out due to the long 0.3 cycle phase lag between the black body variations and the power law variation. They also suggested that this result could instead be caused by changing observed Doppler shifts, again through preferential illumination, but concluded that this only makes sense for emission from a large scale height, such as from a jet (such emission has been associated with the SIMS and type-B QPOs; Fender et al. 2004), due to the small temperature variations.

Together, this work showed that phase-resolved spectroscopy is a very powerful technique for understanding QPOs, which is becoming even more useful as we move to the next generation of X-ray telescopes with large collecting areas, high timing precision, and fine energy resolution.

1.10 Fourier methods

The Fourier transform is a method which decomposes a signal into a series of sine waves. For a continuous series $x(t)$, the Fourier transform and its inverse are defined as

$$\begin{aligned} X(\nu) &= \mathcal{F}[x(t)](\nu) = \int_{-\infty}^{\infty} x(t)e^{2\pi i\nu t} dt \\ x(t) &= \mathcal{F}^{-1}[X(\nu)](t) = \int_{-\infty}^{\infty} X(\nu)e^{-2\pi i\nu t} d\nu, \end{aligned} \quad (1.21)$$

where ν is the Fourier frequency. This holds for an ‘infinite’ series $-\infty < t < \infty$, which covers an infinite range of Fourier frequencies $-\infty < \nu < \infty$, and where $x(t)$ and $X(\nu)$ can be in general complex. The Fourier transform can be split up into a frequency dependent amplitude and phase $X(\nu) = a(\nu)e^{i\phi(\nu)}$.

Alas, the Fourier transform as presented in Eq. 1.21 is not helpful when analysing X-ray data. To begin, observations are unfortunately not infinite, and have both a start and an end. Next, the arrival of X-ray photons at a detector is a Poisson process. We therefore bin photons based upon their arrival time into discrete time bins each with a duration δt and labeled k , such that $t_k = k\delta t$ and $0 \leq k < N$, where the total time of the observation is $T = N\delta t$. This then gives us a discrete light curve $x_k(t_k)$, measured as the count rate, however corresponding to what can only be an integer number of photons in each bin. We can then use the discrete Fourier transform

$$\begin{aligned} X_n(\nu_n) &= \sum_k x_k(t_k)e^{2\pi i \frac{k n}{N}} \\ x_k(t_k) &= \frac{1}{N} \sum_n X_n(\nu_n)e^{-2\pi i \frac{k n}{N}}, \end{aligned} \quad (1.22)$$

which runs for $1 - \frac{N}{2} \leq n \leq \frac{N}{2}$, and has frequency bins $\nu_n = \frac{n}{N\delta t}$ (where ν_n is the observed Fourier frequency). Here the $\nu_0 = 0$ bin is just the mean count rate μ_s , and $\nu_{N/2} = \frac{1}{2\delta t}$ is the Nyquist frequency, the highest frequency that can be probed.

While the negative frequency bins ($n < 0$) are difficult to conceptualise (they result from the Fourier transform’s relation to the complex plane), when $x(t_k)$ is a mathematically real signal (as is the case for a light curve) then the Fourier transform has the property $X(-\nu) = X^*(\nu)$, and these negative frequencies don’t contain any additional information.

1.10.1 Power spectrum

The power spectrum is a representation of the variability in any frequency, defined as

$$P(\nu) = |X(\nu)|^2 = a^2(\nu). \quad (1.23)$$

For the discrete case, we use the periodogram

$$\mathcal{P}_n(\nu_n) = \frac{2T}{\mu_s^2} |X_n(\nu_n)|^2 = \frac{2T}{\mu_s^2} |a_n|^2. \quad (1.24)$$

The normalisation chosen here is the *fractional RMS normalisation* (Lewin et al. 1988; Belloni & Hasinger 1990), such that the integral of the power spectrum gives the total squared fractional RMS. In the discrete case, such an integral becomes the summation of the product of periodogram bins and the frequency bin width, which is related to the variance of the light curve through

$$\sum_{n=1}^{N/2} \mathcal{P}_n \delta\nu = \frac{1}{T} \sum_{n=1}^{N/2} \mathcal{P}_n = \frac{\text{var}[x(t)]}{\mu_s^2}. \quad (1.25)$$

The factor of 2 in the normalisation allows this statement to be true when only considering the power calculated in the positive frequencies (as otherwise half the power would lie in the negative frequencies).

As with any mathematical tool, it is important to understand the statistics of the power spectrum before using it, as discussed in van der Klis (1989). Consider breaking X_n up into its constituent real and imaginary parts, $X_n = X_{R,n} + iX_{I,n}$. Each of $X_{R,n}$ and $X_{I,n}$ are a sum of many random variables, and it therefore follows (through the central limit theorem) that they each follow a Gaussian distribution. The power spectrum $X_n X_n^* = (X_{R,n} + iX_{I,n})(X_{R,n} - iX_{I,n}) = X_{R,n}^2 + X_{I,n}^2$. By definition, \mathcal{P}_n then follows a scaled χ^2 distribution with two degrees of freedom. This however does bring a problem, as the standard deviation of such a distribution is the same as the mean! This can be fixed through averaging, and provided enough realisations are averaged together then the central limit theorem is again valid, meaning $\langle \mathcal{P}_n \rangle$ follows a Gaussian distribution. There are two ways in which to average the power spectrum. We can either use *ensemble averaging*: slice the light curve into multiple (let's say M) segments, calculate a power spectrum for each,

and then find the average for each frequency. The second way involves simply averaging over consecutive (let's say W) frequency bins. Both methods reduce the frequency resolution, and can be combined if needed to produce an average over MW realisations. Provided $MW \gtrsim 25$, then the bin has been averaged over enough realisations such that the central limit theorem holds.

As mentioned, the light curve $x_k(t_k)$ is not only discrete, but each bin contains an integer number of photons following a Poisson distribution. The light curve x_k can then be considered a superposition of the source signal s_k and of a noise term n_k , with the simple relation $x_k = s_k + n_k$, as this photon counting noise is independent of the source signal. The linearity of the Fourier transform means we trivially also have $X_n = S_n + N_n$. When considering the power spectrum of this, we find

$$|X_n|^2 = |S_n + N_n|^2 = |S_n|^2 + |N_n|^2 + S_n N_n^* + S_n^* N_n. \quad (1.26)$$

As the noise and source terms are independent, the cross terms cancel under averaging, and hence $\langle \mathcal{P}_{\text{total}} \rangle = P_{\text{signal}} + P_{\text{noise}}$, where P_{signal} and P_{noise} are the underlying powers within the source signal and of the noise on top. Pure Poisson noise has a white noise spectrum with constant power $P_{\text{noise}} = \frac{2}{\mu_s}$ for fractional RMS normalisation (van der Klis 1989). If there are photons within the light curve which are instead attributed to background⁸, the total count rate is simply the sum of the source and background $\mu_t = \mu_s + \mu_b$. As the Poisson noise is dependent on μ_t , but the normalisation depends on μ_s , the Poisson noise becomes $P_{\text{noise}} = \frac{2\mu_t}{\mu_s^2}$ (van der Klis 1989). However, every real X-ray detector experiences *dead time*, a short period of time after a detection event during which the detector cannot detect another photon if it arrives⁹. This not only suppresses the count rate of the detector, but also changes both the expected value and also frequency response of the Poisson noise in the detected light curve (van der Klis 1989; Bult 2017).

⁸This only applies for photons which cannot be individually distinguished within the light curve as either source or background, and the background can only be seen in aggregate. Any events that can be classified as non-source can be removed during data reduction and do not effect the simple Poisson noise.

⁹This dead time may be a constant, or may depend on the exact nature of the event or other instrument effects

1.10.2 Cross spectrum

The cross spectrum is similar to the power spectrum, however it is computed for two light curves (van der Klis et al. 1987). It contains information on correlated variability between the two. If, as before, we consider two light curves $x(t)$ and $y(t)$, with Fourier transforms $X(\nu) = a_X(\nu)e^{i\phi_X(\nu)}$ and $Y(\nu) = a_Y(\nu)e^{i\phi_Y(\nu)}$. The cross spectrum between them is then, much like the power spectrum in Eq. 1.23, defined as

$$G(\nu) = X(\nu)Y^*(\nu) = a_X(\nu)a_Y(\nu)e^{i(\phi_X(\nu)-\phi_Y(\nu))}. \quad (1.27)$$

Each frequency then has an amplitude, encapsulating the power of the correlated variability at that frequency, and the phase lag of that frequency of y compared to x such that $\arg[G(\nu)] = \phi_X(\nu) - \phi_Y(\nu) > 0$ corresponds to the signal in x lagging that in y .

When considering the cross spectrum between two signals, we can consider the case where they contain correlated signals A_n and B_n , and uncorrelated components T_n and U_n . The Fourier-domain signals are hence $X_n = A_n + T_n$, and $Y_n = \gamma_n B_n + U_n$, where γ_n is the coherence between the signals. The cross spectrum of these signals is

$$G_n = X_n Y_n^* = \gamma_n A_n B_n^* + A_n U_n^* + \gamma_n T_n B_n^* + T_n U_n^*. \quad (1.28)$$

As the last three terms are each effectively the cross spectrum of uncorrelated signals, upon averaging they vanish, such that $\langle G_n \rangle = \gamma_n \langle A_n B_n^* \rangle$. Provided the averaging is sufficient, both the real and imaginary components of $\langle G_n \rangle$ follow Gaussian distributions.

1.10.2.1 Co-spectrum

NuSTAR has a long dead time, the time period after detecting a photon during which it cannot detect another. The effect of dead time is further complicated by different events on the detector causing different dead time properties. This induces a complicated photon counting noise within the light curves and their associated power spectra, which (especially at high frequencies) cannot be easily modelled, except for fitting an empirical function (Bult 2017). However, *NuSTAR*

has two separate Focal Plane Modules (A and B; FMFA, FMPB) which observe simultaneously yet independently. They each therefore see the same source signal but different noise terms. Bachetti et al. (2015) suggested that the issue of the long dead time issue could be solved by finding the *co-spectrum* between the light curves of FMFA and FMPB.

The co-spectrum is the real part of the cross spectrum. We can demonstrate how using the co-spectrum between the FMFA and FMPB eliminates photon counting noise by representing the Fourier transforms of the FMFA and FMPB light curves as $A_n = S_n + N_{A,n}$ and $B_n = S_n + N_{B,n}$ respectively. After averaging, $\Re[\langle G_n \rangle] = \langle S_n S_n^* \rangle$, which is simply an estimate of the power spectrum in the absence of photon counting noise. The imaginary component of the cross spectrum is only noise from the random realisations.

The full co-spectrum is simply the real part of the cross spectrum, i.e. the real part of Eq. 1.28. While under sufficient averaging, the cross terms tend to 0, this is not the case for a single realisation and these cross terms make understanding the statistical distribution of a single realisation of the co-spectrum practically impossible and as such it remains unknown. Huppenkothen & Bachetti (2018) showed that the distribution of the unaveraged co-spectrum between two completely uncorrelated signals (consider the real part of Eq. 1.28 with $\gamma_n = 0$) is the Laplace distribution. While useful for statistically detecting a correlated signal above noise, this is not helpful for the work in this thesis. I considered the distribution of a single realisation of co-spectrum in a simplified case which pertains to *NuSTAR*. In the fictitious case where the FMFA light curve contains only pure signal, such that its Fourier transform is $A_n = S_n$, however FMPB has separately detected the signal but with an additional noise term such that its Fourier transform can be written $B_n = S_n + N_n$. The co-spectrum between FMFA and FMPB is then $\Re[A_n B_n^*] = \Re[S_n (S_n + N_n)^*]$. For simplicity, it is possible to factor out the power of the signals, such that $S_n = \sqrt{P_{S,n}} \tilde{S}_n$ and $N_n = \sqrt{P_{N,n}} \tilde{N}_n$; the new variables have unity power such that $\langle \tilde{S}_n \tilde{S}_n^* \rangle = \langle \tilde{R}_n \tilde{R}_n^* \rangle = 1$. With this notation, the Fourier transforms of the FMFA and FMPB light curves become $A_n = \sqrt{P_{S,n}} \tilde{S}_n$

and $B_n = \sqrt{P_{S,n}} (\tilde{S}_n + \sqrt{\eta_n} \tilde{R}_n)$ respectively, where $\eta_n = \frac{P_{N,n}}{P_{S,n}}$ is the ratio of the source and noise powers. A realisation of the co-spectrum can then be written

$$\Re[A_n B_n^*] = \sqrt{P_{S,n}} \tilde{S}_n \sqrt{P_{S,n}} (\tilde{S}_n + \sqrt{\eta_n} \tilde{R}_n)^* = P_{S,n} Z, \quad (1.29)$$

where $Z = \tilde{S}_n (\tilde{S}_n + \sqrt{\eta_n} \tilde{R}_n)^*$ is a random variable following a distribution with the probability density function

$$P(Z = z) = \frac{e^{[2\eta_n^{-1}(z - |z|\sqrt{1+\eta_n})]}}{\sqrt{1 + \eta_n}}. \quad (1.30)$$

The consideration of the FMPA light curve being noiseless is obviously wrong. If the noise term N_n is instead considered as the overall difference between FMPA and FMPB then the functional form of Eq. 1.30 is potentially reasonable. Indeed, through experimentation this distribution seems to work for the more general $\Re[A_n B_n^*] \Re[(S_n + M_n)(S_n + N_n)^*]$, however the relation of η_n to the noise powers is still unclear.

1.10.3 Bi-spectrum

While the power spectrum is a second order analysis method, the *bi-spectrum* is a third order analysis method, which describes coupling between two frequencies. For the signal $X(\nu)$ it is defined for a pair of frequencies ν_1 and ν_2

$$\mathcal{B}(\nu_1, \nu_2) = X(\nu_1)X(\nu_2)X^*(\nu_1 + \nu_2). \quad (1.31)$$

The phase and modulus of the averaged bi-spectrum, $\langle \mathcal{B}(\nu_1, \nu_2) \rangle$, are respectively known as the *bi-phase* (Maccarone 2013) and *bi-coherence* (Kim & Powers 1979). The bi-coherence is a measure of how well the phases of Fourier components at the two frequencies correlate with one another, and the bi-phase informs on the phase difference between those two components. The bi-spectrum has therefore been used to explore how QPO harmonics couple to each other and to broadband noise frequencies (see e.g.; Arur & Maccarone 2019; Arur & Maccarone 2022). In this thesis I briefly explore using the bi-spectrum to find the phase between two QPO harmonics for the purposes of phase-resolved spectroscopy. However, beyond this

the bi-phase can be used to describe the ‘peakiness’ of a light curve, as discussed by Maccarone (2013), where a bi-phase of 0 produces a peaky ‘pulsar-like’ light curve, but a bi-phase of π produces an ‘eclipser-like’ light curve (see his Fig. 1). One of the most well-known examples of a non-zero bi-coherence is the linear RMS-flux relation.

1.10.3.1 Linear-RMS Flux relation

Uttley & McHardy (2001) showed that the RMS observed in the broadband noise in X-ray light curves from accreting objects is intrinsically non-stationary, as the total RMS within the light curve varies linearly with the overall flux, the linear RMS-flux relation. They showed that this effect occurs across all time scales, and therefore the fluctuations with any Fourier frequency will change the RMS at other frequencies. This leads to the non-zero bi-coherence in the light curve. Heil et al. (2011) showed this effect also occurs for type-C QPOs, with a correlation between the QPO RMS amplitude and mean source flux. Heil et al. (2012) studied the relation for a number of BH XRBs across a number of states, finding it to be an ubiquitous process suggesting that it is tied to core accretion processes. The effect is evidence for the *propagating fluctuations model* (Lyubarskii 1997; Arévalo & Uttley 2006; Ingram & van der Klis 2013), by which variability propagates down through the disc, carrying any variability with it; the largest radii of the disc impart lower frequency variability, on top of which higher frequency can be imprinted by the inner regions of the disc.

Time is an illusion. Lunchtime, doubly so.

— Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

2

Phase-resolved spectroscopy of a quasi-periodic oscillation in the black hole X-ray binary GRS 1915+105 with *NICER* and *NuSTAR*

Contents

2.1	Introduction	48
2.2	Observations	49
2.2.1	<i>NuSTAR</i> Data Reduction	49
2.2.2	<i>NICER</i> Data Reduction	50
2.2.3	Energy spectrum	51
2.2.4	Power spectrum	52
2.2.5	Spectral Timing State	53
2.3	Phase Resolved Spectroscopy	53
2.3.1	QPO frequency tracking	56
2.3.2	Phase lag spectrum	58
2.3.3	Fractional RMS Spectrum	60
2.3.4	Phase difference between harmonics	61
2.3.5	Reconstructed Fourier Transformed Spectra	67
2.4	Theoretical Model	68
2.4.1	Corona	68
2.4.2	Disc	69
2.4.3	The complete model	74
2.5	Model Fits	74
2.5.1	Fitting procedure	74
2.5.2	Results	75
2.6	Discussion	82

2.6.1	Comparison with H1743-322	82
2.6.2	Asymmetric illumination profile	84
2.6.3	Inner Radius	86
2.6.4	Misalignment	88
2.6.5	Biases in the phase-averaged spectrum	90
2.6.6	Thermalisation lags	91
2.7	Conclusions	93

2.1 Introduction

Type-C low frequency quasi-periodic oscillations (QPOs) are often observed in the X-ray light curves from black hole (BH) X-ray binaries (XRBs) (Ingram & Motta 2019). By investigating the behaviour of the QPO, it is possible to learn about the accretion processes which occur close to the BH. One powerful technique to study QPOs is phase resolved spectroscopy (Ingram et al. 2016, 2017), which aims to uncover how the energy spectrum changes over the course of a QPO cycle. By resolving the spectrum at different QPO phases, it is possible to see how the QPO is present in the quasi-thermal disc spectrum; the power law corona spectrum; and the reflection spectrum of coronal photons which have been reprocessed within the disc.

The theorised mechanisms causing QPOs can broadly be classified into *intrinsic* models, in which the intrinsic luminosity of the accretion flow oscillates; or *geometric* in which the observed QPO is caused by a variation of the geometrical structure of the accretion flow. It has been known for a while that the coronal power law spectrum varies with greater RMS than the disc component (Sobolewska & Życki 2006), and recent work by Motta et al. (2015), Heil et al. (2015), and van den Eijnden et al. (2017) showed there was strong evidence of an inclination dependence on the observed properties of the QPOs. This evidence supports a geometrical origin of the QPO. Ingram et al. (2016, 2017) pioneered using phase resolved spectroscopy to reconstruct the waveform of the QPO in narrow energy channels, and used the results to search for a ‘rocking’ Fe K α line as signal of an asymmetric illumination profile shining on the disc and varying with QPO phase such as predicted by

precession-based QPO models. In doing so, they found moderately significant results in favour of this asymmetric profile.

In this chapter, which is adapted from Nathan et al. (2022), I examine observations of the BH XRB GRS 1915+105 as detailed in Section 2.2, with the goal of isolating the asymmetric profile to high statistical significance. I present improvements to the methodology of performing the phase resolved spectroscopy in Section 2.3; and also detail the model which I have used, which is an updated version of the Ingram et al. (2017) model, in Section 2.4. The final results are presented and discussed in Sections 2.5 and 2.6 respectively.

2.2 Observations

The Neutron star Interior Composition ExploreR (NICER; Gendreau et al. 2016) and *the Nuclear Spectroscopic Telescope ARray* (NuSTAR; Harrison et al. 2013) observed GRS 1915+105 quasi-simultaneously on 8th – 9th June 2018¹. The details of the observations are summarised in Table 2.1. In this section, we detail our data reduction procedure and present the basic spectral and timing properties of the data. When analysing the data we make use of HEASARC (2014), and custom code written in Python3 (Van Rossum & Drake 2009).

2.2.1 NuSTAR Data Reduction

We used the *NuSTAR* analysis software, *NuSTARDAS* v1.8.0 with *HEASOFT* v6.22. We generated a cleaned event list with associated list of good time intervals (GTIs) for both focal plane modules (FPMs) – FPMA and FPMB – using *nupipeline*. From this we used *nuproducts* to extract source and background spectra from 49.2 arcsec circular regions and generate spectral response files. We find that the source contributes > 99.9 per cent of the total counts measured by *NuSTAR*. We did not perform a background subtraction when extracting light curves we use for the timing analysis, as the background is not expected to be variable on the QPO period. We use our own custom code to extract FPMA and FPMB source

¹MJD 58277-58278




Figure 2.1: Ratio of the *NICER* 0.24 – 10 keV, and the *NuSTAR* FPMA and FMPB 3 – 75 keV spectra in black, red, and blue respectively, to an absorbed power law model with photon index 2.01, and absorption with hydrogen column density $6.2 \times 10^{22} \text{ cm}^{-2}$.

region light curves from the cleaned event list in 11 broad energy channels in the energy range 3 – 78 keV. We used the `FTOOL rbnrmf` to re-bin the spectral response files into these 11 energy bands.

2.2.2 *NICER* Data Reduction

We used the *NICER* analysis software, *NICERDAS* v2018-04-13_V004 with *HEASOFT* v6.24. We extracted the MPU-merged, uncleaned event lists with the *FTOOL NICER12* for each of the two *NICER* observations, which we then merged together using `nimpumerge` and cleaned with *NICERclean*. We used filter and GTI files which were combined from those of the separate obs IDS using *FTOOLS ftmerge* and `nimaketime`. This results in a single cleaned event list for the two obs IDs combined.

We extracted a flux-energy spectrum from the resulting merged event list using XSELECT, and estimated the instrumental background with the *NICER* `gof.bkg` version 0.5 python script (Remillard et al. 2021). We find that the source contributes 98.6 per cent of the total counts measured by *NICER*. Again, we did not perform a background subtraction when extracting light curves we use for the timing analysis, as the background is not expected to be variable on the QPO period. We used the spectral response files ‘nxtiref20170601v002.rmf’ and ‘nxtiaveonaxis20170601v004.arf’ from CALDB. We extracted light curves from the merged event list in 40 broad energy channels in the energy range 0.3 – 10 keV using our own custom code. We used the FTOOL `rbnrnf` to re-bin the spectral response files into these 40 energy bands.

2.2.3 Energy spectrum

Fig. 2.1 shows the *NICER* (black) and *NuSTAR* (red: FPMA; blue: FPMB) background subtracted flux-energy spectrum plotted as a ratio to a folded absorbed power-law model. We set the hydrogen column density to $N_H = 6.2 \times 10^{22} \text{ cm}^{-2}$ (the absorption model is TBABS with the abundances of Wilms et al. 2000) and the power-law index to $\Gamma = 2.01$ for all three spectra, but allow the three spectra to each have their own normalisation. We see strong reflection features including an iron line at ~ 6.4 keV and a broad Compton hump peaking at ~ 30 keV. We also see that the cross-calibration between *NICER* and *NuSTAR* is excellent in the $\sim 3 - 10$ keV energy range in which their band passes overlap. At energies below ~ 2.7 keV, the *NICER* spectrum includes features that are likely due to calibration uncertainties, and the ‘shelf’ of the response from higher energy photons (NASA 2021) which is dominant below ~ 1 keV due to the astrophysical absorption leaving very few source photons at low energies. We therefore only consider energy channels > 2.7 keV in our spectral analysis. We note there is a cross-calibration discrepancy between *NuSTAR* FMPA & FMPB in the energy range $\sim 3 - 3.5$ keV due to a tear in the Multi Layer Insulation around *NuSTAR*’s FMPA (Madsen et al. 2020), so we also ignore the FMA energy channels < 3.48 keV.




Figure 2.2: The 3-10 keV power spectra of the *NICER* (blue) and *NuSTAR* (orange) observations. The *NICER* spectrum is Poisson-noise subtracted, while the *NuSTAR* spectrum is estimated from the co-spectrum between FPMA and FMPB. The power spectra have been ensemble-averaged and geometrically re-binned to reduce the number of bins by a factor of 25.

2.2.4 Power spectrum

Fig. 2.2 shows the 3-10 keV power spectrum calculated for the merged *NICER* observation (blue) and the *NuSTAR* observation (orange). For both observatories, we extract light curves with time step δt from the cleaned event list. For the purposes of ensemble averaging (e.g. van der Klis 1989), we split the light curves into M segments labelled $1 \leq m \leq M$, each having N time bins and therefore a length $T = N\delta t$. Except for when otherwise stated we use $N = 8192$ and $\delta t = 1/128$ s throughout this paper, so our segments are $T = 64$ s long. The *NICER* power spectrum is calculated in the standard way (the magnitude-squared of the fast Fourier transform (FFT) of the lightcurve), with a constant Poisson noise level subtracted (van der Klis 1989; Uttley et al. 2014). For *NuSTAR* we instead calculate the co-spectrum between the FPMA and FPMB (Bachetti et al. 2015) in order to avoid instrumental features caused by the fairly large *NuSTAR* dead time of $t_D \approx 2.5$ ms. We also correct for the suppression of variability caused by the *NuSTAR* dead time using the simple formula $(\text{RMS}_{\text{det}}/\text{RMS}_{\text{intr}}) \approx 1/(1 + t_D r_{\text{intr}}) = (r_{\text{det}}/r_{\text{intr}})$, where RMS_{det} and RMS_{intr} are respectively the detected and intrinsic RMS variability amplitudes and r_{det} and r_{intr} are the detected and intrinsic count rates (Bachetti et al. 2015). For this observation, the ratio of detected to intrinsic variability is

$\text{RMS}_{\text{det}}/\text{RMS}_{\text{intr}} = 0.754$ (recorded in the *NuSTAR* spectral files as the keyword ‘DEADC’). We see that the *NuSTAR* co-spectrum is very similar to the *NICER* power spectrum. In both, we see a strong low-frequency QPO with a fundamental frequency of ~ 2.2 Hz and a 2nd harmonic at twice that frequency. We can classify this low-frequency QPO as ‘type-C’ based upon its frequency, $\text{RMS} \gtrsim 15\%$, and the presence of the ‘flat-top’ broadband noise (see e.g. Ingram & Motta 2019 for a summary of the characteristic features of type-A, B, and C low frequency QPOs).

2.2.5 Spectral Timing State

Belloni et al. (2000) identified 15 different classes from spectral and variability patterns that *GRS 1915+105* transitions between (also see Huppenkothen et al. 2017), most of which have only to date been observed in one other XRB (Altamirano et al. 2011). From the flux-energy spectrum and power spectrum, it is clear that *GRS 1915+105* was in the designated χ -class during this observation. In the χ -class, *GRS 1915+105* has similar spectral and timing properties as seen in other BH XRBs in the canonical hard state. The state was defined by observations with *RXTE*, however it is typically bright and has a hardness ratio corresponding to the presence of a hard power law. The χ -class can either be radio loud or quiet; during our X-ray observations the source was radio quiet (Motta et al. 2021). This happens to be the dimmest χ -class ever observed, preceding the transition of the source into its current heavily obscured state (Motta et al. 2021). The QPO frequency of ~ 2.2 Hz is also somewhat special, since it is at around this QPO frequency when the QPO phase lags transition from positive (hard photons lag soft photons) to negative (soft photons lag hard photons): the phase lag reduces approximately linearly with the log of QPO frequency, passing through zero at $\nu_{\text{qpo}} \sim 2$ Hz (Reig et al. 2000; Qu et al. 2010; van den Eijnden et al. 2017; Zhang et al. 2020).

2.3 Phase Resolved Spectroscopy

The aim of phase resolved spectroscopy is to investigate how the energy spectrum of the source varies with QPO phase, a task complicated by the ‘quasi-’ nature of

Table 2.1: Details of the simultaneous observations from *NICER* and *NuSTAR* on the 8th – 9th June 2018 (MJD 58277-58278). The *NICER* observation was split into two, however these are merged for our analysis. The net count-rate is reported for the background-subtracted spectra used for the flux-energy fits: 3.5 – 75, 3 – 75 keV for *NuSTAR*; 2.7 – 10 keV for *NICER*.

Mission	<i>NuSTAR</i>		<i>NICER</i>	
	FPM A	FPM B		
ObsID	80401312002		1103010157	1103010158
Start time	12:01:09		11:42:40	23:49:26
End time	05:31:09		22:55:20	05:05:40
Net count rate / s ⁻¹	60.2	59.6	196.8	
Exposure time / s	26166	26512	15386	5033

the oscillation which prevents more direct approaches such as phase-folding. We therefore employ the techniques pioneered by Ingram & van der Klis (2015) to consider the QPO waveform in different energy bands, considering its phase-average and first two harmonics. To do this, we extract three key pieces of information:

- The amplitude (RMS) of each harmonic in each energy band $\sigma_j(E)$. We find this by fitting an estimate of the power spectrum with a multi-Lorentzian model, as described in Section 2.3.3.
- The phase lag of each harmonic in each energy band $\Delta_j(E)$, relative to the phase of the corresponding harmonic in a reference band. This comes from using the cross spectrum between the lightcurve of the subject energy band, and the reference band, as described in Section 2.3.2.
- The phase difference ψ between the first two harmonics measured within the reference band. Using the FFT of the reference lightcurve, this is the difference between the phases in frequency bins containing the two harmonics. We use the bi-spectrum to calculate this, which is described in Section 2.3.4.

Following Ingram & van der Klis (2015) and Ingram et al. (2016), we consider that the count rate $w(E, \gamma)$ in each energy bin denoted by E varies with QPO phase γ as

$$w(E, \gamma) = \mu(E) \left[1 + \sum_{j=1}^J \sigma_j(E) \cos(j\gamma - \Phi_j(E)) \right], \quad (2.1)$$

where $\mu(E)$ is the average count rate in the energy band, $\sigma_j(E)$ is the average RMS of the j^{th} QPO harmonic², and $\Phi_j(E)$ is the phase offset of the j^{th} QPO harmonic. The phase-offset is split into an energy-dependent phase-lag $\Delta_j(E)$, which is the phase lag of a harmonic compared to the same harmonic in the reference band light curve, and the phase difference ψ between the two harmonics in the same reference band light curve. These are combined so that

$$\begin{aligned}\Phi_1(E) &= \Phi_1 + \Delta_1(E) \\ \Phi_2(E) &= 2(\Phi_1 + \psi) + \Delta_2(E),\end{aligned}\tag{2.2}$$

where, following Ingram & van der Klis (2015), we choose to set the arbitrary phase of the first harmonic to $\Phi_1 = \pi/2$.

Putting this together, for $j \geq 1$ we get the Fourier transformed (FT) spectra (Ingram et al. 2016)

$$W_j(E) = \mu(E)\sigma_j(E)e^{i\Phi_j(E)},\tag{2.3}$$

plus the phase-average $W_0(E) = \mu(E)$, which is trivially the flux-energy spectrum. We fit the theoretical model described in the following section simultaneously to the real and imaginary parts of the $j = 1$ and $j = 2$ FT spectra, and the flux-energy spectrum ($j = 0$). We use the full spectral resolution of the instrument for the flux-energy spectrum (grouped to have ≥ 30 counts in each energy channel) and subtract background. For the $j \geq 1$ QPO harmonics, we instead use the broader energy bands defined in Sections 2.2.1 and 2.2.2 and do not perform a background subtraction. This treatment of the background is appropriate because the background is not expected to be variable on the QPO period.

The following subsections which describe these parts of the analysis are each further split into two sub-subsections, with the first describing the method used, and the second presenting the results obtained from the observations analysed in this paper.

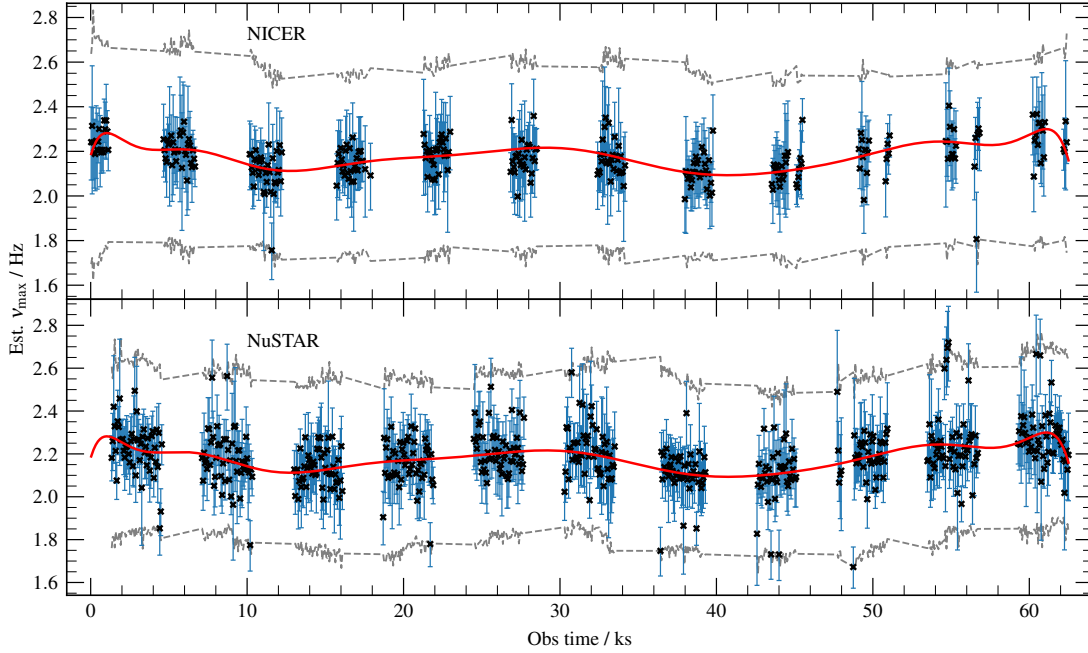


Figure 2.3: The estimated value of the QPO frequency ν_{\max} measured in each 64 s long segment of the two light curves, with the error bars denoting the estimated FWHM. The grey dashed line show the dynamic bounds for the QPO tracking algorithm. The red line is a 15 deg polynomial simultaneously fit to the estimated QPO frequencies, which we then use to get the QPO frequency at arbitrary time.

2.3.1 QPO frequency tracking

During the observations, the frequency of the QPO drifts over time by $\sim \pm 5\%$. To account for this, we identify the frequency of the QPO during each of the M segments so that we can later average Fourier products over the frequency range containing the instantaneous QPO frequency.

2.3.1.1 Method

We determine the QPO frequency for each segment by fitting a model to the power spectrum of each segment of the *NICER* and *NuSTAR* (we sum the FPMA and FPMB counts) light curves. Our model comprises of four Lorentzian functions (van Straaten et al. 2002), two of which are harmonically locked to represent the two QPO harmonics, and a Poisson noise component. The Poisson noise component

²The average RMS is often reported as $\langle \sigma_j(E) \rangle$, with an additional multiple of $\sqrt{2}$ included in Eq. 2.1. This is to make explicit that the average RMS of a sine-wave is $1/\sqrt{2}$. For simplicity we neglect this, keeping in mind our model is normalised to expect the average RMS.

is very simple for *NICER*, taking a constant value of $2/\mu$ where μ is the mean count rate (van der Klis 1989). The Poisson noise is much more complicated for *NuSTAR*, due to the large detector dead time, t_D . Dead time is known to change the shape of the Poisson noise within the power spectrum from a constant value to the general form (Zhang et al. 1995)

$$f(\nu|A, B, t_D) = A - 2Bt_D \text{sinc}(2\pi\nu t_D). \quad (2.4)$$

Since there is currently no accurate dead time model for *NuSTAR*, we follow (Bult 2017) and model the Poisson noise with this sinc function, and determine the parameters A , B and t_D by fitting to the power spectrum averaged over the entire observation. We only consider $\nu > 20$ Hz since this frequency range is Poisson noise dominated. This fit yields $A = (1.783 \pm 0.006) \times 10^{-2}$, $B = 0.15 \pm 0.01$ and $t_D = 3.3 \pm 0.2$ ms. We then freeze A , B and t_D to these best fitting values for the remainder of the analysis.

As each frequency bin of a single power spectrum (without any ensemble or frequency averaging) follows a χ^2 probability distribution with two degrees of freedom³ (van der Klis 1989), we are unable to use χ^2 minimisation for the model fit as this depends on each frequency bin following a Gaussian distribution. We therefore use the maximum likelihood estimation method described in Barret & Vaughan (2012). This method does not require the probability distribution to be Gaussian, it only requires it to be known analytically. We find best fitting model parameters, including the frequency of the QPO fundamental, by maximising the likelihood function calculated assuming a χ^2 probability distribution with two degrees of freedom.

It is important to note that the probability distribution underlying each frequency bin of a un-averaged co-spectrum is not known analytically (Huppenkothen & Bachetti 2018). We are therefore unable to use the maximum likelihood method on the co-spectrum between FPMA and FPMB, which is why we resort to modelling the Poisson noise of the power spectrum, which does have well-understood statistics.

³Apart from the bin at the Nyquist frequency which is a χ^2 distribution with only a single degree of freedom, but this bin was ignored in our calculations for simplicity

2.3.1.2 Results

Fig. 2.3 shows the resulting measurements of QPO frequency (black crosses) as a function of time for *NuSTAR* (top) and *NICER* (bottom). Rather than uncertainties, the error bars shown are the measured full width at half maximum (FWHM) of the Lorentzian function representing the QPO fundamental. In our fits, we restricted the QPO frequency to a range given by the running average of the QPO frequencies of the previous 5 segments $\pm 3/2$ times the running average FWHM of the previous 5 segments, starting with the QPO frequency and FWHM from the average power spectrum. This range is represented by the grey dashed lines. To smooth out the results of our tracking algorithm⁴ we use a degree-15 polynomial to model QPO frequency with time, that we fit simultaneously to the *NICER* and *NuSTAR* data (the solid red line in Fig. 2.3). It is encouraging that the instantaneous QPO frequencies we measure here are very similar to those inferred by Huppenkothen & Bachetti (2021) using a more sophisticated method (see their fig. 23).

2.3.2 Phase lag spectrum

2.3.2.1 Method

In order to calculate the energy-dependent phase lag of the j^{th} harmonic $\Delta_j(E)$, we first calculate the cross-spectrum between the light curve of the energy band centred on energy E (the subject band) and the light curve summed over all energy channels (the reference band) for each of the M segments. For the m^{th} segment, the cross-spectrum as a function of frequency is

$$G_m(\nu, E) \propto S_m(\nu, E)R_m^*(\nu) \quad (2.5)$$

where $R_m(\nu)$ and $S_m(\nu, E)$ are the Fourier transforms⁵ of the m^{th} segment of the reference and subject band light curves respectively, and the constant of proportionality is a normalisation into fractional RMS. For *NICER*, the photons in

⁴The difference in the scatter of QPO frequencies between the lower count rate *NuSTAR* and higher count rate *NICER* data suggests that this is noise in the measurement rather than intrinsic short timescale changes.

⁵As we are using the FFT, we actually use discrete frequency bins $\nu_k = k/T$.

the subject band light curve are also in the reference band light curve (since the reference band consists of all the photons detected by *NICER*). This contributes Poisson noise, which we subtract off following Ingram (2019).

For *NuSTAR*, we instead extract the subject band light curves from the FPMB and the reference band light curve from the FPMA to ensure that the subject and reference band signals are statistically independent of one another, and therefore the cross-spectrum contains no contribution from dead time affected Poisson noise.

We consider the ‘shifted-and-added’ cross-spectrum of the j^{th} QPO harmonic by first averaging over the QPO harmonic in each time segment based upon the tracked QPO frequency. For the m^{th} segment we average over the frequency range $\nu = j\nu_{\text{qpo}}(m)[1 \pm 1/(2Q)]$, where we assume $Q = 8$ for the quality factor (a typical value for a type-C QPO; Ingram & Motta 2019), using the smoothed estimate from our QPO tracking algorithm for $\nu_{\text{qpo}}(m)$. We then average over the time segments to find the overall average value for cross-spectrum for each of the QPO harmonics.

The phase lag for the j^{th} harmonic, $\Delta_j(E)$, is the argument of the averaged cross-spectrum of the corresponding harmonic, of which we estimate the uncertainties using the formula from Ingram (2019, eq 19).

2.3.2.2 Results

We display the measured phase lag spectrum in Fig. 2.4 (bottom) for *NICER* (blue) and *NuSTAR* (orange). We see that the phase lag of the QPO fundamental is almost constant with energy, which is consistent with previous *RXTE* observations showing that the phase lag monotonically reduces from hard lags (a positive lag vs energy gradient) to soft lags (a negative lag vs energy gradient) as the QPO frequency increases, with the cross over occurring for $\nu_{\text{qpo}} \sim 2$ Hz (e.g. van den Eijnden et al. 2017). Note that a slight offset between *NICER* and *NuSTAR* lags results from a phase lag between the *NICER* and *NuSTAR* reference bands. Although the offset is very small for this observation because the energy dependence of the lag is subtle, we account for it in our modelling with a floating phase offset.

2.3.3 Fractional RMS Spectrum

2.3.3.1 Method

To calculate the fractional RMS of the QPO harmonics in each energy band, instead of the power spectrum in that energy band, we boost signal to noise by using the shift-and-added cross-spectrum $G(\nu, E) = \frac{1}{M} \sum_m^M G_m(\nu + \delta\nu_m, E)$, where $\delta\nu_m$ is the difference between the QPO frequency in that segment (from our smoothed tracking) and the average QPO frequency.

Since the QPO harmonics are well correlated between energy bands, we assume unity coherence between the subject bands and the reference band, in which case the shifted-and-added power spectrum of the subject band can be written as (Wilkinson & Uttley 2009; Ingram et al. 2016)

$$P_s(\nu, E) = \frac{|G(\nu, E)|^2 - \hat{b}^2(\nu, E)}{P_r(\nu)}, \quad (2.6)$$

where $\hat{b}^2(\nu, E)$ results from a positive-bias in the calculation of $|G(\nu, E)|^2$ (see Ingram 2019 for details, where b is used instead of \hat{b}) and $P_r(\nu)$ is the Poisson noise subtracted shift-and-added power spectrum of the reference band for *NICER* and the shift-and-added co-spectrum between the FPMA and FPMB for *NuSTAR*.

We fit a multi-Lorentzian model to the resulting power spectral estimate for each energy band. We use three Lorentzian functions: one with $Q = 0$ to represent the broad band noise; and the other two with centroid frequencies and Q tied (such that the centroid frequencies are harmonically related) in order to represent the two QPO harmonics. We normalise our power spectral estimate (Belloni & Hasinger 1990) and Lorentzian functions (van Straaten et al. 2002) such that the best-fitting normalisation of the two Lorentzian components representing the QPO harmonics gives their fractional RMS. We calculate 1σ uncertainties on the RMS by searching parameter space for a marginalised $\Delta\chi^2 = 1$.

2.3.3.2 Results

We display the measured RMS spectrum in Fig. 2.4 (top). Error bars without a lower cap correspond to points consistent with zero within 1σ . We see that, as is typically




Figure 2.4: Top: The fractional RMS of the first and second harmonics of the QPO in both the *NICER* and *NuSTAR* observations in different energy bands, found by fitting lorentzian functions to the power-spectrum of each energy band, which was estimated from a cross-spectrum between a broad reference energy band and a specific energy band, as described in the text. As the fractional RMS must be positive points that have $\Delta\chi^2 < 1$ at zero are shown without an errorbar cap. The fractional RMS is slightly diluted by background photons in the very highest *NICER* energy bands. Bottom: The phase lag of the first and second harmonics of the QPO, measured against a reference band. For *NICER* this reference band is the full energy lightcurve, while for *NuSTAR* it is the full energy lightcurve of FMPB.

the case for type-C QPOs, the fractional RMS increases with energy for $E \lesssim 10$ keV before levelling off. We note good agreement between *NICER* and *NuSTAR*.

2.3.4 Phase difference between harmonics

2.3.4.1 Method

As we have the phase lag of each QPO harmonic in the energy bands compared to their counterpart in the reference band, we now need to find the phase lag between the harmonics in the reference band ψ . For this, we use the bispectrum⁶ as

⁶We note this method can find the phase-difference between any doublet harmonics, e.g. also between the 2nd and 4th harmonics. In much the same way, higher order polyspectra could be used to find the phase difference between other harmonics, e.g. the trispectrum could be used to

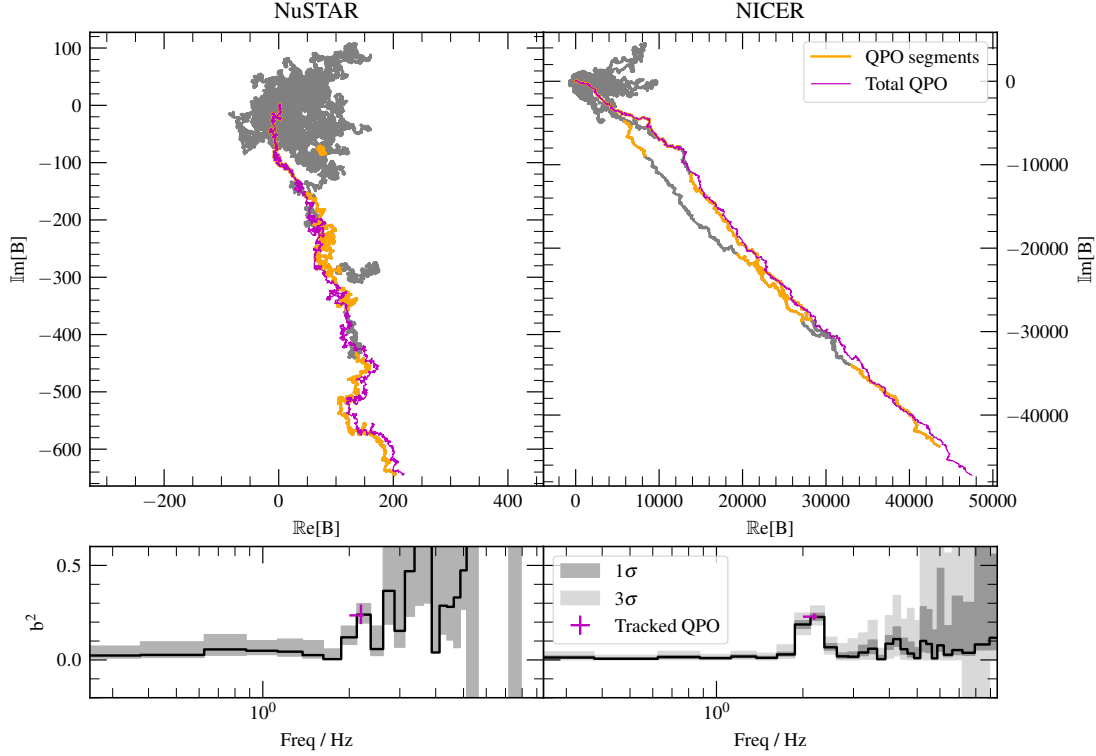


Figure 2.5: The random walk of the auto-bispectrum $B(\nu_k)$ for frequency bins ν_k which cover the frequency range $\nu_k < 10$ Hz. Between 6-10 Hz the frequency bins are geometrically re-binned to reduce the number of bins by a factor of 4. The upper panels show traces of $B_m(\nu_k)$ in three colours: grey, orange, and magenta. Each grey trace corresponds to a frequency, and is made up of M short segments, each corresponding to a segment in the light curve. Starting from the origin, the calculated values of $B_m(\nu_k)$ are consecutively appended in the complex plane (such that $B_{m+1}(\nu_k)$ begins at the end of $B_m(\nu_k)$). For each trace, each segment corresponds to a particular time in the observation, and therefore the frequency of the QPO can be identified. The segments which correspond to a time where the frequency bin contains the QPO $B_m([\nu_{\text{QPO},m}]_k)$ are highlighted in orange. Finally, these QPO segments are chained together, as if they were a single frequency bin, which is displayed in magenta. On the lower panel, the auto-bicoherence $b^2(\nu_k)$ of each chain is given for each frequency. The grey uncertainties, and also the vertical 1σ uncertainty of the tracked QPO auto-bicoherence, are calculated from percentiles of a bootstrapped population.

suggested by Arur & Maccarone (2019). This yields similar results to the method of Ingram & van der Klis (2015) used in Ingram et al. (2016), Ingram et al. (2017) and de Ruiter et al. (2019), but is more statistically robust. In particular there is no Poisson noise correction in the Ingram & van der Klis (2015) method, whereas the effect of Poisson noise on the bispectrum is well covered in the literature (e.g. Wirnitzer 1985; van der Klis 1989; Kovach et al. 2018). The bispectrum method should therefore be more robust to low count rate observations. The bispectrum is defined as a function of two frequencies, ν_1 and ν_2 , such that the bispectrum of the reference band light curve in the absence of Poisson noise is

$$\mathcal{B}(\nu_1, \nu_2) = \frac{1}{M} \sum_{m=1}^M R_m(\nu_1) R_m(\nu_2) R_m^*(\nu_1 + \nu_2). \quad (2.7)$$

The phase-difference between harmonics can be retrieved from the ‘auto-bispectrum’, $B(\nu) = \mathcal{B}(\nu_1 = \nu, \nu_2 = \nu)$, which is

$$B(\nu) = \frac{1}{M} \sum_{m=1}^M B_m(\nu) = \frac{1}{M} \sum_{m=1}^M R_m(\nu) R_m(\nu) R_m^*(2\nu). \quad (2.8)$$

Since $R_m(\nu) = |R_m(\nu)| e^{i\Phi_m(\nu)}$, we see that

$$B(\nu) = \frac{1}{M} \sum_{m=1}^M |R_m(\nu)|^2 |R_m(2\nu)| e^{i[2\Phi_m(\nu) - \Phi_m(2\nu)]}. \quad (2.9)$$

Therefore the bi-phase $\equiv \arg[B(\nu)] = 2\Phi(\nu) - \Phi(2\nu) = -2\psi(\nu)$, where we take the phase difference between frequency ν and frequency 2ν to be $\psi(\nu) = \frac{1}{2}\Phi(2\nu) - \Phi(\nu)$ following Ingram & van der Klis (2015). The phase difference between the two QPO harmonics is therefore $\psi = -\frac{1}{2}\arg[B(\nu_{\text{qpo}})]$.

In order to calculate $B(\nu_{\text{qpo}})$, we adopt the same shift-and-add technique for the auto-bispectrum as described in the previous section for the cross-spectrum, again employing our smoothed estimate for the instantaneous QPO frequency from Fig. 2.3. The QPO is only coherent on timescales of $\sim Q$ cycles (e.g. van den Eijnden et al. 2016)], and so it is reasonable to use segments of duration $T \sim Q/\nu_{\text{qpo}}$ (Ingram & van der Klis 2015). As the QPO frequency in our observations is ~ 2.2 Hz, and with an assumed quality factor $Q = 8$, we use 4 s long segments. Using $\delta t = 1/128$ s, this

find the phase difference between the 1st and 3rd harmonics.

gives segments with $N = 512$ time bins. We correct the *NICER* auto-bispectrum for Poisson noise as described in Appendix A.1.1. For *NuSTAR*, we avoid dead time effects by using FPMA data for $R(\nu)$ and FPMB data for $R(2\nu)$.

2.3.4.2 Results

Using the bi-spectrum method described above, we measure $\psi/\pi = 0.20 \pm 0.02$ for *NuSTAR* and $\psi/\pi = 0.125 \pm 0.006$ for *NICER*. The slight difference in ψ between observatories is statistically significant, and indicates that the QPO waveform depends on photon energy.

We compare the results using this method to the method used in the literature (e.g. Ingram et al. 2016) where $\Phi(\nu)$ and $\Phi(2\nu)$ are taken directly from the FFT, and a minimisation is used to find ψ . For *NuSTAR* this gives $\psi/\pi = 0.21 \pm 0.04$, and for *NICER* this gives $\psi/\pi = 0.12 \pm 0.01$, which are consistent with the values from our updated method. Here, we again avoid the dead time affected *NuSTAR* Poisson noise by taking $\Phi(\nu)$ from the FFT of the FMPA light curve, but $\Phi(2\nu)$ from the FFT of the FMPB light curve.

As discussed in Ingram & van der Klis (2015), it only makes sense to measure the phase difference between harmonics if the phases of those two harmonics are correlated. In such a case, the QPO has some well-defined underlying waveform, and it makes sense to do phase-resolved spectroscopy. Otherwise, the spectrum does not vary in shape in a systematic way with QPO phase, and a QPO phase-resolved analysis would thus be meaningless. It is therefore important to measure how well correlated the two dominant QPO harmonics are before continuing. We explore this in the following sub-section.

2.3.4.3 Phase correlation between QPO harmonics

The auto-bispectrum can also be used to measure the extent to which the phases of the two harmonics are correlated. Specifically, the auto-bicoherence is the modulus of the auto-bispectrum, re-normalised in some useful way. We use the Kim & Powers (1979) normalisation for which the auto-bicoherence, $b^2(\nu)$, is unity if the phases of the ν and 2ν components are perfectly correlated. In the opposite case of

a completely uncorrelated signal, $b^2(\nu) \rightarrow 0$ as the number of light curve segments used to calculate b^2 tends to $M \rightarrow \infty$, and it becomes meaningless to measure the phase difference between the harmonics at ν and 2ν .

In the absence of Poisson noise the auto-bicoherence is given by (Kim & Powers 1979)

$$b^2(\nu) = \frac{\left| \sum_{m=1}^M B_m(\nu) \right|^2}{\sum_{m=1}^M |R_m(\nu) R_m(\nu)|^2 \sum_{m=1}^M |R_m(2\nu)|^2}. \quad (2.10)$$

We describe how we account for Poisson noise and dead time effects in the denominator of the above equation in Appendix A.1.1, and also describe a bootstrapping technique (following Stevens & Uttley 2016) we use to calculate the errors on the auto-bicoherence in Appendix A.1.2.

The auto-bicoherence as a function of frequency is shown in Fig. 2.5 (bottom) for *NuSTAR* (left) and *NICER* (right). The black stepped lines show the measured values and the shaded regions represent the 1 and 3 σ confidence regions (calculated using a bootstrapping method). We see that the auto-bicoherence is consistent with zero for all frequencies except for around the QPO fundamental frequency. This indicates that the phase of the first harmonic is well correlated with that of the second harmonic, and that there is therefore an underlying QPO waveform (Ingram & van der Klis 2015; de Ruiter et al. 2019). In contrast, all other pairs of frequencies are uncorrelated.

The width of the feature around the QPO frequency in these plots of the bispectrum could be because the QPO frequency drifts during the observation, and therefore the frequency bin that contains the QPO frequency may not be the same for each segment. We therefore calculate $b^2(\nu_{\text{qpo}})$ using the same shift-and-add techniques as described in the previous sub-sections. The result is marked in magenta. The vertical error bar is 1σ and again calculated by bootstrapping (as discussed in Appendix A.1.2), and the horizontal error bar shows the frequency range covered by the QPO fundamental during the observation. We see that this new shifted-and-added bicoherence is consistent with the black stepped line, and therefore the shift-and-add does not enhance the coherence of the QPO for this particular observation. While the coherence between the QPO harmonics is

statistically non-zero, it is still not unity. This can be partly explained by the presence of uncorrelated broad band noise at the QPO frequency. However, the QPO dominates over the broad band noise in the power spectrum for $\nu \approx \nu_{\text{qpo}}$. It is therefore likely that ψ is not constant in time – as would be the case for a perfectly periodic oscillation – but instead varies around a well-defined mean value.

We visualise the auto-bispectrum in Fig. 2.5 with ‘jellyfish plots’ (top panels). Here, for each Fourier frequency, we plot $B(\nu)$ as a vector sum on the complex plane. The vector sum for each frequency is plotted as a grey line (and there are a total of 29 frequencies after geometric re-binning above 6 Hz). We see that this forms a random walk on the complex plane. For most frequencies, this random walk forms a blob that never gets far from the origin, indicating that the phase at ν is poorly correlated with that at 2ν . For a narrow range of frequencies, the random walk instead forms a much straighter path that extends far from the origin, indicating good correlation between ν and 2ν . The auto-bicoherence is a measure of how far from the origin the vector sum ends up, and the biphase is a measure of the orientation on the complex plane of the summed vector. The jellyfish plots therefore demonstrate that $b^2(\nu)$ is large for $\nu \approx \nu_{\text{qpo}}$ because each segment has a similar bi-phase, and so the segments all line up well on the complex plane.

We also demonstrate how the drifting of the QPO frequency during the observation leads to the frequency bin that contains the QPO changing from one segment to another. To do this, we mark segments on the jellyfish plots that contain the instantaneous QPO frequency by colouring them orange. We see that the QPO frequency does indeed jump from one frequency bin to another. This is particularly noticeable in the *NICER* jellyfish plot. The magenta line instead shows the vector sum of $B(\nu_{\text{qpo}})$ that we use to calculate $b^2(\nu_{\text{qpo}})$; i.e. we take only the segments that contain the instantaneous QPO frequency and add them on the complex plane. Consistent with the bicoherence plots, we see that the magenta line reaches a comparable distance from the origin to the two grey lines around the QPO frequency.

2.3.5 Reconstructed Fourier Transformed Spectra

The above analysis extracted key information about the two QPO harmonics in different energy bands. In order to examine how the energy spectrum is modulated with QPO phase, we use Eqs. 2.2 and 2.3 to combine the three quantities we have extracted from the energy resolved light curves: the phase lag spectra $\Delta_j(E)$, the fractional RMS spectra $\omega_j(E)$, and the phase difference between the harmonics ϕ . This produces the Fourier transformed spectra $W_j(E)$ for the first two harmonics ($j = 1, 2$). As $W_j(E)$ are complex quantities, we separate these into real and imaginary parts $\Re[W_j(E)]$ and $\Im[W_j(E)]$. Therefore, for both of our *NuSTAR* and *NICER* observations we have calculated 4 spectra

1. $\Re[W_1(E)]$, the real part of the FT spectrum of first QPO harmonic,
2. $\Im[W_1(E)]$, the imaginary part of the FT spectrum of first QPO harmonic,
3. $\Re[W_2(E)]$, the real part of the FT spectrum of second QPO harmonic,
4. $\Im[W_2(E)]$, the imaginary part of the FT spectrum of second QPO harmonic,

giving a total of 8 FT spectra. We also have the phase-average flux-energy spectra ‘ $W_0(E)$ ’ for each of *NuSTAR*’s FMPA and FMPB, plus *NICER*, bringing our total to 11 spectra which we will simultaneously fit with the model described in the following section.

These 11 spectra are shown in Fig 2.6. The left panel shows the phase-averaged flux-energy spectrum, $W_0(E)$, observed by *NICER* (black), FPMA (red) and FPMB (blue). The derived FT spectra are plotted on the right hand side as grey and black points. The first and second harmonic (i.e. $j = 1$ and $j = 2$, or in other words the fundamental and the first overtone) are plotted respectively in the first and second panels from the top (as labelled), and grey and black points correspond to the real ($\Re[W_j(E)]$) and imaginary ($\Im[W_j(E)]$) parts respectively. Open circles correspond to *NuSTAR* and the points with no marker to *NICER*. Note that each part of each harmonic has only one *NuSTAR* FT spectrum, and not one for the FPMA and another for the FPMB. This is because the *NuSTAR* FT spectra are derived by

extracting the subject bands of the cross-spectrum from the FPMB and the reference band from the FPMA. Both FPMs are therefore used for this single measurement.

2.4 Theoretical Model

Our model for the QPO FT calculates the X-ray spectrum as a function of QPO phase, γ , before Fourier transforming to output the real and imaginary parts of the QPO FT for the zeroth, first and second harmonics. The model is similar to the one described in Ingram et al. (2017), but with some extra features. As in Ingram et al. (2017), we assume that the accretion flow has two components: a thin accretion disc and a corona. We assume that the disc is stationary with inner and outer radii r_{in} and r_{out} . We make no assumptions about the shape of the corona, but we assume that its intrinsic bolometric luminosity is constant in time, and variation in the observed flux is caused by us viewing the corona from different directions at different QPO phases.

The spectrum from the corona is described in Section 2.4.1. Section 2.4.2 describes the spectrum from the disc, which we split into a thermal component (Section 2.4.2.2) and a non-thermal component (Section 2.4.2.3).

2.4.1 Corona

We represent the corona spectrum with the model NTHCOMP (Zdziarski et al. 1996), which is a power-law (index Γ – such that the *photons* emitted per unit energy is $\propto E^{-\Gamma}$) between low and high energy cut-offs that are respectively governed by the seed photon temperature T_{bb} and the electron temperature T_e . We parameterise the QPO phase-dependent bolometric flux observed from the corona as

$$N_c(\gamma) = N_0 + A_{1N} \sin[\gamma - \phi_{1N}] + A_{2N} \sin[2(\gamma - \phi_{2N})], \quad (2.11)$$

where N_0 , A_{1N} , A_{2N} , ϕ_{1N} and ϕ_{2N} are left as free parameters. We parameterise the photon index $\Gamma(\gamma)$ and electron temperature $T_e(\gamma)$ in a similar way (see e.g. Eq. 5 of Ingram et al. 2017). We tie T_{bb} to the peak disc temperature, which we discuss at the end of Section 2.4.2.2.

2.4.2 Disc

The corona irradiates the disc, and since the disc is very optically thick, all of the irradiating flux is reprocessed in the disc atmosphere and re-emitted. The disc is also heated by viscous dissipation of gravitational potential energy, generating intrinsic disc flux. The total radiated flux is the sum of intrinsic and reprocessed flux. We assume that all of the intrinsic flux plus some fraction of the reprocessed flux is in thermal equilibrium with the disc, thus contributing a blackbody component to the emitted spectrum. The total restframe specific intensity emergent from the disc coordinate (r, ϕ) at QPO phase γ is therefore the sum of this blackbody component, and a non-thermal component which includes well-known ‘reflection’ features such as the iron line and Compton hump

$$I(E_d, r, \phi, \gamma) = I_{\text{bb}}(E_d, r, \phi, \gamma) + I_{\text{nt}}(E_d, r, \phi, \gamma), \quad (2.12)$$

where E_d is photon energy in the restframe of disc coordinate (r, ϕ) . From this, we calculate the observed specific disc flux by tracing rays from the disc to the observer following null-geodesics in the Kerr metric (using the code YNOGK, which is based on GEOKERR: Yang & Wang 2013; Dexter & Agol 2009). A summary of the ray tracing procedure can be found in Appendix A.2.1 (or see e.g. Ingram et al. 2019 for a more detailed description.)

2.4.2.1 Illumination of the disc

A precessing corona will preferentially illuminate different disc azimuths at different phases of its precession cycle. Instead of making assumptions about the shape of the corona, we follow Ingram et al. (2017) by parameterising the QPO phase-dependent illuminating flux as a function of disc radius and azimuth with the emissivity function, $\epsilon(r, \phi, \gamma)$, such that

$$\frac{I_{\text{nt}}(E_d, r, \phi, \gamma)}{D^2} = f_{\text{R}}(\gamma) N_{\text{c}}(\gamma) \epsilon(r, \phi, \gamma) \mathcal{R}(E_d), \quad (2.13)$$

where D is the distance from the observer to the BH. Here, we normalise $\epsilon(r, \phi, \gamma)$ and $\mathcal{R}(E_d)$ (Eqs. A.8 and A.7 in Appendix A.2.2) such that $f_{\text{R}}(\gamma)$ is the *observer’s*

reflection fraction. This is defined by Ingram et al. (2019) as the observed bolometric reflected flux divided by the directly observed bolometric coronal flux in the simplified case in which the disc re-emits the incident radiation isotropically. In this case, since $N_c(\gamma)$ is defined as the directly observed bolometric coronal flux, the observed bolometric reflected flux is simply $f_R(\gamma)N_c(\gamma)$. In reality, the reflected flux is *not* emitted isotropically. The function $\mathcal{R}(E_d)$, which we discuss below, includes this subtlety.

We employ the following form for the emissivity function

$$\epsilon(r, \phi, \gamma) = \mathcal{N}_\epsilon \epsilon(r) \left\{ 1 + A_1 \cos^2 \left[\frac{1}{2}(\gamma - \phi + \phi_1) \right] + A_2 \cos^2[\gamma - \phi + \phi_2] \right\}, \quad (2.14)$$

where the radial dependence is given by a twice broken power-law (Wilkins & Fabian 2011, 2012)

$$\epsilon(r) = \begin{cases} (r/r_{\text{br},1})^{-q_1} & \text{if } r \leq r_{\text{br},1} \\ (r/r_{\text{br},1})^{-q_2} & \text{if } r_{\text{br},1} < r \leq r_{\text{br},2} \\ (r_{\text{br},2}/r_{\text{br},1})^{-q_2} (r/r_{\text{br},2})^{-3} & \text{if } r > r_{\text{br},2}. \end{cases} \quad (2.15)$$

This way, if $A_1 = A_2 = 0$, the reflection spectrum will not depend at all on QPO phase, since the illumination pattern on the disc becomes axi-symmetric. When the asymmetric illumination (‘asymmetry’) parameters A_1 and A_2 are non-zero, there are instead bright patches on the disc that rotate around the disc rotation axis once per QPO cycle, with the location on the disc of the peak brightness set by the phase parameters ϕ_1 and ϕ_2 . Specifically, $A_1 > 0$ and $A_2 = 0$ will lead to one bright patch rotating about the disc surface normal, and $A_1 = 0$ and $A_2 > 0$ will lead to two identical bright patches (see Ingram et al. 2017 for more details). The normalisation constant, \mathcal{N}_ϵ , is set by Eq. A.8. We also parameterise the reflection fraction as a sum of sinusoids in the form of Eq. 2.11.

2.4.2.2 Thermal flux

We assume that the blackbody component of the disc specific intensity is given by

$$\frac{I_{\text{bb}}(E_d, r, \phi, \gamma)}{D^2} = N_d B(E_d, T(r, \phi, \gamma)), \quad (2.16)$$

where N_d is a constant model parameter, $B(E, T)$ is the Planck function, and

$$T(r, \phi, \gamma) = \left[T_{\text{visc}}^4(r) + T_{\text{irr}}^4(r, \phi, \gamma) \right]^{1/4}. \quad (2.17)$$

Here, $T_{\text{visc}}(r)$ is the ‘intrinsic’ disc temperature; i.e. the temperature in the absence of irradiation. $T_{\text{irr}}(r, \phi, \gamma)$ is the ‘irradiation’ disc temperature; i.e. the temperature in the absence of viscous dissipation. This parameterisation of the temperature stems from the Stephen-Boltzmann law, as if the thermal disc flux is from thermalised gas being heated by both the viscous heating (with flux $\sigma_{\text{SB}}T_{\text{visc}}$), and coronal photons thermalising in the disc atmosphere (with flux $\sigma_{\text{SB}}T_{\text{irr}}$).

We set the intrinsic disc temperature as

$$T_{\text{visc}}(r) \propto \frac{T_{\text{visc,max}}}{(r^2 dA/dr)^{1/4}}, \quad (2.18)$$

where the constant of proportionality is set to ensure that the maximum temperature is $T_{\text{visc,max}}$ and the relativistic expression for dA/dr is given by e.g. Equation A1 in Ingram et al. (2019). In Newtonian gravity, $dA/dr = 2\pi r$, and so the familiar $T_{\text{visc}} \propto r^{-3/4}$ emissivity of a simple disc model is recovered⁷. Note that any colour-temperature correction factor accounting for the spectrum not being strictly blackbody is simply swallowed up into the definition of $T_{\text{visc,max}}$.

The irradiation temperature is related to the thermalised portion of the illuminating flux via the Stefan-Boltzmann law. We parameterise it as

$$T_{\text{irr}}^4(r, \phi, \gamma) \propto f_{\text{R}}(\gamma - \Delta\gamma) N_c(\gamma - \Delta\gamma) \epsilon(r, \phi, \gamma - \Delta\gamma), \quad (2.19)$$

as through our construction the product of these three terms is the illuminating flux on the disc, although only a small fraction of this flux will thermalise. Defining $T_{\text{irr,max}}(\gamma)$ as the maximum value of $T_{\text{irr}}(r, \phi, \gamma)$ for a given QPO phase, we set the constant of proportionality in the above equation to ensure that the QPO phase-averaged value of $T_{\text{irr,max}}(\gamma)$ is proportional to the model parameter kT_i . With this parameterisation, kT_i effectively sets the fraction of the illuminating flux that thermalises in the disc atmosphere. The model parameter $\Delta\gamma$ accounts

⁷Note that we do not employ a stress free inner boundary condition, which is appropriate if the corona is located inside of the disc, providing a torque.

for the *thermalisation timescale*, which is the time it takes for the irradiating flux to thermalise in the disc. This timescale is currently poorly understood, but it should lead to the thermal component responding to changes in the illuminating flux with a delay compared to the non-thermal component (e.g. emission lines: García et al. 2013b). Here we parameterise this thermalisation timescale such that the current disc temperature depends on what the irradiating flux was some time $\Delta\gamma/(2\pi\nu_{\text{qpo}})$ ago. This is an extremely simplified formalism. In reality, we may expect $T_{\text{irr}}(\gamma)$ to be smeared as well as delayed, and we may also expect the delay itself to depend on e.g. disc radius.

Finally, we set the seed photon temperature in NTHCOMP to

$$T_{\text{bb}}(\gamma) = \left[T_{\text{visc,max}}^4 + T_{\text{irr,max}}^4(\gamma) \right]^{1/4}. \quad (2.20)$$

2.4.2.3 Non-thermal flux

We use the model XILLVERCP to calculate the non-thermal component of the restframe emergent disc spectrum, $\mathcal{R}(E_d)$. XILLVERCP calculates the emergent spectrum from a passive ($T_{\text{visc}} = 0$), constant density slab (electron number density $n_e = 10^{15} \text{ cm}^{-3}$) being irradiated by an NTHCOMP spectrum⁸. The output spectrum includes emission lines (most prominently the iron $K\alpha$ line), absorption edges (most prominently the iron K edge) and the Compton hump. It also includes a quasi-thermal component caused by some fraction of the irradiating photons thermalising in the disc, which we must ignore because we have already accounted for the thermalised illuminating flux in our blackbody component (see previous section). For XILLVERCP, this component peaks in the UV and is entirely below our bandpass, and so is simple to ignore.

An important input parameter of XILLVERCP is the ionisation parameter $\xi = 4\pi F_x/n_e$, where F_x is the illuminating X-ray flux. We set the ionisation parameter

⁸Note that XILLVERCP is calculated for an irradiating spectrum given by NTHCOMP with the seed photon temperature hardwired to $kT_{\text{bb}} = 0.05 \text{ keV}$, whereas we allow our continuum spectrum to have a seed photon temperature that is free to vary with QPO phase.

from our existing parameterisation of the illuminating flux such that⁹

$$\xi(r, \phi, \gamma) \propto f_{\text{R}}(\gamma) N_{\text{c}}(\gamma) \epsilon(r, \phi, \gamma) / n_{\text{e}}(r). \quad (2.21)$$

Defining $\xi_{\text{max}}(\gamma)$ as the maximum value of $\xi(r, \phi, \gamma)$ for a given QPO phase, we set the constant of proportionality in the above equation to ensure that the QPO phase-averaged value of $\xi_{\text{max}}(\gamma)$ is equal to the model parameter ξ_0 . For $n_{\text{e}}(r)$ we adopt the form corresponding to Zone A of the Shakura & Sunyaev (1973) disc model, following e.g. Ingram et al. (2019); Mastroserio et al. (2019); Shreeram & Ingram (2020): $n_{\text{e}}(r) \propto r^{3/2} [1 - \sqrt{r_{\text{in}}/r}]^{-2}$. This Zone A is the solution of the disc density profile in a region where the pressure within the disc is radiation dominated, and the opacity is dominated by electron scattering

2.4.2.4 Approximations

Our treatment of the restframe spectrum emergent from a given disc patch is very approximate. First of all, the XILLVERCP model that we use has $n_{\text{e}} = 10^{15} \text{cm}^{-3}$ hardwired, whereas we would theoretically expect the density of the disc in GRS 1915+105 to be closer to $n_{\text{e}} \sim 10^{20} \text{cm}^{-3}$ (Shakura & Sunyaev 1973). Second, the emergent spectrum is not strictly the sum of a thermal and a non-thermal component. In reality, the disc atmosphere is being irradiated from below by the intrinsic disc emission and from above by the corona, and the true emergent spectrum would need to be calculated by solving the radiative transfer equation with these boundary conditions (Reis et al. 2008). However, although high density XILLVER models are now available (García et al. 2016; Mastroserio et al. 2021), they only consider irradiation by the corona and ignore the intrinsic disc emission. Moreover, the XILLVER solutions are calculated in steady-state, and so there is no thermalisation timescale. Our treatment is therefore the best time-dependent approximation of a high density irradiated disc with strong intrinsic emission currently available.

⁹Computing the ionisation in this way does mean the restframe reflection spectrum is non-uniform across the disc, thus is formally $\mathcal{R}(E_{\text{d}}, r, \phi, \gamma)$.

2.4.3 The complete model

The total observed QPO phase-dependent specific flux, $F(E, \gamma)$, is the sum of the disc and coronal contributions. Our model calculates $F(E, \gamma)$ for 16 QPO phases and Fourier transforms to get the QPO FT for the zeroth (phase-average), first, and second harmonics.

Additionally, on top of the phase-dependent model described above, we include an extra `XILLVERCP` component to account for distant reflection. This component is fixed to be constant with QPO phase, as variations on the timescale of the QPO period should be strongly washed out by light-crossing delays for a distant reflector. Finally, we account for line-of-sight absorption with the model `TBABS`. The hydrogen absorption column N_{H} is a free parameter, and we adopt the abundances of all other elements relative to hydrogen of Wilms et al. (2000).

2.5 Model Fits

2.5.1 Fitting procedure

We fit our model simultaneously to 11 spectra overall. As described in Section 2.3.5, these are: the flux-energy *NICER*, FPMA and FPMB *NuSTAR* spectra, plus the first and second harmonics of the real and imaginary parts of the QPO FT measured separately by *NICER* and *NuSTAR*. Using `XSPEC` version 12.10.1f, we applied the same *NICER* response matrix to all 5 *NICER* spectra, we used the FPMA response for the flux-energy FPMA spectrum, and the FPMB response for both the FPMB flux-energy spectrum and the *NuSTAR* QPO FT. This is because the *NuSTAR* QPO FT was calculated using FPMB channels as the subject bands, and the full band FPMA light curve as the reference band. We see from Fig. 2.1 that the cross-calibration between *NICER* and the *NuSTAR* modules is good, except for the normalisation. We account for this by multiplying our model by a floating constant, which is fixed to unity for *NICER* and left free for the *NuSTAR* FPMA and FPMB.

As we used different reference bands for the calculation of the *NICER* and *NuSTAR* QPO FTs (the full band *NICER* and *NuSTAR* FPMB light curves

respectively), there is a phase difference between them caused by the phase lag between the two reference bands (see Section 2.3.2). We account for this by setting a phase offset ϕ_c for the first harmonic (the phase offset for the second harmonic is exactly $2\phi_c$ ¹⁰). Similarly to floating calibration constants employed for flux-energy spectra, we set $\phi_c = 0$ for *NICER* and leave ϕ_c as a free parameter for *NUSTAR*.

In our fits, we leave the truncation radius r_{in} as a free parameter and fix the spin to $a = 0.998$. While we do not necessarily expect that this choice reflects the true value of the spin (see e.g. Mills et al. 2021), this value enables the widest possible range of the r_{in} parameter to be explored without it becoming smaller than the ISCO. Although the spin does affect the geodesics, this only has a very subtle effect on the spectrum compared with the location of the disc inner radius.

The full list of the free parameters in our model is included in Table 2.2, with those modulated with QPO phase γ in the manner of Eq. 2.11 grouped together, and are labeled in the manner of ‘ $N_c(\gamma)$ ’. We begin by finding a global best fit through minimising the χ^2 fit statistic, before running a long Markov Chain Monte Carlo (MCMC) simulation around the best fit to explore the parameter space and understand the significance of the best fitting parameter values.

2.5.2 Results

Our global best fit, which is shown in Fig. 2.6, has $\chi^2 = 3374.6$ for 3179 degrees of freedom (DoF). The phase-average flux-energy spectrum in the upper-left of the figure shows the good fit, especially to the Fe $K\alpha$ line. The right-hand side of the figure shows the real and imaginary parts of the FT spectra of the first two harmonics. The features of these spectra are markedly less intuitive, and so we rely on analysis of the parameter space to understand the model fit. As the distant reflector is constant on the QPO timescale, the FT spectra do not include this component. However, as they are normalised to the observed flux (see Eq. 2.3), the absorption column does matter and therefore is included.

¹⁰This follows as the phase of each harmonic in every energy band is measured in reference to the phase of the first harmonic in the reference band. ϕ_c is the phase offset between the first harmonics in the reference band between the two instruments, and the same offset is $2\phi_c$ when instead measured at frequency of the second harmonic. See Appendix A.3 for the derivation.

Parameter	Unit	Chain mean	Description
N_{H}	10^{22}	6.8 ± 0.3	TBABS col. density
$A_{1\Gamma}$		0.04 ± 0.01	
$A_{2\Gamma}$		0.05 ± 0.02	
$\Gamma(\gamma)$	cyc	0.42 ± 0.04	Photon index
	cyc	0.13 ± 0.03	
Γ		1.93 ± 0.01	
A_{1e}	keV	$38. \pm 13.$	
A_{2e}	keV	$19._{-9}^{+8}.$	
$kT_e(\gamma)$	cyc	0.44 ± 0.04	Electron temp.
	cyc	0.97 ± 0.03	
kT_e	keV	$57._{-10}^{+9}.$	
$\log \xi^\dagger$		2.5 ± 0.2	Ionisation of dist. refl.
norm^\dagger	10^{-3}	0.7 ± 0.2	Normalisation of dist. refl.
A_{Fe}	$A_{\text{Fe}\odot}$	$7. \pm 1.$	Accreting Fe abundance
Incl.	deg	$75.1_{-0.3}^{+0.5}$	Inclination of source
r_{in}	r_{g}	$1.43_{-0.02}^{+0.01}$	Inner truncation radius
q_1		$13.6_{-0.7}^{+0.8}$	Inner emissivity index
$r_{\text{br},1}$	r_{in}	2.4 ± 0.5	1 st break radius
q_2		$4._{-4}^{+5}.$	Outer emissivity index
$r_{\text{br},2}$	$r_{\text{br},1}$	$14._{-10}^{+11}.$	2 nd break radius
A_1		0.5 ± 0.3	
A_2		1.3 ± 0.7	
Asym.	cyc	0.9 ± 0.1	Asymmetric illumination
	cyc	$0.99_{-0.05}^{+0.06}$	
A_{1f}		0.3 ± 0.1	
A_{2f}		0.3 ± 0.1	
$f_{\text{R}}(\gamma)$	cyc	0.45 ± 0.05	Reflection fraction
	cyc	0.94 ± 0.05	
f_{R}		$0.96_{-0.09}^{+0.08}$	
$\log \xi_{\text{max}}$		$4.26_{-0.09}^{+0.08}$	Max radial ionisation
$\Delta\gamma$	cyc	0.17 ± 0.07	Thermalisation phase lag
$kT_{\text{v,max}}$	keV	0.24 ± 0.05	Max radial viscous heating
kT_{i}	keV	0.64 ± 0.04	Heating from irradiation
N_{d}		$25. \pm 11.$	Disc normalisation
A_{1N}		1.2 ± 0.2	
A_{2N}		0.8 ± 0.3	
$N_{\text{c}}(\gamma)$	cyc	0.96 ± 0.03	Coronal Normalisation
	cyc	0.70 ± 0.03	
N_{c}		5.8 ± 0.2	
FMPA		0.944 ± 0.001	<i>NuSTAR</i> FMPA norm.
FMPB		0.951 ± 0.001	<i>NuSTAR</i> FMPB norm.
ϕ_{c}	cyc	0.018 ± 0.004	<i>NuSTAR</i> phase offset

Table 2.2: The mean and $\pm 1\sigma$ credible interval of the posterior distributions of each parameter, calculated with a MCMC. The two parameters marked with \dagger are those solely relating to the phase-constant reflected component assumed to come from a distant reflector. The ‘Asym.’ parameters are those that govern the asymmetric illumination profile. The FMPA and FMPB norm parameters are floating calibration constants.

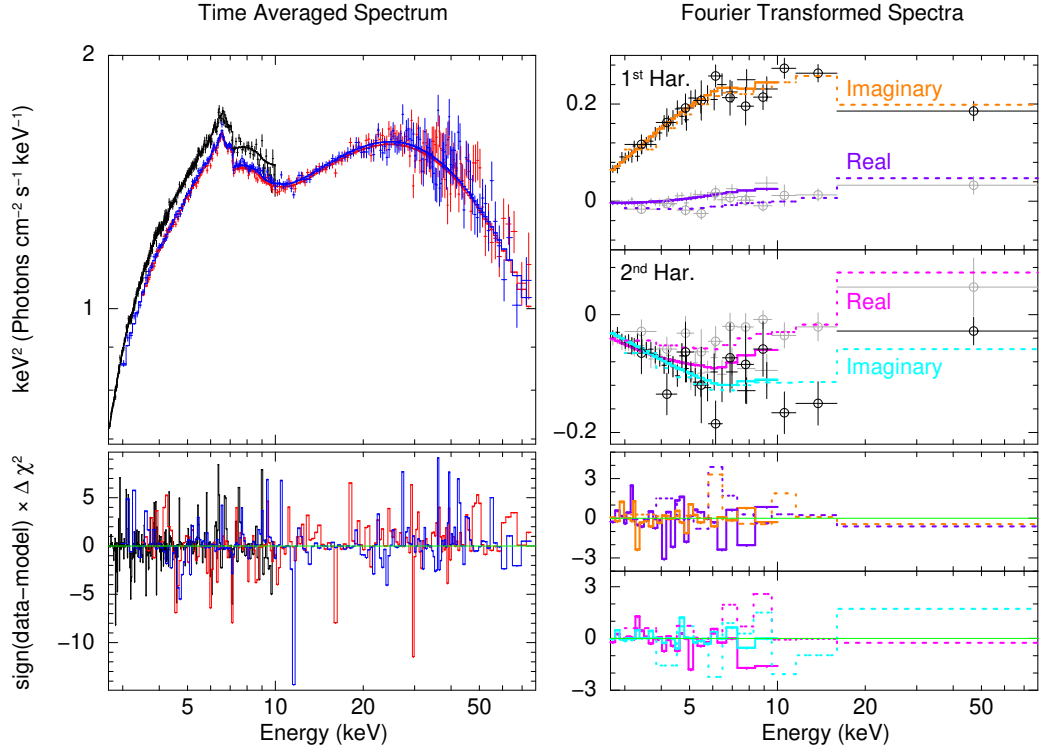


Figure 2.6: The unfolded best fit model, and the ratio of the data to the model, for the flux-energy spectra (left) and the harmonics of the QPO (right). The flux-energy spectra from *NICER* is in black, *NuSTAR* FPMA is in red, and *NuSTAR* FPMB is in blue, which have been rebinned for plotting purposes. The spectra of the real and imaginary components of the 1st QPO harmonic are in purple and orange respectively. Likewise, the 2nd QPO harmonic components are shown in magenta and light blue. Both harmonics have the *NuSTAR* data points with circles and a dotted line, with the circle-less points and solid lines corresponding to *NICER*. The markers in grey show the real component of the FT spectra, whereas those in black show the imaginary component. The axis labels apply to both the left and right hand sides of the figure, including the units; this follows from the Fourier transformed spectra being the RMS multiplied by a unitless phase term.

We explored parameter space by running a long MCMC simulation with 120 walkers which each run for 250,000 steps from an initial distribution based on the covariance matrix of the best fit. Every parameter had a uniform prior with bounds considerably far from the range required, except where parameters must be non-negative¹¹. We burn the first 50,000 steps of each walker, enough to ensure that the Geweke convergence diagnostic (Geweke 1992) is within ± 0.3 for every

¹¹The amplitude ‘*A*’ parameters of phase-modulated quantities are such examples, including the asymmetry parameters A_1 and A_2 .

parameter and we therefore only consider steps where the MCMC has converged. Finally, we thin the MCMC down, only taking every 100th step from each walker. We show the posterior means and 1σ credible intervals from the MCMC in Table 2.2.

In Fig. 2.7, we visualise our results by reconstructing parameter modulations from the thinned chain. For each step in the chain (there are 240,000 steps altogether), we use the parameter values corresponding to that step in order to calculate each of the 7 quantities plotted in Fig. 2.7 as a function of QPO phase. From these 240,000 functions of QPO phase, we create a 2D histogram, which we plot as a probability map (black represents the largest probability).

Panels 2-5 in Fig. 2.7 show the parameters that are allowed to vary with QPO phase via a sum of sinusoids (e.g. Eq. 2.11); from top to bottom: $f_R(\gamma)$, $\Gamma(\gamma)$, $kT_e(\gamma)$ and $N_c(\gamma)$. We consider the posterior distributions in Fig. A.2, and we see that all the parameter modulations are significant to at least 3σ , with $N_c(\gamma)$ likely at a much higher significance. The bottom two panels in the figure are $\log \xi_{\max}(\gamma)$ and $kT_{\text{irr}}(\gamma)$. The modulations in these parameters are not free to vary in our fit, they are instead calculated from $N_c(\gamma)$ and $f_R(\gamma)$ (see Eqs. 2.19, and 2.21). These are remarkably constant with QPO phase, as they are both $\propto N_c(\gamma)f_R(\gamma)$ (neglecting the phase shift $\Delta\gamma$), whose modulations are approximately out of phase.

The top panel is iron line centroid energy. In order to determine this, we first calculate the observed QPO phase-dependent reflection spectrum (Eq. A.6) assuming a δ -function iron line in the disc restframe (i.e. $I(E_d) = \delta(E_d - 6.4\text{keV})$) and then calculate the centroid energy of the resulting QPO phase-dependent line profile (using Equation 7 from Ingram et al. 2017). Any modulations of this function with QPO phase are caused entirely by QPO phase dependence of the emissivity function $\epsilon(r, \phi, \gamma)$, which in turn is driven exclusively by the asymmetry parameters A_1 and A_2 . We see that the line centroid energy is strongly modulated with QPO phase, implying that the emissivity function is required by the fit to vary with QPO phase. We show the posterior distribution of the asymmetry parameters from the MCMC in Fig. 2.8. The point $A_1 = A_2 = 0$ lies outside of the 2σ

contour (in red), therefore we are able to reject the axi-symmetric null-hypothesis of $A_1 = A_2 = 0$ at the 2σ significance level.

We show the QPO phase-dependent spectrum of the best-fitting model in Fig. 2.9, without the line-of-sight absorption and distant reflector. We see large changes in spectral shape over the course of the 8 QPO phases pictured, most obviously the change in continuum normalisation $N_c(\gamma)$ following the same trend as shown in Fig. 2.7. As the normalisation of the reflection spectrum is broadly constant with phase, its key features of the Fe $K\alpha$ and Compton hump are less pronounced when $N_c(\gamma)$ is larger. The temperature of the Compton hump itself does change, most noticeably at its lowest value of $\gamma \approx 1/4$ cycles.

When we consider the covariance between the modulated variables (see Fig. A.1), we see the strongest correlation is between the reflection fraction f_R and continuum normalisation N_c . Interestingly, the phase-averaged values are negatively correlated, however the size and phases of the modulations are positively correlated. We can see in Fig. 2.7 that these modulations are also in anti-phase, and so their modulations work against each other to keep the incident flux onto the disc approximately constant, as can be seen in the waveforms of $\log \xi_{\max}(\gamma)$ and $kT_{\text{irr}}(\gamma)$.

We find that the posterior mean of the phase lag $\Delta\gamma$, intended to represent the time it takes photons to thermalise in the disc atmosphere, is $\Delta\gamma \approx 0.17$ QPO cycles. Fig. 2.10 shows the posterior distribution of $\Delta\gamma$ from the MCMC, including a conversion from QPO cycles to time lag for a QPO frequency of $\nu_{\text{QPO}} = 2.2$ Hz. This phase lag can be seen in the bottom two panels of Fig. 2.7, since the dip in $kT_{\text{irr}}(\gamma)$ occurs ~ 0.17 QPO cycles after the dip in $\log \xi_{\max}(\gamma)$. It can also be seen in Fig. 2.9: e.g. the orange line for QPO phase = 0.125 cycles corresponds to a peak in iron line centroid energy but not to a peak in disc peak temperature. Converting $\Delta\gamma$ to a time lag gives ~ 75 ms, which is rather large. This large value is possibly due to model systematics, since if the thermalisation timescale really were this long, then we would see much longer thermal reverberation lags than have been observed (e.g. Uttley et al. 2011; Kara et al. 2019).




Figure 2.7: Each panel shows a parameter which is modulated with QPO phase in our best fit. The heatmap is constructed from the MCMC to visualise the confidence of the modulation. For each step in the MCMC, the modulation curve of each quantity was reconstructed, and is over-plotted to produce the heatmap. The top panel shows the iron-line centroid energy E_c which was calculated from the varying illumination profile, assuming a rest-frame δ -function profile (see text for details). The reflection fraction, Γ , kT_e , and N_c are the modulations straight from their parameters in the model, whereas $\log \xi_{\text{max}}$ and kT_{irr} are calculated from the irradiating flux, as described in the text.




Figure 2.8: A_1 and A_2 contour based around the MCMC. The 1, 2, and 3 σ credible intervals are shown as contours (purple, orange, red), with the blue lines highlighting the values at the best-fit. Within the 1σ contour the density is shown as a gray-scale 2D-histogram. Outside the 3σ contour individual points are shown as grey points. The marginalised histograms also show the $\pm 1\sigma$ credible interval with purple dashed lines.




Figure 2.9: The QPO phase dependent component of our best-fitting model; without line-of-sight absorption or the distant reflector component. This is comprised of the NTHCOMP cut-off powerlaw; a ray-traced XILLVERCP reprocessed spectrum with relativistically smeared atomic features ($5 - 8$ keV) and Compton hump ($\gtrsim 20$ keV); plus a blackbody thermal spectrum at $\lesssim 2$ keV. This shows the model at 8 phases, although 16 phases are used for the calculation of the Fourier transformed components.




Figure 2.10: Histogram of $\Delta\gamma$ from the MCMC, with the corresponding time-lag for a frequency of 2.2 Hz. The $\pm 1\sigma$ credible interval is outlined with purple dashed lines.

2.6 Discussion

We have conducted a phase-resolved spectral analysis of a ~ 2.2 Hz QPO from GRS 1915+105 observed simultaneously by *NICER* and *NuSTAR*. We found that the continuum normalisation $N_c(\gamma)$, photon index $\Gamma(\gamma)$, reflection fraction $f_R(\gamma)$, and electron temperature $kT_e(\gamma)$ are all required to be modulated with QPO phase to $> 3\sigma$ confidence, plus we rule out that the asymmetric illumination parameters $A_1 = A_2 = 0$ at the 2σ confidence level. Alongside this, we found a small inner truncation radius r_{in} , and a large thermalisation phase lag $\Delta\gamma$. We now discuss the implications of these results.

2.6.1 Comparison with H1743-322

We compare our results with the previous similar study of a ~ 0.25 Hz type-C QPO in H 1713-322 by Ingram et al. (2017) (hereafter I17). The first thing to note is the strikingly similar profiles of the continuum normalisation, $N_c(\gamma)$ (compare our Fig. 2.7 with the Fig. 5 in I17), which is a proxy for the X-ray flux. Whereas the fractional amplitude of the $N_c(\gamma)$ modulation is larger here, the profiles – and therefore the QPO waveforms – are similar for the two observations, despite the

QPO frequencies being very different. This similarity is confirmed by the measured phase difference between harmonics, ψ , which defines the QPO waveform (Ingram & van der Klis 2015). Here we measure $\psi/\pi \approx 0.2$ for the *NuSTAR* observation, whereas the *NuSTAR* measurement for H1743-332 was $\psi/\pi \approx 0.3$ (Ingram et al. 2016). This indicates that the $N_c(\gamma)$ modulation for H1743-322 should be similar to what we measure here for GRS 1915+105 but with the peak slightly delayed, which can indeed be seen in the $N_c(\gamma)$ profiles. It is perhaps surprising that the QPO waveform of the two observations is so similar given that Ingram & van der Klis (2015) observed it to change dramatically between two *RXTE* observations of GRS 1915+105 with QPO frequencies of ~ 0.5 Hz and ~ 2.25 Hz. Fig. 5 (right panel) in de Ruiter et al. (2019) reveals that this is because ψ reduces more steeply with QPO frequency in GRS 1915+105 than it does in H1743-322, such that its value for GRS 1915+105 at $\nu_{\text{qpo}} \sim 2$ Hz is close to the corresponding value for H1743-322 at $\nu_{\text{qpo}} \sim 0.2$ Hz. The measurement of ψ presented here is therefore consistent with the trends found by de Ruiter et al. (2019). Beyond the overall shape, we also see a much wider spread in our $N_c(\gamma)$ histogram than that presented by I17. This is because we are now using a much more flexible model than in previous work, including modulations in kT_e , and so it is better able to compensate when $N_c(\gamma)$ deviates from its best fitting functional form.

As for the modulation in iron line centroid energy, $E_c(\gamma)$, we see a strong second harmonic (evidenced by there being two maxima per QPO phase as opposed to one) both in our observation and in the I17 analysis of H1743-322. This property was also previously observed for GRS 1915+105 by (Ingram & van der Klis 2015), albeit with a low statistical significance. We however see that the phase of the $E_c(\gamma)$ modulation is shifted here with respect to the H1743-322 observation: here the maxima in E_c occur at QPO phase $\gamma \sim 0$ and 0.5 cycles, whereas for H1743-322 they occur at $\gamma \sim 0.2$ and ~ 0.7 cycles. A similar evolution in the $E_c(\gamma)$ waveform is seen between the two GRS 1915+105 observations presented in Ingram & van der Klis (2015): their $E_c(\gamma)$ waveforms for $\nu_{\text{qpo}} \sim 0.5$ Hz and $\nu_{\text{qpo}} \sim 2.25$ Hz QPOs are respectively similar to the ~ 0.2 Hz QPO in H1743-322 and the ~ 2.2 Hz QPO

in GRS 1915+105 presented here. This gives a potential hint that the phase of the line centroid energy modulation evolves systematically with QPO frequency, as is the case for the flux modulation (de Ruiter et al. 2019). We additionally see that the $E_c(\gamma)$ modulation presented here has a lower mean and a higher amplitude than that presented by I17. This is consistent with the disc inner radius reducing as the QPO frequency increases: the mean is reduced by increased gravitational redshift and the amplitude is increased by faster orbital motion closer to the BH.

I17 found modulations of the reflection fraction, f_R , and power law index, Γ , to have 3.52σ and 0.95σ significance respectively, whereas here we find both modulations to have $> 3\sigma$ significance. As the reflected spectrum is spectrally harder than the directly observed spectrum, an increase in spectral hardness can be caused either by a reduction in Γ or an increase in f_R . Whereas the $\Gamma(\gamma)$ and $f_R(\gamma)$ modulations presented by I17 are broadly in phase, suggesting they somewhat compensate for each other. Here we find they are broadly in anti-phase suggesting they are both contributing to the modulation of the spectral hardness.

2.6.2 Asymmetric illumination profile

Our model requires an asymmetric illumination profile with $A_1 = 0.5 \pm 0.3$ and $A_2 = 1.3 \pm 0.7$ which is consistent with two, nonidentical bright patches rotating around the disc with QPO phase. The QPO phase dependence of the iron line profile could, to some extent, be reproduced by changes in the ionisation state of the disc atmosphere, which would cause changes in the shape and centroid energy of the iron line in the restframe reflection spectrum (since e.g. Compton broadening of the line increases with the number of free electrons and more ionised species are more tightly bound and therefore produce higher energy fluorescence lines). The parameters that affect the shape of the restframe reflection spectrum and are modulated with QPO phase in our model are Γ , kT_e , and $\log \xi$. All three affect the ionisation state of the disc, with larger kT_e , larger $\log \xi$, and smaller Γ increasing the number of irradiating photons with a high enough energy to ionize neutral iron ($E > 7.1$ keV). In the model employed by I17, only Γ was allowed to vary with QPO phase, whereas

here we also allow kT_e and $\log \xi$ to vary, with $\log \xi$ tied to the illuminating flux. We also note that here we employ a radial $\log \xi$ profile, whereas only a single ionisation parameter was used for the entire disc in I17. Given the flexibility of our model and the conservative nature of our analysis, we find that the asymmetric illumination parameters A_1 and A_2 are only required to be non-zero at $> 2\sigma$ confidence level.

Such an asymmetric, QPO phase-dependent illumination pattern is naturally expected if the corona illuminating the disc is precessing (Ingram & Done 2012). Alternatively, precession of the disc itself (i.e. precession of the reflector and not the illuminator) could potentially explain our data (Schnittman et al. 2006). However, disc precession is not expected theoretically (Bardeen & Petterson 1975; Liska et al. 2019) unless the frame dragging effect is strong enough to tear the disc into a number of discrete, independently precessing rings (Nixon & King 2012; Liska et al. 2021). Such a configuration could in principle cause QPOs if the number of independent rings is small enough to produce a coherent oscillation. However, disc tearing requires a large misalignment between the binary and BH rotation axes, which is expected to be rare since the biggest natal kicks, which would produce the largest misalignments, are also the most likely to completely disrupt the binary system (Fragos et al. 2010). This is in contrast to the near-ubiquity of type-C QPOs. Line profile modulations could also result from c-mode disco-seismic waves (e.g. Kato & Fukue 1980; Kato 2001). However, this would not explain the QPO being much stronger in the Comptonised spectrum than in the disc spectrum, and the line profile variations produced by c-modes are quite subtle compared with what we observe (see e.g. Fig 5 in Tsang & Butsky 2013), due to the c-mode oscillation only occurring in a narrow range of disc radii. The spiral waves of the AEI model are also consistent with the asymmetric illumination profile we measure here (Varnière & Tagger 2002). However, the AEI model is not consistent with a reflection fraction modulation, which we measure with $> 3\sigma$ significance, whereas the precessing corona model is.

A new QPO diagnostic will soon be available in the form of X-ray polarization. Whereas the precession model predicts the polarization degree and angle to be modulated on the QPO period (Ingram et al. 2015), alternative models such as the

AEI do not. After its launch later this year, the *Imaging X-ray Polarimetry Explorer* (*IXPE*) should be capable of detecting QPOs in the X-ray polarization properties in a long exposure observation of a bright X-ray binary, as long as high inclination X-ray binaries turn out to be moderately polarized ($\gtrsim 4\%$; Ingram & Maccarone 2017).

2.6.3 Inner Radius

Our best-fitting model also requires a small disc inner radius of $r_{\text{in}} \simeq 1.43_{-0.02}^{+0.01} R_{\text{g}}$, where $R_{\text{g}} = GM/c^2$ and the uncertainty (from the MCMC posterior distribution*) only reflects the statistical error around the best fit within our model. While this small value of r_{in} is compatible with previous reflection spectroscopy results (Miller et al. 2013; Zhang et al. 2019; Shreeram & Ingram 2020), it does not leave much room for a corona to precess within this truncation radius in order to give rise to the best fitting asymmetric illumination profile. Moreover, it is very much incompatible with the precession being specifically at the Lense-Thirring precession frequency, which requires $r_{\text{in}} \approx 41 R_{\text{g}}$ to reproduce a QPO frequency of $\approx 2.2 \text{ Hz}^{12}$.

We therefore first consider if our analysis has returned an artificially small disc inner radius due to the use of an inadequate continuum model. Including a second, softer power law component can yield a larger truncation radius in fits to the flux-energy spectrum, since the new component takes the place of the broad red wing characteristic of small r_{in} (Yamada et al. 2013; Mahmoud & Done 2018; Mahmoud et al. 2019; Zdziarski et al. 2021b; Zdziarski et al. 2021a). Such a treatment could therefore yield a larger truncation radius in our analysis if our measured disc inner radius is driven primarily by the time-averaged iron line profile. However, our analysis is also sensitive to the QPO phase-dependence of the iron line profile. For a large truncation radius, the red wing can dominate over the blue wing when the receding disc material is illuminated, whereas for a small truncation radius the variability is entirely in the blue wing because the red wing is always suppressed by gravitational redshift (Ingram & Done 2012; You et al. 2020).

¹²Using Eq. 2 of Ingram et al. (2009) with $\zeta = 0$, as seen in numerical simulations (Fragile et al. 2007), $r_{\text{in}} = r_{\text{isco}}(a)$, $a = 0.998$, and $M = 10.1 M_{\odot}$ (Steehns et al. 2013).

By excluding the iron line from the flux-energy spectrum, but leaving in those bands for the Fourier transformed spectra, we are able to investigate whether this small r_{in} is required by only the phase-average spectrum or also by the QPO phase dependence of the spectrum. We therefore run a new phase-resolved fit in which we ignore the 5 – 8 keV data from the flux-energy spectrum but leave in those bands for the Fourier transformed spectra. For this new fit, the measured r_{in} value will not be driven by the time-averaged shape of the iron line, but by how its shape *changes* throughout the course of each QPO cycle. This new fit also returns a small disc inner radius of $r_{\text{in}} = 1.49_{-0.04}^{+0.03} R_{\text{g}}$. It is still possible that including a more complex continuum could allow for a slightly larger r_{in} , but this would add many more degrees of freedom to our already extremely flexible model.

Our fits therefore indicate (although not definitively) asymmetric illumination of a disc extending to a very small inner radius. It could be that the corona is not radially extended but instead a vertically extended structure, such as a precessing jet (Kalamkar et al. 2015; Stevens & Uttley 2016; Kylafis et al. 2020) (if the jet base is sufficiently X-ray bright; Fender et al. 1999; Markoff et al. 2005). GRMHD simulations show that jet precession can be induced by the frame dragging effect (Liska et al. 2018), but only so far in the presence of a thick disc. Indeed, our best fitting model includes a very strongly peaked emissivity function, with $q_1 = 13.6_{-0.7}^{+0.8}$ for disc radii $r < 3.3 \pm 0.6 R_{\text{g}}$. This emissivity profile is roughly compatible with one created by a vertically extended corona raised slightly above the thin accretion disc (Wilkins & Fabian 2012).

It is, however, unclear exactly how the frame dragging effect could drive such slow precession for such a small disc inner radius. Perhaps the corona is vertically extended with the density increasing with distance from the BH. This weighting of the density to larger distances from the BH will slow down precession compared with a precessing ring at the same r_{in} . The torque exerted on the corona by the outer disc, which is not currently considered in calculations of the precession frequency, will also slow down precession. It is not, however, clear whether or not these two effects are sufficient to solve the discrepancy we find here. Even if they are, and

Lense-Thirring precession really is the true type-C QPO mechanism, their presence will make it much more difficult to infer BH mass and spin from the QPO frequency than previously hoped. A clear counter-argument to this point is provided by the QPO triplet in GRO J1655-40, the frequencies of which can be explained very well by the relativistic precession model, returning a precise mass prediction that agrees with the dynamical value (Motta et al. 2014b; Fragile et al. 2016). This result instead argues that our very small disc inner radius is instead the result of modelling systematics such as employing an overly simplified continuum, as discussed above.

It is important to note that all well-studied QPO models in the literature assume that the QPO frequency increases during the state transition primarily due to the disc inner radius moving inwards (Ingram & Motta 2019). There is therefore no published QPO model that can reproduce our observed QPO frequency for our measured disc inner radius without significant modification.

2.6.4 Misalignment

Our best fitting disc inclination angle is $i = 75.1_{-0.3}^{+0.5}$ degrees, whereas the jet inclination angle, inferred from radio observations of superluminal ejections (Fender et al. 1999), is $\theta \approx 60^\circ$ (using the radio parallax distance of $D \approx 8.6$ kpc; Reid et al. 2014). We therefore infer a misaligned system, consistent with QPO models that invoke Lense-Thirring precession (Stella & Vietri 1998; Ingram et al. 2009). Following I17, we can estimate the misalignment angle β between the disc and BH spin axes by taking the large-scale jet as a proxy for the BH spin axis¹³. This is not necessarily simply given by $\beta = i - \theta$, since the azimuthal disc viewing angle, Φ , is unknown. Rather, it can be found by solving the geometry (Veledina et al. 2013; Ingram et al. 2017)

$$\cos \theta = \sin i \sin \beta \cos \Phi + \cos i \cos \beta, \quad (2.22)$$

which gives the possible values for the misalignment angle β , given the observed inclinations of the disc i and the BH spin axis θ (assuming the jets are ejected

¹³This of course might not be correct, as jets have been observed to precess as in e.g, Miller-Jones et al. (2019).

along this axis), and any possible but unknown azimuthal disc viewing angle Φ . Fig. 2.11 shows all of the solutions for β for the full range of Φ and θ values (colour scale), and assuming $i = 75^\circ$. The red line represents the solutions corresponding to $\theta = 60^\circ$, which cover the range $15.1^\circ < \beta < 135.1^\circ$ for $|\Phi| < 63.7^\circ$ (i.e. for some values of Φ there is no solution). The minimum misalignment compatible with our results is therefore $\beta \approx 15^\circ$, which is a large enough misalignment to produce the observed $\sim 15\%$ RMS QPO amplitude with a corona precessing around the BH spin axis (Ingram et al. 2017).

This misalignment will greatly influence the BH spin inferred from disc continuum fitting in the soft state. In the most recent such analysis of GRS 1915+105, Mills et al. (2021) report a best fit of $a \approx 0.86$ ($r_{\text{isco}}(a) = r_{\text{in}} \approx 2.57 R_g$), with the very high spin required by reflection spectroscopy measurements (including our own) also within uncertainties. However, they assume a completely aligned system, with the binary inclination δ equal to the disc inclination i equal to the jet inclination θ . Ignoring relativistic corrections, the disc inner radius inferred from disc fitting is $r_{\text{in}} \propto D/(M\sqrt{\cos i}) R_g$, and so adopting $i = 75^\circ$ in place of $i = \theta = 60^\circ$ (but still assuming $\delta = \theta$ and $D = 8.6$ kpc) would instead give $r_{\text{in}} \approx 3.63 R_g$. The inferred inner radius is pushed even further out if we set $\delta = i$ as is assumed in the precession model (Ingram et al. 2009; Ingram et al. 2015), since the BH mass is related to the binary mass function $f(M)$ as $M \propto f(M)/\sin^3 \delta$. For $\delta = i = 75^\circ$ and still assuming $D = 8.6$ kpc, the inner radius becomes $r_{\text{in}} \approx 5 R_g$, which is incompatible with our measured value of $r_{\text{in}} \approx 1.4 R_g$. If we instead assume that δ is unknown, we find that for $i = 75^\circ$ and $D = 8.6$ kpc the disc continuum fitting inner radius is equal to our value if $\delta \approx 39^\circ$, implying a BH mass of $M \approx 32 M_\odot$ which is much greater than the dynamical measurement of $10.1 \pm 0.6 M_\odot$ (Steehns et al. 2013).

We note, however, that a thin misaligned disc is expected to form a so-called Bardeen-Petterson configuration (Bardeen & Petterson 1975), whereby the outer and inner regions align respectively with the binary and BH spin axes. GRMHD simulations show that the transition from misaligned to aligned can be close to the BH ($r \sim 6 R_g$; Liska et al. 2019), but still further out than our very small inner




Figure 2.11: The BH spin axes inclination θ , for different values of misalignment between it and the disc, for the different values of disc azimuthal angle Φ of the misalignment. The red line highlights $\theta = 60^\circ$, the observed inclination large-scale jet, which could be assumed to be aligned with the BH spin-axis.

radius of $r_{\text{in}} \sim 1.4 R_g$. In contrast, our simplified model only considers a planar disc. It could be that failing to account for a more realistic warped disc geometry has introduced a bias into our measurement of i , or even of r_{in} .

2.6.5 Biases in the phase-averaged spectrum

The shape of our best fitting phase-resolved spectrum changes dramatically with QPO phase (Fig. 2.9). These non-linear variations over each QPO cycle may cause biases in analyses that only fit a steady-state model to the time-averaged flux-energy spectrum. Following I17, we investigate these potential biases in Fig. 2.12 by plotting the phase-averaged spectrum of our best fitting phase-resolved model (black) alongside the spectrum calculated by setting all modulated parameters to their phase-averaged values (red), as well as the percentage difference with respect to the phase-averaged model (bottom panel). We see that the difference between the two spectra is only $\sim 1\%$ in the $\sim 5 - 20$ keV region, but is much larger at lower and higher energies, indicating that ignoring spectral variability does indeed cause a bias. We investigate further by fitting the observed flux-energy spectrum with a steady-state model (blue lines), which reproduces the phase-averaged model very well (except for a narrow feature at ~ 6.7 keV, which *NICER* and *NuSTAR* cannot resolve, but future missions such as *ATHENA* will be able to). The parameters

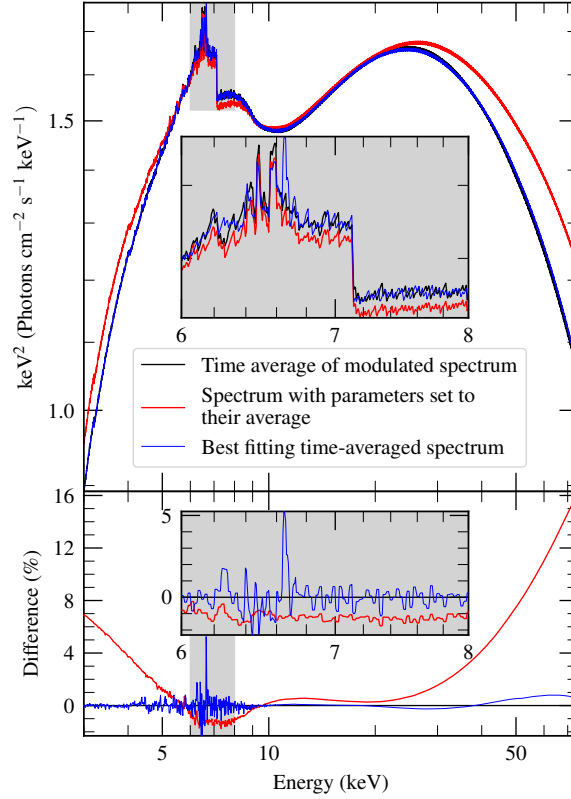


Figure 2.12: Top: The best fitting flux-energy spectrum (black); the same however with all the modulations turned off so each parameter is set to its phase-average value (red); the best-fitting model when there is no phase-dependant parameter modulations or changing illumination profile (blue). Bottom: The percentage difference between the two. The grey inset panels show a 6-8 keV zoomed-in version of their respective lines, with the corresponding areas of the main graph also highlighted in grey.

of the new fit are consistent with those of the phase-resolved model except kT_{irr} , kT_{visc} and kT_e are all cooler and the disc normalisation is larger. We therefore conclude that these parameters are biased by spectral variability, but not other key parameters such as i and r_{in} .

2.6.6 Thermalisation lags

In our model we include a phase lag, $\Delta\gamma$, between variations in the illuminating flux and corresponding variations in disc heating. This is to allow for the finite time it takes for incident radiation to thermalise in the disc atmosphere. This timescale is poorly understood theoretically, but is surely longer than the corresponding timescales for other processes such as fluorescence and scattering, which are almost

instantaneous (García et al. 2013b). Leaving $\Delta\gamma$ as a free parameter in principle enables us to empirically measure this thermalisation timescale. Our best fitting value of $\Delta\gamma = 0.15_{-0.05}^{+0.06}$ cycles corresponds to a time lag of $\sim 70_{-25}^{+26}$ ms given that the QPO frequency is 2.2 Hz, which is much larger than expected. This very long thermalisation time delay is incompatible with observations of $\lesssim 1$ ms soft lags in e.g. MAXI J 1820+070 (Kara et al. 2019) and GX 339-4 (Uttley et al. 2011), and so it is very likely that the $\Delta\gamma$ parameter is accounting for some other over-simplification in our model. This is perhaps not surprising, since our prescription is so simple. For instance, we assume a single thermalisation delay, whereas in reality it is likely dependent on disc radius (Frank et al. 2002) and on ionisation parameter (which varies with both disc radius and azimuth in our model). Moreover, in reality there will be some smearing such that e.g. very fast fluctuations in the irradiating flux will not be efficiently transferred into fluctuations in disc temperature, which we currently completely ignore.

In the context of the precessing corona model, there are a number of effects that we do not account for here which would lead to modulations of the observed disc temperature. For instance shadowing: the precessing corona and/or jet obscuring different disc azimuths at different precession phases leading to a variation in the shape of the overall observed disc spectrum. We also note that in the precessing corona model, the misalignment between the disc and BH spin axes, β , remains constant but as the corona precesses around the BH spin axis its misalignment with the disc axis varies between 0 and 2β (e.g. Ingram et al. 2015). This would lead to the illumination profile being axisymmetric once per QPO cycle, which is not accounted for by the illumination profile employed here. There are also potential sources of systematic error in our spectral model. For instance in the XILLVERCP grids that we use, the seed photon temperature is hardwired to $kT_{\text{bb}} = 0.05$ keV, whereas the seed photon temperature of the NTHCOMP component in our model varies with QPO phase. Finally, we ignore light crossing delays in our parameterisation of the disc illumination profile, which can be the order of milliseconds for reflection from

$r \sim 10s$ of R_g . This is $\lesssim 1\%$ of the $\sim 0.45s$ QPO period, but will become increasingly important for the study of QPOs with increasingly higher centroid frequency.

2.7 Conclusions

We have conducted a phase-resolved spectral analysis of a type-C QPO in a simultaneous *NICER* and *NuSTAR* observation of GRS 1915+105. We used a QPO phase-resolving technique, following Ingram et al. (2016) but significantly improved on their work, by including a QPO frequency tracking algorithm that makes it possible to analyse long observations over which the QPO frequency may change, and using the bispectrum to constrain the phase difference between QPO harmonics which has the advantage of enabling a proper Poisson noise correction. We have also used a significantly more advanced version of the Ingram et al. (2017) tomographic model with which to fit the phase-resolved spectra in the Fourier domain. Our new model allows more parameters to be simultaneously modulated and includes self-consistent modulations to the ionisation parameter and disc heating due to irradiation.

Our fit requires the asymmetric illumination parameters A_1 and A_2 to be > 0 with $> 2\sigma$ confidence. Similar to the results of I17 for H1743-322, this is consistent with that expected for precession of the illuminator, but our measurement has only moderate statistical significance. We also detect a $> 3\sigma$ significant modulation of the reflection fraction, indicating that the geometry of the inner accretion disc changes systematically with QPO phase. We inferred a high disc inclination ($i \approx 75^\circ$), which implies that the disc is misaligned with the previously observed jet ejections ($\theta \approx 60^\circ$). This is consistent with the precessing corona model for the QPO (Ingram et al. 2009). However, our fit also favoured a small disc inner radius, which requires the corona to be vertically rather than radially extended and is inconsistent with the precession frequency being set purely by the frame dragging effect. We discussed some possible effects that could slow down Lense-Thirring precession compared with the simplest prediction. We note that no QPO model currently in the published literature can reproduce the observed QPO frequency

for our measured disc inner radius (Ingram & Motta 2019). Therefore, either our measured radius is affected by modelling systematics – such as the continuum in reality being more complex than we assume – or new or modified theories must be developed to model the QPO frequencies.

We also recovered a large thermalisation delay, which implies that irradiating photons take ~ 70 ms to thermalise in the disc atmosphere. We argued that this is infeasibly large and discussed potential sources of systematic error that could be contributing to the measurement.

We compared this analysis and the work of Ingram et al. (2016, 2017) on H1743-332, and in particular note they both show a strong second harmonic in the modulation of the iron line centroid energy $E_c(\gamma)$. We found hints that the $E_c(\gamma)$ modulation evolves systematically with QPO frequency, as would be expected for example if the disc inner radius reduces as the QPO frequency increases. Further work is required to test this hypothesis.

"Space," says the introduction to *The Hitchhiker's Guide*, "is big. Really big. You just won't believe how vastly, hugely, mind-bogglingly big it is," and so on.

— Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

3

Proof-of-Principle mass measurement of the black hole in X-ray binary H1743-322

Contents

3.1	Introduction	96
3.2	Observations and Data Reduction	97
3.2.1	Observation Log	97
3.2.2	<i>NuSTAR</i> Data Reduction	97
3.2.3	<i>XMM-Newton</i> Data Reduction	97
3.2.4	Cross-Calibration	99
3.2.5	Timing Analysis	101
3.3	Spectral Fits	103
3.3.1	The model	103
3.3.2	Model fitting procedure	105
3.3.3	Results	107
3.4	Discussion	111
3.4.1	Proof-of-Principle	112
3.4.2	Mass Measurement	114
3.4.3	Higher Spin	114
3.4.4	Inner radius and QPOs	115
3.4.5	Density and inclination	116
3.5	Conclusions	117

3.1 Introduction

Black holes (BHs) in X-ray binaries (XRBs) typically have their mass measured when the system is in quiescence. In this state, the luminosity of the accretion system is very low, and the BH mass can be measured dynamically by tracking the orbital phase-dependent Doppler shifts of absorption lines in the optical spectrum of the companion star (e.g. Casares & Jonker 2014), as has been done for ~ 20 of the ~ 70 BHs in XRBs (Corral-Santana et al. 2016). However, these optical techniques can be inhibited by extinction, which is a significant effect for sources within the Galactic plane and bulge. One example of a system without a dynamical mass measurement is the BH XRB H1743-322. Steiner et al. (2012) was able to measure the source distance $D = 8.5 \pm 0.8$ kpc by measuring the observed time delay between approaching and receding super-luminal jet ejecta using radio observations. Using multiple flares to track the evolution and motion of the jet, they were able to associate a physical distance to the projected separation observed, and therefore calculate a distance to the source. This distance is consistent with the likely scenario of the source lying in the Galactic bulge, and indeed it experiences too much extinction for the optical measurement.

Jonker et al. (2021) hypothesised that this inability to measure the masses of systems lying in the Galactic disc and bulge could cause a bias to the measured BH mass population: assuming XRBs are likely to form in the Galactic plane or bulge, lower mass systems are more likely to receive a large enough ‘kick’ during the BH formation to push the system away from these regions with higher extinction (Fryer & Kalogera 2001), therefore suggesting that we are likely more able to measure the BH mass in lower mass systems. They suggest that this could be a reason why the population of BHs detected through gravitational waves (during mergers) is more massive than those detected in XRBs (Farr et al. 2011; Corral-Santana et al. 2016; Abbott et al. 2021).

Ingram et al. (2022, hereafter I22) proposed a model that is able to use the X-ray reflection spectrum to constrain the ratio of BH mass M to its distance from us D . In this chapter I present my work considering a set of observations of the BH

XRB H1743-322, which are described in Section 3.2. In Section 3.3 I summarise the details of how the I22 model works, and how I fit the model to the data. I discuss the mass measurement I obtain from my fits and the existing distance measurement, along with the relevance of this proof-of-principle demonstration of the technique in Section 3.4.

3.2 Observations and Data Reduction

3.2.1 Observation Log

We use observations of three outbursts of the BH XRB H1743-322, in 2014, 2016, and 2018. The observations are a combination of five with *the Nuclear Spectroscopic Telescope ARray* (*NuSTAR*; Harrison et al. 2013) and five with *X-ray Multi-Mirror Mission* (*XMM-Newton*; Jansen et al. 2001). We consider these as five epochs, which we label 14, 16a, 16b, 18a, and 18b, based upon the *NuSTAR* observation. The details of the ten different observations used are listed in Table 3.1. The data used in epoch 14 was included in Ingram et al. (2016, 2017) (the *NuSTAR* data and the "XMM2a" and "XMM2b" data within those papers).

3.2.2 *NuSTAR* Data Reduction

We used the *NuSTAR* data reduction software *NUSTARDS* v2.0.0, with *CALDB* v20200726. We first used the *FTOOL* *NUPIPELINE* to produce cleaned event lists, before extracting source and background products with the *FTOOL* *NUPRODUCTS*. For each observation, we used 60" circular source extraction region, and a 90" circular background extraction region, chosen from an area of the detector without source counts. For each observation the source regions were chosen per detector, but with a mutual background region. Finally, we used the *GRPPHA* to bin the energy spectra so that they contain a minimum of 25 counts per channel.

3.2.3 *XMM-Newton* Data Reduction

The *XMM-Newton* observations all used the RGS (Reflection Grating Spectrometers) and EPIC-pn (European Photon Imaging Camera) in timing mode, however with

Table 3.1: Details of the observations used, with the start and end time formats given as "YY-MM-DD hh:mm" in UTC. The observations are split into five epochs of simultaneous observations, each with one *NuSTAR* and between zero to two *XMM-Newton* observations, as listed.

Epoch	ObsID	<i>NuSTAR</i>		ObsID	<i>XMM-Newton</i>	
		Start	End		Start	End
14	80001044004	14-09-23 18:21	14-09-25 08:51	0724401901	14-09-23 20:08	14-09-24 18:00
				0740980201	14-09-24 18:54	14-09-25 08:35
16a	80202012002	16-03-13 01:36	16-03-14 19:11	0783540201	16-03-13 04:18	16-03-14 19:00
16b	80202012004	16-03-15 06:46	16-03-17 00:21	0783540301	16-03-15 05:00	16-03-16 18:52
18a	90401335002	18-09-19 02:36	18-09-19 01:01		—	
18b	80202012006	18-09-26 08:36	18-09-27 21:31	0783540401	18-09-26 14:39	18-09-28 03:04

Table 3.2: The difference in the power law index needed to cross-calibrate simultaneous *NuSTAR* and *XMM-Newton* observations, when a simple absorbed power law model is fit to the overlapping 3 – 10 keV band. The representative parameter uncertainty is 3×10^{-3} .

Epoch	<i>NuSTAR</i>	<i>XMM-Newton</i>	$\Delta\Gamma$
14	80001044004	0724401901	0.106
		0740980201	0.111
16a	80202012002	0783540201	0.088
16b	80202012004	0783540301	0.094
18b	80202012006	0783540401	0.085

no EPIC-mos data recorded. We used the *XMM-Newton* Science Analysis Software (SAS) version 20.0.0 for this data reduction.

3.2.3.1 EPIC-PN Timing Mode

We used the SAS task EPPROC to extract the cleaned events list, before extracting source products from a region defined with $31 \leq \text{RAWX} \leq 45$, using only single and double events ($\text{PATTERN} \leq 4$), and ignoring bad pixels ($\text{FLAG} = 0$). We rebinned the spectra to have a minimum of 25 counts per channel using SPECGROUP. We generated the response using RMFGEN using the new calibration published in CCF XRT3_XAREAEF_0014.CCF available with the setting APPLYABSFLUXCORR=YES, and we generated the ancillary file with ARFGEN.

3.2.3.2 RGS

We used the SAS task RGSPROC to produce the source and background spectra, along with the response, for the first spectral order of each RGS1 and RGS2. The RGS spectra were also rebinned using SPECGROUP to a minimum of 25 counts for channel.

3.2.4 Cross-Calibration

The observations we consider come from five separate instruments, and from each we make use of a different energy range:

- *NuSTAR*'s FPMA and FPMB (3-78 keV),
- *XMM-Newton*'s EPIC-pn in timing mode (2.5-10 keV),




Figure 3.1: The 3 – 20 keV flux-energy spectra from *NuSTAR*'s FMPA (black) and FMPB (red) detectors, along with *XMM-Newton*'s EPIC-pn (blue) which was in timing mode, from the 16a epoch. These were simultaneously fit to an absorbed power-law, each with a free normalisation, and with the power law index tied between the FMPA and FMPB, however with the EPIC-pn index otherwise free. The top panel shows the data and model unfolded about the response matrix, while the lower panel shows the ratio of the data to the model.

- *XMM-Newton*'s RGS1 and RGS2 (0.9-2.1 keV).

As the energy-range of *XMM-Newton*'s EPIC-pn camera and *NuSTAR*'s FPMs overlap, it is important that we check the cross-calibration between them. To do this, we fit the spectra from the three instruments with a simple absorbed power-law model, considering only the overlapping energy range (3 – 10 keV). For line-of-sight absorption, we use TBABS with the abundances of Wilms et al. (2000), and we tie the hydrogen column density N_H to be the same for all three spectra. The power law normalisation is left as a free parameter for each instrument to account for known discrepancies in absolute flux calibration. We tie the power law index between the

NuSTAR FPMA and FPMB, however allow it to be different for the EPIC-pn.

The top panel of Fig. 3.1 shows the three spectra from the 16a epoch unfolded around the best-fitting absorbed power-law model, and the bottom panel shows the data to (folded) model ratio. The EPIC-pn spectrum is slightly harder than the simultaneous *NuSTAR* spectrum. This is seen across the observations, with the difference between the fitting power-law indices listed in Table 3.2. We also see a discrepancy in the residuals at ~ 6 keV corresponding to the Fe $K\alpha$ line, with the line feature peaking at higher energies in EPIC-pn consistently across all of the epochs. This discrepancy is most likely due to calibration issues with EPIC-pn timing mode. The EPIC-pn timing mode has previously been found to measure harder spectra than other observatories in addition to *NuSTAR*. For example, an observation of GX 339-4 with a simultaneous *RXTE* exposure (Kolehmainen et al. 2014), and an observation of the Rapid Burster in which *Swift* XRT and *NuSTAR* were in agreement, but both disagreed strongly with the EPIC-pn (van den Eijnden et al. 2017).

Here, we have used the most up to date (as of August 2022) EPIC-pn calibration files. Although they are an improvement on the old files, for example reducing the discrepancy in power-law indices between the EPIC-pn and the *NuSTAR* FPMs, problems do remain. We therefore discard the EPIC-pn observations in this analysis.

3.2.5 Timing Analysis

All the observations are from when H1743-322 was in the hard spectral state. We briefly consider the timing properties by presenting the Poisson noise-subtracted power-spectra¹ of the 3-10 keV light curves of each observation in Fig. 3.2, using $\delta t = 1/32$ s time binning, and ensemble averaging over 128 s segments. In each epoch, the estimate of the power-spectrum from the *NuSTAR* observation is show in blue, with the *XMM-Newton* observations shown in orange and green. Slight differences in the shapes can be put down to the observations not being completely simultaneous and, to a lesser extent, the different effective area curves of EPIC-pn and the

¹For the *NuSTAR* observations, the power spectrum is estimated from the co-spectrum between the FMPA and FMPB light curves to circumvent dead time effects (Bachetti et al. 2015), as described in Section 2.2.4 of this thesis.




Figure 3.2: The power spectra of each epoch estimated for the *NuSTAR* (blue) and *XMM-Newton* (orange and green) observations. In each, on top of the broadband noise a strong type-C QPO is seen, spanning frequencies $0.2 \text{ Hz} \lesssim \nu_{\text{QPO}} \lesssim 1.2 \text{ Hz}$, with a second harmonic at twice the frequency.

NuSTAR FPMs. Each epoch shows a strong type-C low-frequency quasi-periodic oscillation (QPO) on top of the broad-band noise, with often a second harmonic showing. The QPO frequency ranges from $\sim 0.2 \text{ Hz}$ to $\sim 1 \text{ Hz}$. The asymmetric shape of some of the QPOs (for example, 16b) also hint at the frequency evolution of the QPO during the course of the observation. Similarly, discrepancies between the QPOs seen with different telescopes are also a result of this evolution, as can be seen in epochs 14 and 16b. These QPOs will be studied in greater detail in a chapter 4.

3.3 Spectral Fits

3.3.1 The model

3.3.1.1 rtdist

Here we use the version of RELTRANS which was presented in I22, called RTDIST. In this section we summarise the key concepts of the model, but the interested reader should refer to I22 for the full details. The direct continuum is represented by the model NTHCOMP (Zdziarski et al. 1996), which considers Compton up-scattering of a thermal population of seed photons by a thermal population of electrons. The spectrum is a power law with photon index Γ between low and high frequency breaks that are set respectively by the seed photon and electron temperature kT_{bb} and kT_e . The restframe reflection spectrum is represented by XILLVERDCP (García et al. 2013a, 2016), which is computed assuming an NTHCOMP illuminating spectrum. The photon index and electron temperature are free parameters, whereas the seed photon temperature is hardwired to $kT_{\text{bb}} = 0.05$ keV. RTDIST accounts for general relativistic effects by tracing rays in the Kerr metric using the code YNOGK (Yang & Wang 2013). The corona is assumed to be a point-like source located on the BH spin axis a distance $h R_g$ from the BH (i.e. the lamppost model). The angular emissivity of the corona is set by the parameters b_1 , b_2 and the so-called boost parameter $\text{boost} = 1/\mathcal{B}$. If $b_1 = b_2 = 0$ and $\text{boost}=1$, the corona emits isotropically. For $\text{boost} > 1$, emission is beamed towards the BH and for $\text{boost} < 1$ it is beamed away from the BH. The parameters b_1 and b_2 influence the angular emissivity in a more subtle way (see I22 for details). The disc extends from an inner radius r_{in} to an outer radius r_{out} and has a constant scale height h_d/r . We freeze r_{out} to $1000 R_g$ and leave r_{in} as a free parameter. We also freeze the BH spin in our fits to $a_* = 0.998$ to enable the disc truncation radius r_{in} to vary across the widest possible range without necessarily being fixed at the innermost stable circular orbit (ISCO) of the BH, and also never being smaller than the ISCO. The disc is assumed to be in the BH equatorial plane, and we view it from an inclination angle i .

The shape of the restframe reflection spectrum depends on the shape and normalisation of the illuminating spectrum, plus the electron number density of the disc, n_e . XILLVERDCP takes as input parameters Γ , kT_e , the iron abundance relative to solar A_{Fe} , n_e and the ionisation parameter

$$\xi = 4\pi \frac{F_x}{n_e}, \quad (3.1)$$

where F_x is the illuminating X-ray flux. It has long been recognised that ξ is a key parameter for determining ionisation balance in the disc (e.g. Ross & Fabian 2005), but n_e was always frozen to $n_e = 10^{15} \text{ cm}^{-3}$ in older reflection models. García et al. (2016) showed that n_e is also an important parameter in itself, and both ξ and n_e have since been simultaneously constrained with reflection spectroscopy for both AGN and XRBs (see e.g. Tomsick et al. 2018; Jiang et al. 2019a,b; García et al. 2019). As is theoretically expected (Shakura & Sunyaev 1973), n_e is found to be larger for XRBs than for AGNs. It is the simultaneous constraint of ξ and n_e that ultimately means that F_x can be constrained, thus making the spectrum sensitive to the ratio M/D . RTDIST assumes a radial density profile² $n_e(r) \propto r^{3/2}(1 - \sqrt{r_{\text{in}}/r})^{-2}$ (Shakura & Sunyaev 1973) and takes as input parameters the normalisation of this relation, hereafter referred to as n_e , and the distance D as input parameters. The ionisation parameter as a function of radius $\xi(r)$ is then calculated from input parameters including D , M and n_e (see equation 10 of I22). The parameter n_e is defined as $n_e(r)$ at the radius at which $\xi(r)$ reaches its peak value. We use a grid of 50 radial bins for this calculation.

3.3.1.2 Complete model

The complete model within XSPEC we use is

$$\text{CONST} \times \text{TBABS} \times (\text{DISKBB} + \text{RTDIST} + \text{XILLVERCP}). \quad (3.2)$$

In order, these components are used as follows:

²The code also provides the option to assume that n_e is constant with radius, but here we use a radial density profile.

- **CONST:** We use a scalar value (which is constant with the spectral energy) to ensure an accurate absolute flux calibration. For *NuSTAR*, this constant has been calibrated to within a $\sim 4\%$ systematic uncertainty via observations of the Crab nebula in ‘stray light’ (Madsen et al. 2017), which is light that bypasses the mirror assembly. Following the analysis of Madsen et al. (2017), we use $\text{CONST} = 0.86$ for *NuSTAR*’s FPMA, and we consider an uncertainty on this of 4% when it comes to exploring parameter uncertainties. We leave this spectral constant as a free parameter for *NuSTAR*’s FPMB and the *XMM-Newton*’s RGS.
- **TBABS:** The effects of a line of sight absorption column, using the abundances of Wilms et al. (2000).
- **DISKBB:** This is used to model the intrinsic thermal disc spectrum.
- **RTDIST:** The bulk of our model, described in the preceding sub-sub-section.
- **XILLVERCP:** A non-relativistic distant reflector, with a spectrum unmodified by relativistic effects.

3.3.2 Model fitting procedure

Our strategy for model fitting considers *individual fits*, in which we fit the model to each of our five epochs individually, and then a *joint fit* when we fit the five epochs simultaneously, tying some key parameters to be the same for each epoch. For each of these, we first find a best fitting model using χ^2 minimisation, before using a Markov chain Monte Carlo (MCMC) simulation to explore parameter space. We use XSPEC v. 12.12.0 throughout, for both χ^2 minimisation and for running the MCMC. For all fits, we consider data from the *XMM-Newton* RGS1 and RGS2 and the *NuSTAR* FPMA and FPMB.

3.3.2.1 Individual fits

In order to find an initial fit, we freeze $D = 8.5$ kpc (Steiner et al. 2012). This is a sensible choice since the spectral model is only dependent on M/D , with a complete degeneracy between M and D because we are not considering timing properties in this study. Therefore, we are effectively simplifying the model to have one model parameter $M_{8.5} = M \frac{8.5 \text{ kpc}}{D}$ in place of the two parameters M and D . We also freeze $\text{CONST} = 0.86$ for the FPMA but leave it as a free parameter for FPMB, RGS1 and RGS2. With the exception of CONST , the rest of the model parameters are tied to be the same for all instruments. Epoch 14 includes two separate *XMM-Newton* observations. We tie the parameters for these two observations to be the same as one another.

Early analysis found that the model didn't require the quadratic angular emissivity b_1, b_2 parameters, or the disc-scale height h_d/r to be non-zero at anything other than very low significance, hence we freeze these to zero for our model fitting. This corresponds to considering an emissivity profile modified from isotropic only by the 'boost' parameter, and considering only a flat disc.

We initialise the MCMC simulation from the best fitting model found via χ^2 minimisation. We use an isotropic prior on the inclination, and a Jeffreys ($\propto 1/p$) prior on the normalisations of the model components and $M_{8.5}$ in order to ensure that the parameter space is equally considered. At the end of the analysis we provide a distance to each step of the MCMC by sampling randomly from the marginalised MCMC posterior of Steiner et al. (2012), and calculate the mass value associated with the value of $M_{8.5}$ for that step³. Along with considering the distance, we also wrap the uncertainty on the FMPA calibration in a similar way, by assigning a new random value $\sim \mathcal{N}(\mu = 0.86, \sigma = 4\%)$ for the FMPA normalisation, and scaling the model normalisations to calculate a new mass value. As M^2 scales with the observed flux, this propagates a 2% error into M . Our final posterior distribution

³The value of 8.5 kpc comes from the mode of the distance distribution. Here we present parameters using the mean, so due to the skew of distance posterior this is closer to 8.1 kpc. However, for completeness the full posterior for the mass will be presented.

on mass therefore includes uncertainty on the distance, as well as uncertainty on the absolute flux calibration of the instruments used.

3.3.2.2 Joint fit

We initialise our joint fit from the results of the individual fits. We then tie $M_{8.5}$, inclination i and iron abundance A_{Fe} to be the same for all 5 epochs and find a new best fit by minimising χ^2 . All other previously free parameters remain free. As with the individual fits, we then move to an MCMC. Again, we allow the instrument normalisations to be free, with the same priors as before.

3.3.3 Results

3.3.3.1 Individual fits

The spectra of the best fitting models to each epoch are shown on the left-hand side of Fig. 3.3, and so is the ratio of the data to the best fitting model.

For each individual fit we ran a MCMC until convergence, and then running 500 walkers each for 7500 steps. We trimmed the MCMC by taking every 100th step from each walker, to produce a final sample with 37500 parameter combinations. From this, we show the posterior distributions of some key parameters, in the upper panels of the subplots in Figs. 3.4 and 3.5, and the posterior mean and 1σ credible interval for each parameter is listed in Table 3.3.

3.3.3.2 Joint fit

The spectra of the best fitting joint model, with $M_{8.5}$, inclination and iron abundance tied across all five epochs, are shown on the right side of Fig. 3.3. For this joint fit we ran an MCMC until convergence, and then running 200 walkers each for 52500 steps. We trimmed the MCMC by taking every 100th step from each walker, to produce a final sample with 105000 parameter combinations. The lower panels of the subplots in Figs. 3.4 and 3.5 show the posterior distributions of key parameters for the joint chain, including parameters which are tied between epochs, and those that are still free in each epoch. Table 3.4 lists the posterior mean of each parameter, along with the the 1σ credible interval.

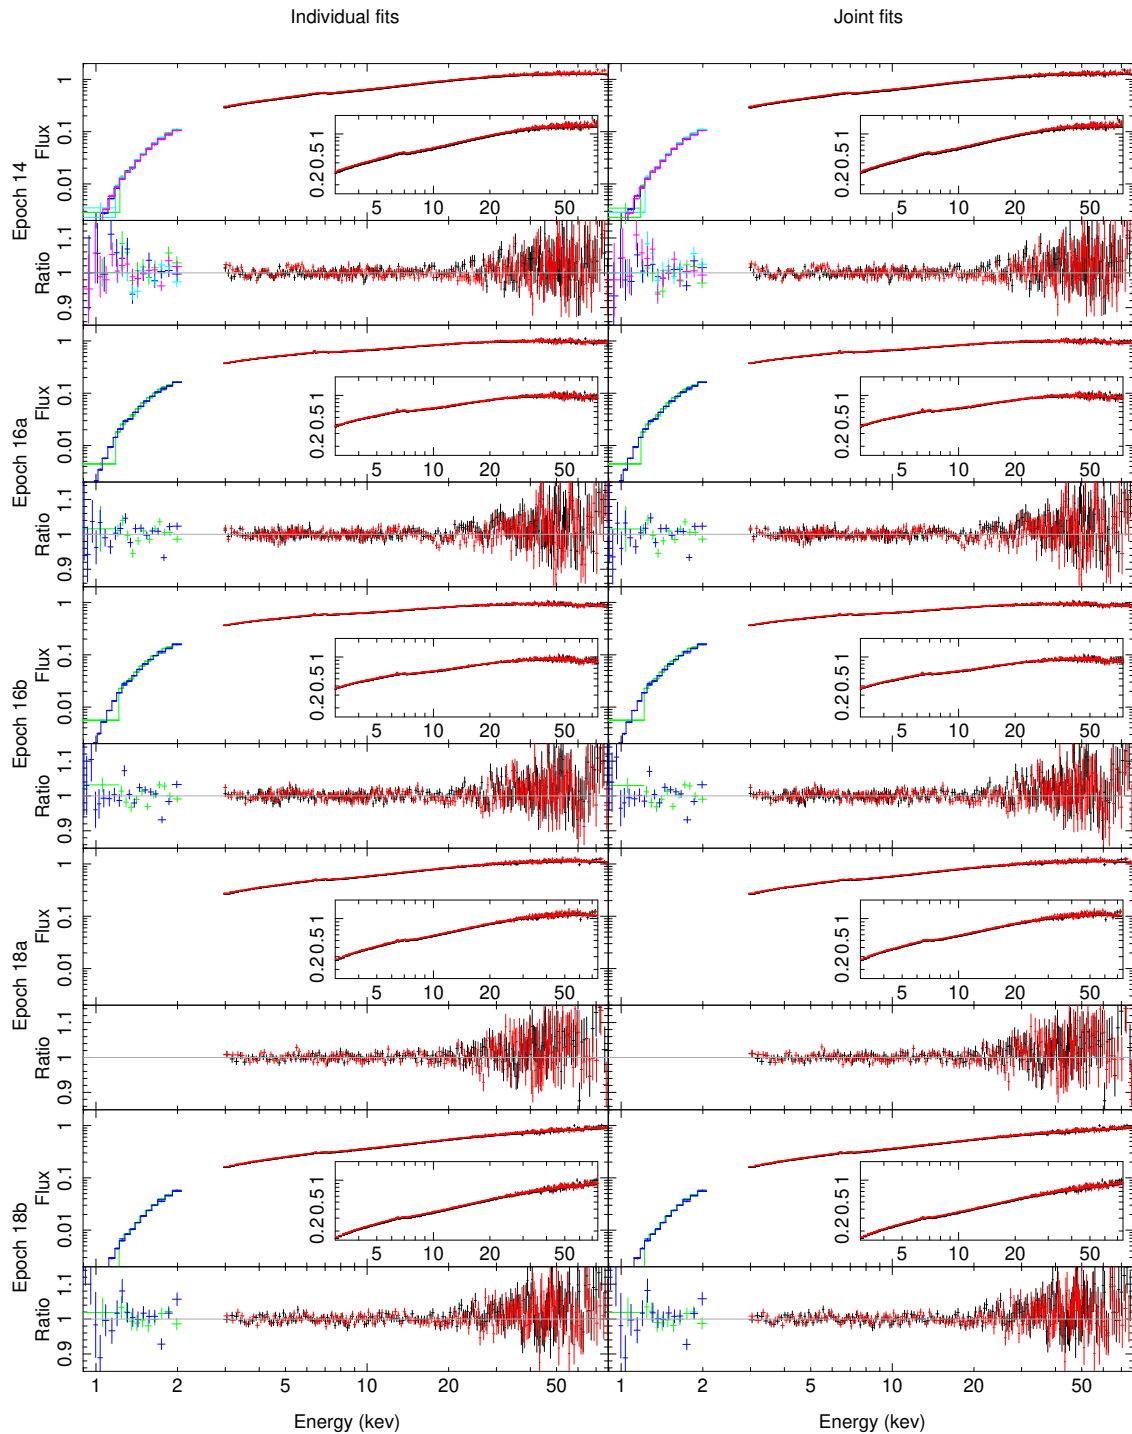


Figure 3.3: The observations and best fitting models unfolded around the telescope response matrices, plus the ratio of the data to the best fitting models. The flux is given in units of $\text{keV}^2 \text{ photons/s/cm}^2/\text{keV}$. On the left is the best fitting model to each epoch individually, while on the right is each epoch compared to the model found in the fit with the mass, inclination, and iron fraction tied. For comparison, the axes all have the same scale. The *NuSTAR* data are shown in black (FPMA) and red (FPMB), while *XMM-Newton* data are shown in green, cyan (RGS1), and blue, magenta (RGS2). The insets show a zoom-in of the 3 – 75 keV region to show the *NuSTAR* data closely.

Variable	Epoch 14	Epoch 16a	Epoch 16b	Epoch 18a	Epoch 18b
FMPB calibration	0.8974 ± 0.0009	0.8781 ± 0.0008	0.8785 ± 0.0009	0.8854 ± 0.0013	$0.8889^{+0.0014}_{-0.0013}$
RGS1 calibration	0.671 ± 0.009	$0.7547 \pm +0.01$	$0.731^{+0.009}_{-0.010}$	—	$0.682^{+0.011}_{-0.010}$
RGS2 calibration	0.646 ± 0.009	0.725 ± 0.010	$0.700^{+0.009}_{-0.010}$	—	0.657 ± 0.011
$M_{8.5}$ (M_{\odot})	14^{+6}_{-7}	15 ± 5	17 ± 5	80 ± 40	36 ± 15
M (M_{\odot})	14 ± 6	14 ± 5	16 ± 5	70 ± 40	35^{+15}_{-14}
i (degrees)	72 ± 4	73 ± 3	66^{+7}_{-6}	69 ± 7	59^{+6}_{-7}
A_{Fe}	0.69 ± 0.07	0.52 ± 0.02	0.54 ± 0.03	$0.68^{+0.15}_{-0.14}$	0.72 ± 0.07
RTDIST norm	$(1.5 \pm 0.4) \times 10^{-2}$	$(0.90 \pm 0.09) \times 10^{-2}$	$(0.90^{+0.12}_{-0.11}) \times 10^{-2}$	$(1.7 \pm 0.4) \times 10^{-2}$	$(4 \pm 2) \times 10^{-2}$
DISKBB norm	9^{+7}_{-6}	$(8 \pm 3) \times 10^4$	$(2.2^{+1.3}_{-1.2}) \times 10^4$	$0.5^{+0.3}_{-0.4}$	$(2 \pm 2) \times 10^7$
Dist. refl. norm	$(7 \pm 2) \times 10^{-3}$	$(1.8 \pm 0.5) \times 10^{-3}$	$(0.9 \pm 0.4) \times 10^{-3}$	$(0.9^{+0.5}_{-0.6}) \times 10^{-3}$	$(0.4^{+0.3}_{-0.4}) \times 10^{-3}$
Γ	1.605 ± 0.006	1.741 ± 0.011	1.729 ± 0.012	1.642 ± 0.014	1.582 ± 0.011
T_{in} (keV)	0.83 ± 0.14	$0.201^{+0.007}_{-0.008}$	$0.224^{+0.015}_{-0.014}$	1.5 ± 0.4	0.091 ± 0.007
boost	6 ± 2	4 ± 2	3 ± 2	$2.0^{+1.4}_{-1.3}$	4 ± 2
h (R_g)	21 ± 6	16 ± 2	12 ± 2	8 ± 2	4.5 ± 0.9
kT_e (keV)	58 ± 8	130 ± 50	290 ± 150	39^{+7}_{-6}	330^{+120}_{-130}
$\log[N_e/\text{cm}^{-3}]$	19.5 ± 0.4	19.7 ± 0.3	19.5 ± 0.3	$18.20^{+1.09}_{-0.94}$	$18.2^{+0.5}_{-0.4}$
$\log[\xi/(\text{erg cm s}^{-1})]$	3.75 ± 0.05	0.6 ± 0.4	$0.8^{+0.5}_{-0.6}$	$1.2^{+0.8}_{-0.9}$	1.1 ± 0.8
N_{H} (cm^{-2})	3.00 ± 0.02	3.33 ± 0.06	3.10 ± 0.07	$3.16^{+0.14}_{-0.15}$	3.16 ± 0.04
r_{in} (R_g)	120 ± 30	6 ± 2	7 ± 2	$31.4^{+10.5}_{-9.8}$	28 ± 7

Table 3.3: A summary of the parameters in the individual spectral fits, with the results provided by the MCMCs given as the posterior mean with the 1σ credible interval. The calibration constants are given relative to the fixed FMPA calibration of 0.86, and as described in the text the BH Mass is calculated from $M_{8.5}$, with random sampling from a distance distribution and FPMA normalisation.

Variable	Epoch 14	Epoch 16a	Epoch 16b	Epoch 18a	Epoch 18b
FMPB calibration	$0.8974^{+0.0007}_{-0.0008}$	0.878 ± 0.001	0.8786 ± 0.0009	$0.8853^{+0.0013}_{-0.0012}$	$0.8889^{+0.0011}_{-0.0012}$
RGS1 calibration	0.675 ± 0.007	$0.758 \pm +0.01$	0.740 ± 0.007	—	0.677 ± 0.014
RGS2 calibration	0.650 ± 0.007	$0.727^{+0.010}_{-0.009}$	0.709 ± 0.006	—	$0.652^{+0.013}_{-0.014}$
$M_{8.5}$ (M_{\odot})			12.4 ± 1.0		
M (M_{\odot})			12 ± 2		
i (degrees)			73.3 ± 1.4		
A_{Fe}			$0.54^{+0.02}_{-0.03}$		
RTDIST norm	$(1.40^{+0.12}_{-0.13}) \times 10^{-2}$	$(8.3 \pm 0.3) \times 10^{-3}$	$(8.3 \pm 0.4) \times 10^{-3}$	$(1.26^{+0.08}_{-0.07}) \times 10^{-2}$	$(2.1 \pm 0.3) \times 10^{-2}$
DISKBB norm	$3.4^{+1.3}_{-1.1}$	$(1.0 \pm 0.3) \times 10^5$	$(4 \pm 1) \times 10^4$	$1.1^{+0.4}_{-0.5}$	$(1.8^{+1.6}_{-1.4}) \times 10^7$
Dist. refl. norm	$(7.1 \pm 0.9) \times 10^{-3}$	$(2.0 \pm 0.4) \times 10^{-3}$	$(1.6 \pm 0.4) \times 10^{-3}$	$(0.9 \pm 0.3) \times 10^{-3}$	$(1.8 \pm 0.3) \times 10^{-3}$
Γ	1.610 ± 0.004	1.752 ± 0.007	1.736 ± 0.006	1.631 ± 0.006	$1.605^{+0.007}_{-0.006}$
T_{in} (keV)	0.95 ± 0.07	0.196 ± 0.008	0.207 ± 0.007	$1.37^{+0.16}_{-0.15}$	$0.0929^{+0.011}_{-0.008}$
boost	$5.5^{+0.6}_{-0.5}$	1.1 ± 0.4	$3.1^{+0.8}_{-0.7}$	0.6 ± 0.2	1.9 ± 0.6
h (R_g)	17 ± 2	15 ± 2	12.2 ± 0.9	9 ± 1	5.2 ± 0.5
kT_e (keV)	59 ± 7	80 ± 20	100 ± 20	$34.4^{+1.4}_{-1.3}$	480 ± 20
$\log[N_e/\text{cm}^{-3}]$	19.90 ± 0.08	19.83 ± 0.11	$19.93^{+0.05}_{-0.06}$	19.8 ± 0.2	19.4 ± 0.2
$\log[\xi/(\text{erg cm s}^{-1})]$	$3.70^{+0.05}_{-0.04}$	$0.17^{+0.15}_{-0.13}$	$0.5^{+0.5}_{-0.4}$	$2.6^{+0.3}_{-0.4}$	1.8 ± 0.3
N_{H} (cm^{-2})	2.99 ± 0.02	3.35 ± 0.06	3.18 ± 0.05	$3.43^{+0.15}_{-0.16}$	3.17 ± 0.05
r_{in} (R_g)	83 ± 4	$5.4^{+1.3}_{-1.4}$	4.9 ± 0.5	26 ± 4	26 ± 4

Table 3.4: A summary of the parameters in the joint spectral fit, with the results provided by the MCMCs given as the posterior mean with the 1σ credible interval. The calibration constants are given relative to the fixed FMPA calibration of 0.86. As described in the text the BH Mass is calculated from $M_{8.5}$, with random sampling from a distance distribution and FPMMA normalisation. The parameters $M_{8.5}$ (and therefore BH mass), inclination, and A_{Fe} are tied between the five epochs in this joint fit, so there is only one value for each of them.

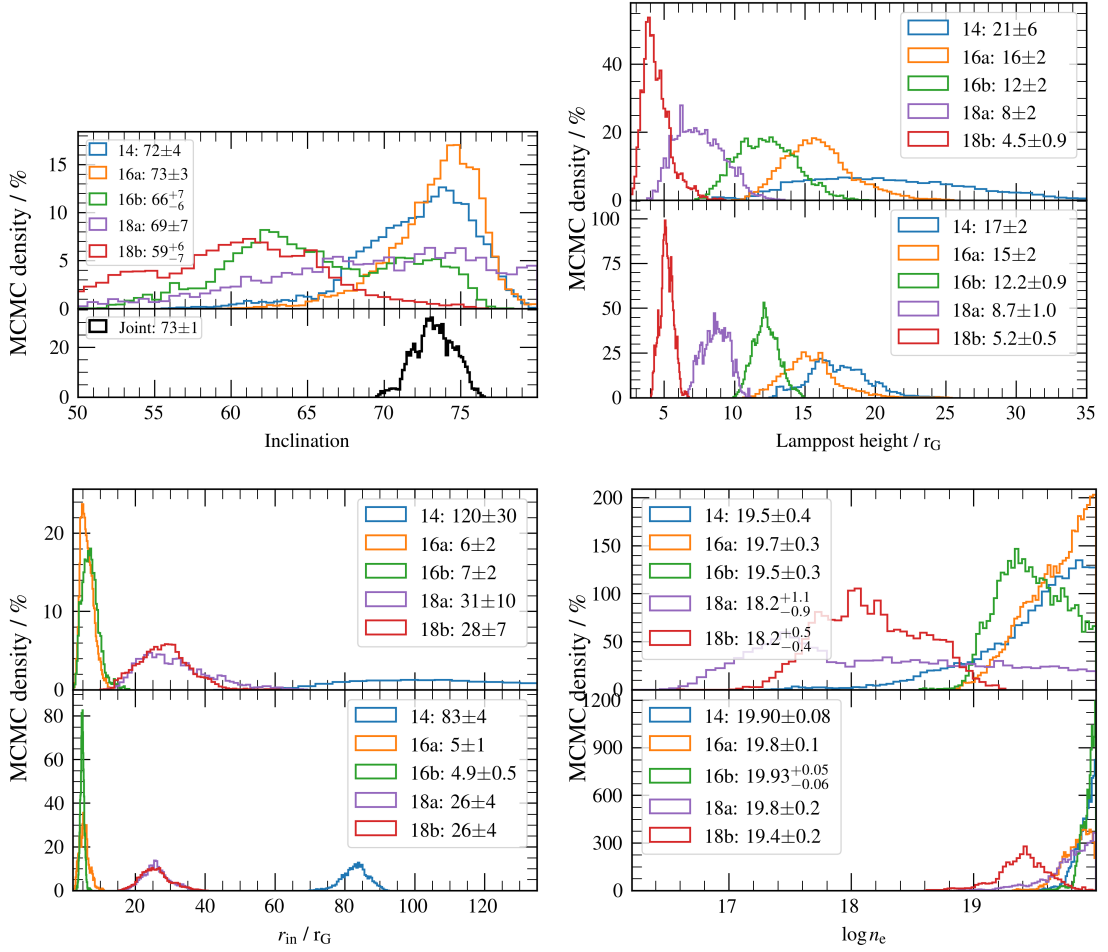


Figure 3.4: Histograms showing the MCMC posterior densities (in units of %) for the four parameters: inclination, lamppost height, disc inner radius, and disc electron density. Each subfigure is split into an upper and lower panel. The upper shows the results for the MCMCs run on the individual spectral fits, whereas the lower shows the result for same parameters in the MCMC run on the joint spectral fit. The five epochs are colour-coded as labelled. As the inclination was tied between the epochs in the joint spectral fit, this is show in black. The legends list the posterior means along with the 1σ credible interval. It should be noted that the model was capped with $\log N_e \leq 20$.

3.4 Discussion

We have modelled the spectra of H1743-322 across 5 long-look observations with a reflection model that is sensitive to the ratio of BH mass to distance. From inputting a previous distance measurement as a Bayesian prior, we obtain a mass measurement of $M = 12 \pm 2 M_\odot$. This is the first mass measurement of H1743-322, since extinction has prevented the detection of absorption lines in the optical




Figure 3.5: The MCMC posterior density of the model parameter $M_{8.5}$, and then with this converted to the BH mass. The histograms are stylistically the same as those in Fig. 3.4. $M_{8.5}$ was the raw model parameter, from which the BH mass was calculated from $M_{8.5}$ by assigning each step within the MCMC a random distance and FPMA efficiency as described in the text.

spectrum of the companion star and thus it has not been possible to constrain a radial velocity curve. Here we discuss the implications of our results, along with caveats and suggested improvements to the model.

3.4.1 Proof-of-Principle

This study is among the first to measure BH mass from the X-ray spectrum alone. The first to our knowledge was Parker et al. (2016), who estimated the BH mass of and distance to GX 339-4 by simultaneously fitting the disc and reflection components with the disc inner radius tied between the two. Mastroserio et al. (2019) used an earlier version of RELTRANS to measure the mass of Cygnus X-1 by simultaneously modelling the X-ray spectrum and timing properties. This X-ray reverberation mass measurement was slightly higher than the existing optical measurement (Orosz et al. 2011), but Miller-Jones et al. (2021) have since revised the optical measurement (based on a revised distance estimate) to agree remarkably well with Mastroserio et al. (2019).

Here, we have used a newer version of RELTRANS (RTDIST: Ingram et al. 2022) to measure BH mass from the X-ray spectrum alone, without considering timing properties. The new model is sensitive to the ratio M/D because the ionisation parameter used to calculate the restframe reflection spectrum is calculated self-consistently as a function of disc radius from the observed reflected flux, the disc

density, the BH mass and the distance (see equation 10 of Ingram et al. 2022). We constrain the mass by inputting a previous measurement of the distance (Steiner et al. 2012). This study represents the first test of the model. It is encouraging that we obtain a reasonable mass measurement of $M = 12 \pm 2 M_{\odot}$. It is by no means trivial that this analysis should have yielded a BH mass in the expected range – the model is free to predict any BH mass. We see our result as a proof-of-principle of the basic validity of our assumptions. Similar proof-of-principle tests were conducted by Shreeram & Ingram (2020), Zdziarski et al. (1996), and Connors et al. (2021) who modelled spectra of XRBs by fitting for the ionisation parameter, but then checked that their result was consistent with the ionisation parameter theoretically expected given the observed reflected flux, the measured disc density and the estimated BH mass and distance.

There are a number of simplifying assumptions in the model that could, in principle, have prevented it from returning a reasonable mass estimate. First, the XILLVERDCP model used to calculate the restframe reflection spectrum ignores the effect of intrinsic disc emission on the ionisation balance, instead assuming that the disc atmosphere is exclusively heated from above by the corona. This assumption is better for lower disc temperatures, such as those associated with the hard state observations we analyse here. Second, the XILLVERDCP tables hardwire the seed photon temperature to $kT_{\text{bb}} = 0.05$ keV, which is lower than the expected seed photon temperature for these observations, if they are provided by the disc as is commonly assumed. The broadband shape of the illuminating spectrum influences the ionisation balance in the disc atmosphere, and thus the assumed value of kT_{bb} could well influence the BH mass inferred from our method. Ideally, future versions of the model will have kT_{bb} as a free parameter, but it is at least encouraging that freezing it to a low value does not appear to return a wildly incorrect result. Finally, the radial disc emissivity profile in RELTRANS is calculated in the lamppost geometry, which is presumably a highly simplified representation of the true, extended corona. Again, it is very encouraging that this assumption does not appear to prevent us from obtaining a reasonable result.

In future, we will conduct further, more stringent tests of the model by additionally considering X-ray timing properties. These tests will be particularly useful if we consider XRBs with existing mass and distance measurements, which can be compared with the values of these properties yielded by a full X-ray reverberation mapping analysis. Successful tests of the model are positive for the prospects of using the same model to constrain the Hubble constant by measuring the distance to a sample of AGNs (Ingram et al. 2022).

3.4.2 Mass Measurement

Özel et al. (2010, but also see Farr et al. 2011) used known mass measurements of 16 XRBs to fit an exponentially decaying mass function (which was motivated by theoretical expectations), and estimated the mass function

$$P(M) = \begin{cases} \frac{1}{1.57M_{\odot}} \exp\left[\frac{6.30M_{\odot}-M}{1.57M_{\odot}}\right], & M > 6.3M_{\odot}, \\ 0, & M \leq 6.3M_{\odot}. \end{cases} \quad (3.3)$$

Our measured mass of $M = 12 \pm 2 M_{\odot}$ is towards the high end of this distribution. With the caveat that this work is primarily a proof-of-principle test of a new method, we do note that it is perhaps theoretically expected for a high-extinction source such as H1743-322 to have a larger BH mass than the low-extinction sources for which optical mass measurements have been possible (Jonker et al. 2021). This follows because XRB progenitor stars likely formed in the dense regions of the Galaxy, with some being ejected to their current high Galactic latitudes by a supernova kick (Blaauw 1961; Verbunt et al. 2017). The prediction that the largest BHs formed via direct collapse with no associated kick (Fryer & Kalogera 2001; Fryer et al. 2012) therefore implies that the XRBs that remained in dense regions after BH formation are systematically more massive than those that were ejected.

3.4.3 Higher Spin

Putting concerns aside, we consider the implications of this mass measurement on the spin of the BH in H1743-322. Steiner et al. (2012) considered soft-state X-ray observations of H1743-322 (in which it can be assumed that $r_{\text{in}} = r_{\text{ISCO}}$) to

estimate the BH spin a_* . For this estimate, they used the distance measurement that we also use here, as well as inputting the BH mass distribution from eqn 3.3 as a mass prior, in the absence of a mass measurement for this source. Since their soft-state spectral modelling combined with their distance measurement constrains the disc inner radius in km, a larger assumed mass decreases the inferred r_{in} in units of R_g and thus increases the inferred spin. From their Figure 3, we can see that our mass measurement implies a spin of $a_* \sim 0.5 - 0.6$.

3.4.4 Inner radius and QPOs

Fig. 3.2 demonstrates the type-C QPO seen in each of the observational epochs, spanning the frequency range $0.2 \text{ Hz} \lesssim \nu_{\text{QPO}} \lesssim 1.1 \text{ Hz}$. In Fig. 3.6 we compare the QPO frequency in each epoch with our measured disc inner radius, r_{in} (from the joint fit). We see that the inner radius is consistent with reducing as the QPO frequency increases, as is predicted for a large array of QPO models (Ingram & Motta 2019). We additionally plot this predicted relationship for the precessing inner flow model of Ingram et al. (2009). In this model, the entire corona precesses due to the frame dragging effect. The precession frequency depends on the BH mass and spin, the inner radius of the corona, the outer radius of the corona (assumed to be the inner radius of the disc) and the power law index ζ characterising the surface density profile within the corona, $\Sigma(r) \propto r^{-\zeta}$. We calculate the precession frequency for a range of disc inner radii by fixing M to our best fitting value, the spin to $a = 0.5$ as implied by our new mass estimate, and each line represents a different assumption about the corona inner radius and ζ . We consider three values for ζ : $\zeta = 0.5$ (solid lines) from classical advection dominated accretion flows; $\zeta = 0.0$ (dashed lines) as approximately seen in numerical simulations; and $\zeta = -0.45$ (dotted lines) close to a thin disc. The corona inner radius was either taken to be the ISCO of the BH (black lines) or the bending wave radius for a disc scale height $h_d/r = 0.15$ (grey lines; Lubow et al. 2002; Ingram et al. 2009).

We see that the models all reproduce the reduction in disc inner radius with QPO frequency, as expected, but our measured r_{in} values for high QPO frequency

are smaller than the model predictions – as is common for reflection modelling of bright hard state and intermediate state observations (e.g. García et al. 2015; Fürst et al. 2016). This could of course be due to the precession model being incorrect or incomplete, or could instead be due to systematic errors associated with reflection modelling. For example, here we represent the corona as a single temperature region, whereas an extended, stratified corona would be better represented by a sum of Comptonised spectra each with different power-law indices and high energy cut offs. The resulting curvature of the continuum spectrum typically leaves a weaker red wing in the iron line residuals, thus indicating a larger disc inner radius (e.g. Basak et al. 2017; Zdziarski et al. 2022). Alternatively the discrepancy could be because the formula we have used for the precession frequency ignores the torque that the precessing corona presumably in reality feels from the outer disc. Since this torque acts to slow down precession, accounting for it would decrease the disc inner radius inferred from the precession model.

3.4.5 Density and inclination

We find that a high disc density is required in our fits, particularly in the joint fit (Figure 3.4). This is theoretically expected for an XRB ($n_e \propto 1/M$ for a radiation pressure dominated disc; Shakura & Sunyaev 1973). However, we see that we are unable to sample the high density end of the n_e posterior due to the XILLVERDCP grid having a maximum density of $n_e = 10^{20} \text{ cm}^{-3}$. This is due to numerical issues associated with larger densities. In the model, $M \propto 1/n_e$ and so our inability to sample the highest densities means that we very likely under predict the low mass region of the mass posterior. This model systematic could potentially be very important, and our result has to be considered in the context of the particular reflection models we have used, including all the relevant assumptions these models include. We cannot investigate the high density until reliable models are available, however the statistical uncertainty of our result demonstrates how this method could be used when using future versions of reflection models.




Figure 3.6: The disc inner radius (mean and 1σ credible interval) of each epoch from the MCMC of the joint fit, as compared to the QPO frequency measured in the 3 – 10 keV band (see Fig. 3.2). The lines are examples of the QPO frequency vs disc inner radius from Eq. 2 of Ingram et al. (2009), using our best fit mass of $12 M_{\odot}$ and an assumed higher spin $a = 0.5$. The black lines have the precessing flow’s inner radius set by the bending radius with $h/r = 0.15$, whereas the grey lines show the flow down to the ISCO. Three different surface density power laws were used: $\zeta = 0.5$ (solid line); $\zeta = 0.0$ (dashed line); and $\zeta = -0.45$ (dotted line).

From the joint fit, we find a disc inclination of $[73.3 \pm 1.4]^{\circ}$, which is consistent with jet inclination of $[75 \pm 3]^{\circ}$ found by Steiner et al. (2012), and also with a $\sim 15^{\circ}$ misalignment (given the unknown projection of the disc rotation axis on the sky) as found for Ingram et al. (2017).

3.5 Conclusions

Here we present the first mass-measurement of the BH in the XRB H1743-322, among the first to use the X-ray spectrum alone. By using the X-ray reflection spectrum shape to infer the X-ray flux incident onto (and therefore reflected from) the disc, and comparing that to the flux we observe, we are able to constrain the BH mass-distance ratio M/D . Past measurements of the source distance using the radio jets by Steiner et al. (2012) gave a distance of $D = 8.5 \pm 0.8$ kpc. We use our model to measure $M_{8.5}$, the BH mass if it were at 8.5 kpc, before folding

in the measurement uncertainty using an MCMC analysis to obtain a posterior mass distribution with mean $M = 12 \pm 2 M_{\odot}$. Since our posterior implies a higher mass than the Galactic XRB mass distribution that was previously used by Steiner et al. (2012) to measure the BH spin, it also implies a higher spin of $a \sim 0.5 - 0.6$ than that inferred by Steiner et al. (2012). This larger value is more in line with the spins previously inferred from reflection spectroscopy.

‘While I’m still confused and uncertain, it’s on a much higher plane, d’you see, and at least I know I’m bewildered about the really fundamental and important facts of the universe.’

— Terry Pratchett, *Equal Rites*

4

Phase-resolved spectroscopy of quasi-periodic oscillations in H1743-322 across three outbursts

Contents

4.1	Introduction	119
4.2	Data Analysis	120
4.2.1	Data selection and reduction	120
4.2.2	Phase-resolving method	121
4.2.3	Results	123
4.3	Phase-resolved fits	130
4.3.1	Fitting procedure	130
4.4	Discussion and Conclusions	137

4.1 Introduction

The observations of the black hole (BH) X-ray Binary (XRB) H1743-322 which I analysed in Chapter 3 all showed a strong type-C low frequency quasi-periodic oscillation (QPO) in their power spectrum. While, during that work, I briefly considered these QPOs I did not analyse them in detail. Some of the data, that of the 2014 outburst, had been previously analysed by Ingram et al. (2016) and Ingram et al. (2017), who not only found strong evidence (3.7σ) of the iron line energy varying

with QPO phase, but also detected the variation of the reflection fraction with QPO phase (3.5σ), and were able to rule out the null hypothesis of the axi-symmetric, QPO phase independent emissivity profile at a moderate significance (2.4σ).

I therefore used the machinery I had built during my work from chapter 2 to phase-resolve the QPO in the H1743-322 observations I have, and fit them with the updated QPO spectral model also from that chapter.

4.2 Data Analysis

4.2.1 Data selection and reduction

We consider the same data used in chapter 3, as each observation shows a strong type-C QPO (see Fig. 3.2). We see that the *XMM-Newton* EPIC-pn data provide higher quality timing information than the *NuSTAR* focal plane modules (FPMs). We therefore include the EPIC-pn data in our analysis despite the cross-calibration difficulties that this incurs, and ignore the lower signal to noise RGS data. This consideration also leads us to analyse only epochs with simultaneous *XMM-Newton* and *NuSTAR* exposures: epochs 14, 16a, 16b, and 18b (i.e. we do not consider epoch 18a, for which there is only a *NuSTAR* observation). Note that the *XMM-Newton* observation in epoch 14 is split into two observations for programmatic reasons (obs IDs 0724401901 and 0740980201), which we treat separately in our analysis. *XMM-Newton* also observed H1743-322 during the orbit prior to our epoch 14 observations (obs ID 0724400501). Ingram et al. (2016) included this earlier data in their analysis to increase signal to noise, whereas here we include only the latter orbit to ensure simultaneous coverage by *NuSTAR*.

The data reduction for both the *XMM-Newton* and *NuSTAR* was the same as described in Sections 3.2.3.1 and 3.2.2. We address cross-calibration discrepancies (Section 3.2.4 and Fig. 3.1) by including a 0.5% systematic error on the EPIC-pn phase-averaged spectra using the FTOOL GRPPHA.

4.2.2 Phase-resolving method

We apply the same phase-resolving method as in chapter 2 (based upon the method of Ingram et al. 2016), which is described in detail in Section 2.3. As a summary, we consider the light curve in every energy channel to vary with QPO phase γ , with a two-harmonic waveform given by

$$w(E, \gamma) = \mu(E) \left[1 + \sum_{j=1}^2 \sigma_j(E) \cos(j\gamma - \Phi_j(E)) \right]. \quad (4.1)$$

Such a waveform is described by five parameters: the average count rate of that energy channel $\mu(E)$; the fractional variability amplitude (i.e., the RMS) of each of the two harmonics $\sigma_j(E)$; and the phase offset of each of the harmonics $\Phi_j(E)$. We measure every phase as relative to the first harmonic in a *reference band* taken as the light curve covering the entire energy range of the instrument, and we follow Ingram et al. (2016) by defining this as $\Phi_1 = \pi/2$. To find the phases $\Phi_j(E)$ in each energy band, we measure the phase lag of each of harmonic relative to the same harmonic in the reference band $\Delta_j(E) = \Phi_j(E) - \Phi_j$. We also find the phase difference between the harmonics within the reference band $\psi = \frac{1}{2}\Phi_2 - \Phi_1$ (where the half is to the harmonic relation of the frequencies each phase is measured at). The phase offsets of the two harmonics are then

$$\begin{aligned} \Phi_1(E) &= \Phi_1 + \Delta_1(E) \\ \Phi_2(E) &= 2(\Phi_1 + \psi) + \Delta_2(E). \end{aligned} \quad (4.2)$$

The RMS and phase lag spectra of the two harmonics are calculated the same way as in Sections 2.3.2 and 2.3.3, through the use of the cross-spectrum between the subject and reference band $G(\nu, E) = \langle S(\nu, E)R^*(\nu) \rangle$. The phase lag associated with each harmonic is calculated by averaging the cross-spectrum over the full width at half maximum (FWHM) of that harmonic and then taking the argument. This assumes that the QPO dominates strongly over the broad band noise. In a small improvement to the method upon that described in Section 2.3, we now calculate the FWHM of each harmonic from a simple three Lorentzian fit to the reference band power spectrum whereas before we had simply assumed a quality factor $Q = 8$ typical of type-C QPOs.

To find the RMS of each QPO harmonic, we estimate the power spectrum of the subject band through (Wilkinson & Uttley 2009; Ingram et al. 2016)

$$P_s(\nu, E) = \frac{|G(\nu, E)|^2 - \hat{b}^2(\nu, E)}{P_r(\nu)}, \quad (4.3)$$

where $\hat{b}^2(\nu, E)$ is a bias term relating to the squaring of the cross-spectrum (see e.g. Ingram 2019). For subject bands highly correlated with the reference band, as is the case here, the above expression is equal to the Poisson noise subtracted power spectrum of the subject band light curve, except with far smaller statistical uncertainties. By fitting a three Lorentzian model to these power spectra we then find the amplitude of both QPO harmonics as a function of energy. As before, for the *NuSTAR* data we use FMPA as the subject band and FMPB as the reference band to eliminate Poisson noise and dead time effects, whereas for *XMM-Newton* we rely on the very short dead time to allow us to simply subtract the Poisson noise.

Similarly, the phase difference between the harmonics in the reference band is calculated using the same method as before (Section 2.3.4). We use the ‘auto-bispectrum’, $B(\nu)$, which is equal to the bi-spectrum $\mathcal{B}(\nu_1, \nu_2)$ (e.g. Maccarone et al. 2011) in the limit of the two frequencies ν_1 and ν_2 being equal to one another,

$$B(\nu) = \mathcal{B}(\nu, \nu) = \langle R^2(\nu)R^*(2\nu) \rangle. \quad (4.4)$$

We then calculate the phase difference ψ from the auto bi-spectrum by first correcting for Poisson noise, then taking the argument, then dividing by -2 .

The frequency of the QPO is known to drift over the course of an observation. To account for this, we track the QPO frequency throughout the observation and average the power, cross and bi- spectra over the instantaneous QPO frequency corresponding to each data segment, to create ‘shift-and-added’ versions of these spectra (e.g. Méndez et al. 1998; de Avellar et al. 2013). To track the QPO frequency with time, we use the same technique as is described in Section 2.3.1. This involves first calculating the power spectrum of the reference band light curve for segments of duration 256 s, and then fitting each one with a model comprised of four Lorentzians plus Poisson noise. For *XMM-Newton*, the Poisson noise model is simply constant

with frequency, whereas for *NuSTAR* we use a sinc function due to dead time effects (Bult 2017). We fit using maximum likelihood estimation to account for the non-Gaussian (χ^2 with two degrees of freedom) uncertainties associated with these single-epoch power spectra (Barret & Vaughan 2012). The necessary use of the maximum likelihood method is also why we consider the power spectrum of the combined FPMA and FPMB light curve instead of the co-spectrum between the two FPMs: the statistical probability distribution of the co-spectrum is still unknown (Huppenkothen & Bachetti 2021). We obtain an initial estimate of the instantaneous QPO frequency from the centroid of the corresponding Lorentzian function in our model, which we finally smooth out by simultaneously fitting a cubic spline to the *XMM-Newton* and *NuSTAR* ν_{qpo} estimates. The use of a spline is an improvement on the previous analysis, in which we used a degree-15 polynomial, as there are fewer design decisions.

4.2.3 Results

Fig. 4.1 shows the results of the QPO tracking, with each epoch in its own subfigure, and each observation within its own panel. We see that it is common for the QPO frequency to vary by at least $\sim 10\%$ over the course of an observation, with the most dramatic change seen in epoch 16*b*. The change in QPO frequency combined with the slightly different observing times between the instruments is likely responsible for the mismatch in the exact QPO shape seen in Fig. 3.2, although the difference in instrument response over the 3 – 10 keV energy band also may contribute if the QPO doesn't necessarily have constant frequency with energy.

The RMS and lag spectra are shown for the fundamental and second harmonics in Fig. 4.2, with each observation in its own column. In each case, the *NuSTAR* data are shown in blue, with the *XMM-Newton* EPIC-pn data in orange (and also green for epoch 14, which includes two *XMM-Newton* obs IDs). In all four epochs, the RMS spectrum of the first harmonic (top row) rises with energy before leveling off at $E \sim 3 - 4$ keV, as is typical for type-C QPOs, (e.g. Yan et al. 2013; Huang et al. 2018). The phase lags of the fundamental (third row) generally rise

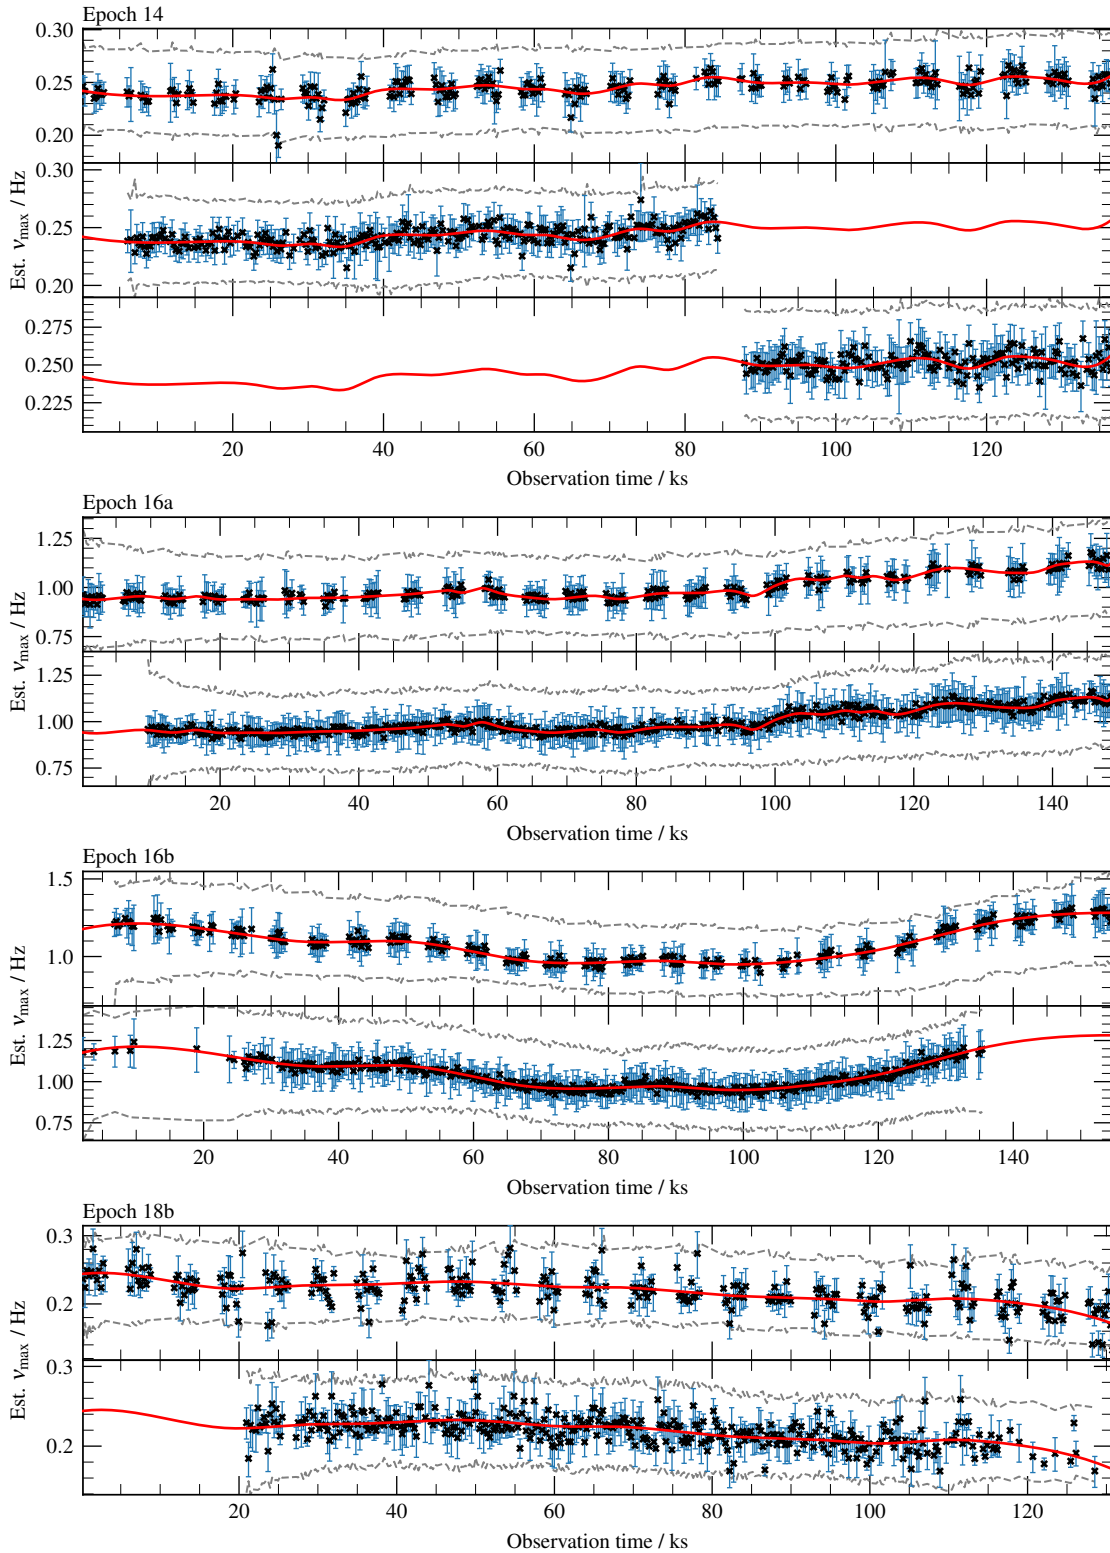


Figure 4.1: How the QPO frequency was tracked in each epoch. For each observation, the QPO frequency was estimated in short time segments, and then a cubic spline was jointly fit to the observations in each epoch. In each subfigure, the uppermost panel shows the *NuSTAR* data, whereas the remaining panels show the *XMM-Newton* data. The grey dashed lines show the running bounds for the QPO frequency. The x-axis is scaled from the start of the first of the simultaneous observations.

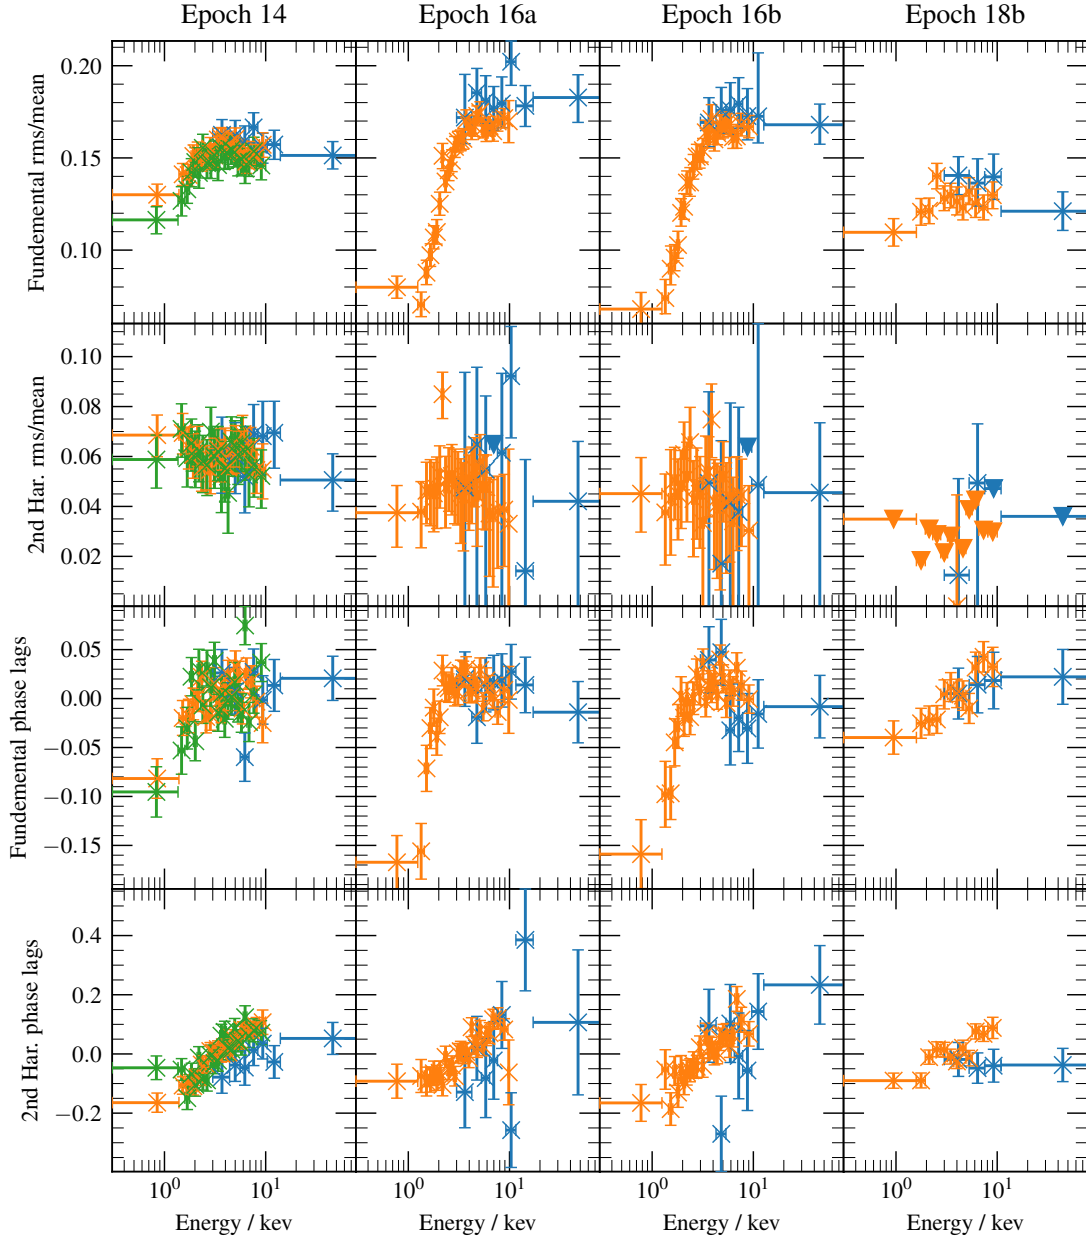


Figure 4.2: The RMS and phase lag spectra of the two harmonics, for each epoch 14, 16a, 16b, and 18b. The *NuSTAR* observation in each epoch is shown as blue markers, whereas the *XMM-Newton* observations are shown in orange (and green for the second *XMM-Newton* obsID in epoch 14). For some of the second harmonics, only an upper bound can be given for the RMS, which is shown as a triangle rather than a cross. We show these spectra for the 0.3 – 78 keV energy range, despite our further analysis only considering 2.5 – 78 keV.




Figure 4.3: The bi-coherence b^2 measured within each observation, split into the *XMM-Newton* and *NuSTAR* observations in each epoch. The red line in epoch 14 corresponds to obsID 0740980201, with the black line corresponding to 0724401901. The solid line shows the calculated $b^2(\nu)$, while the dark and light shaded regions correspond to the 1σ and 3σ confidence regions based on a bootstrapping algorithm. The purple cross shows the bi-coherence for the QPO frequency after being tracked over the the observation, with error bars showing the 1σ confidence interval on $b^2(\nu_{\text{QPO}}$, and also the frequency range which has been considered.

steeply with energy for $E \lesssim 2$ keV but rise much more gradually with energy for $E \gtrsim 2$ keV, a trend common in all the epochs but most noticeable for the 16a and 16b. Note that this corresponds to a small hard lag between bands typically used for *RXTE* data, meaning that our results are consistent with *RXTE* population studies of QPO phase lags (van den Eijnden et al. 2017), given that the QPO frequency is $\nu_{\text{qpo}} < 2$ Hz for all epochs considered here. While the second harmonic doesn't show prominent structure above the noise level in the RMS spectra, a trend of hard phase lags can be seen in all the observations (again consistent with the population study of van den Eijnden et al. 2017).

We calculate the phase difference between the two harmonics in the reference

Epoch	<i>XMM-Newton</i>	<i>NuSTAR</i>
14	$0.315 \pm 0.004^{(a)}$	0.355 ± 0.006
	$0.311 \pm 0.004^{(b)}$	
16a	0.248 ± 0.005	0.2881 ± 0.02
16b	0.251 ± 0.007	0.2842 ± 0.02
18b	0.249 ± 0.008	0.2991 ± 0.02

Table 4.1: Phase difference between QPO harmonics in the reference band in units of π radians, ψ/π . For epoch 14, the result labeled *a* and *b* are obs IDs 0724401901 and 0740980201 respectively. Differences between the *XMM-Newton* and *NuSTAR* values are expected if the phase difference is not constant with energy.

band ψ using the auto-bispectrum method. Fig. 4.3 shows the auto-bicoherence (the modulus of the bi-coherence is normalised to take values between 0 and 1, and encapsulates the strength of the coupling of the variability at two Fourier frequencies – in this case between the two QPO harmonics; see Section 2.3.4) for each epoch. The black and red stepped lines have been calculated by averaging the bi-spectrum over light curve segments without accounting for variations in the QPO frequency. The peaks in these lines at the QPO frequency indicate that the phase of the first two QPO harmonics are correlated with one another. For epoch 14, we see this correlation (which was already demonstrated by Ingram et al. 2016 using a different method) very clearly for both *XMM-Newton* and *NuSTAR*. For epochs 16a and 16b, it is very clear only in *XMM-Newton* and the data are much more noisy in epoch 18b. The purple crosses instead represent the bi-coherence of the QPO calculated by accounting for time-evolution of the QPO frequency using our frequency tracking algorithm. We see that this bi-coherence is large for all epochs except for 18b, with the improvement being particularly striking for *NuSTAR* data in epochs 16a and 16b. This indicates that the bi-coherence of the QPO itself is large, but the variation of the QPO frequency over time leads to the bi-coherence calculated for a given frequency being washed out by contamination from the broad band noise. Note that, although the QPO bi-coherence is not especially high for epoch 18b, it is non-zero at least at the 1σ level.

Since the phases of the first two QPO harmonics are correlated with one another, it makes sense to calculate the mean phase difference between the harmonics

ψ , which we present in Table 4.1 for each observation. As the phase difference between the harmonics is energy dependent, the different energy band-pass of each instrument means that light curves from *NuSTAR* and *XMM-Newton* are not expected to necessarily show the same phase difference, and we indeed see that *NuSTAR* measures a slightly higher phase difference in every epoch. We note that Ingram et al. (2016) previously measured ψ for the epoch 14 data using a slightly different method. Our results are very similar to theirs ($\psi/\pi = 0.336 \pm 0.004$, 0.347 ± 0.006 and 0.347 ± 0.005 for their *XMM-Newton* epochs 1a & 1b, their *XMM-Newton* 2a & 2b, and *NuSTAR* respectively, where our epoch 14 does not include their 1a & 1b), although not formally consistent within statistical errors. The main differences are that the new method, unlike the old one, accounts for Poisson noise; for time variability of the QPO frequency; and also weights light curve segments by variability amplitude. We therefore see the new method as more accurate, but it is encouraging that the measurement appears to be reasonably robust to these assumptions. The statistical errors quoted for both methods are smaller than the systematic errors associated with algorithm design decisions, but we do not use these error estimates for any later part of our analysis.

To demonstrate the bi-spectrum method, we show ‘jellyfish’ plots for epoch 16a as a representative example in Fig. 4.4 (we present the same plots for the other epochs in Appendix B.1). Here, we sum the bi-spectrum for each light curve segment on the complex plane. Each grey line is this vector sum for a different Fourier frequency. Segments coloured orange are ones that are identified with the QPO frequency by our frequency tracking algorithm, and the purple line is the vector sum of all of these segments that contain the QPO fundamental frequency (i.e. if the QPO frequency were constant, each purple segment would lie directly on top of an orange segment, which in turn would lie on top of a grey segment). The length of each grey line relates to the bi-coherence as a function of frequency, and the length of the purple line relates to the QPO bi-coherence represented by the purple crosses in Fig. 4.3. We see that the QPO segments line up on the complex plane much better than those associated with broad band noise frequencies (which is equivalent to the




Figure 4.4: The auto-bispectrum $B(\nu_k)$, calculated for the *XMM-Newton* (left) and *NuSTAR* (right) observations of Epoch 16a. Each grey trace in the upper panels corresponds to a frequency $\nu_k < 10$ Hz (between 6 – 10 Hz the frequency bins from the fast Fourier transform are geometrically rebinned). As the light curve is broken into short segments during the analysis (labeled m), the auto-bispectrum $B_m(\nu_k)$ is calculated for each segment. Each grey trace is composed of small lines of the relevant frequency bin of auto-bispectrum within the complex plane, each calculated from the consecutive light curve segments. As the frequency bin containing the QPO can be tracked throughout the observations, the segments of each grey frequency track are highlighted in orange at times when that frequency corresponds to the QPO. Finally, the magenta lines correspond to a ‘shifted-and-added’ bi-spectrum, comprised of the frequency bin of the QPO from each successive light curve segment. The lower panels show the auto-bicoherence b^2 , effectively the length of grey each frequency track, normalised to the power at the two relevant frequencies (ν_k and $2\nu_k$). The grey shaded regions shows the 1σ and 3σ uncertainties, estimated from bootstrapping. The magenta cross represents b^2 (and 1σ uncertainty) of the tracked QPO, along with the frequency range it has been tracked over.

peaks in the bi-coherence seen in Fig. 4.3). The phase difference ψ is simply related to the orientation of the purple line on the complex plane. For the case of *NuSTAR*, we see that accounting for QPO frequency variations is very important, since QPO segments line up well with one another, but they are spread across multiple Fourier frequencies. The result is that the QPO bi-coherence is large (purple cross in the right hand, second from top panel) even though the bi-coherence for each individual frequency has a large uncertainty due to the large Poisson noise level.

With the results thus far, we are able to see that the data quality of epoch 18b is not suitable for the fitting a phase-resolved model, with the strength of

the 2nd harmonic hard to detect above the broad-band noise (although we can see its presence in the power spectra), and only a poor energy resolution can be achieved for the fundamental harmonic RMS spectrum. As the methodology depends on the QPO harmonics being dominant over the broadband noise in order to extract the phase information, we elect not to proceed further with this epoch, and instead focus on epochs 14, 16a, and 16b.

4.3 Phase-resolved fits

4.3.1 Fitting procedure

We use the same QPO phase dependent model as we did in chapter 2, which is described in detail in Section 2.4. This model is comprised of NTHCOMP (Zdziarski et al. 1996) to represent the observed coronal continuum flux; the disc flux is locally computed at different patches of the disc with a black body plus a XILLVERCP ‘reflected’ component (García et al. 2013a), with the emission from the entire disc ray-traced to the observer’s frame. The model has two features that allow it model the QPO spectra. Firstly, it allows certain key parameters to be modulated with QPO phase as a two-harmonic variation, such as the continuum normalisation which is modulated as

$$N_c(\gamma) = N_0 + A_{1N} \sin[\gamma - \phi_{1N}] + A_{2N} \sin[2(\gamma - \phi_{2N})] \quad (4.5)$$

with N_0 , A_{1N} , A_{2N} , ϕ_{1N} , and ϕ_{2N} being free parameters of the model. The model also allows for other parameters to be modulated with QPO phase in the same way: the coronal photon index $\Gamma(\gamma)$; the reflection fraction $f_R(\gamma)$; and the coronal electron temperature T_e .

The other key element of the model is that the emissivity profile of the coronal flux onto the disc is allowed to be axi-asymmetric and QPO modulated. It is parameterised such that the emissivity at disc coordinate (r, ϕ) at QPO phase γ is

$$\epsilon(r, \phi, \gamma) = \epsilon(r) \left\{ 1 + A_1 \cos^2 \left[\frac{1}{2}(\gamma - \phi + \phi_1) \right] + A_2 \cos^2[\gamma - \phi + \phi_2] \right\} \quad (4.6)$$

where the radial dependence is a simple twice broken power-law (e.g. Wilkins & Fabian 2011, 2012). When $A_1 = A_2 = 0$, the emissivity profile axi-symmetric and is independent of the QPO phase. However, when either of these parameters are non-zero then the emissivity profile has either one or two bright patches which move around the disc with QPO phase, which will cause changes in the shape of the relativistic reflection spectrum we see. The model goes further, and also calculates the modulations in the local profiles of both the disc ionisation parameter and the disc blackbody temperature based upon the relative change in the received coronal flux. The model calculates the local blackbody temperature from two components: the intrinsic disc ‘viscous’ temperature $T_{\text{visc}}(r)$ that is independent of QPO phase and the temperature from irradiation $T_{\text{irr}}(r, \phi, \gamma)$, such that the local disc temperature is

$$T^4(r, \phi, \gamma) = T_{\text{visc}}^4(r) + T_{\text{irr}}^4(r, \phi, \gamma). \quad (4.7)$$

The relative amplitude of the blackbody temperature variations is set by the emissivity, but the QPO phase-averaged value of $T_{\text{irr}}^4(r, \phi, \gamma)$ is set by the model parameter T_i . It is therefore possible to artificially switch off temperature variations in the model by fixing $T_i = 0$. The temperature modulations are also subject to a phase delay $\Delta\gamma$ as in theory there may be a thermalisation timescale which is non-negligible on the QPO frequency.

The complete model we fit includes a multiplicative power law for cross-calibration; a TBABS column absorption component using the abundances of Wilms et al. (2000); and a XILLVERCP component representing a distant reflector. We fit this model simultaneously to the phase-averaged spectrum and the real and imaginary parts of the QPO first and second harmonic Fourier transform spectra, as described in chapter 2. We treat each epoch completely separately. We tie physical parameters between the two *XMM-Newton* obs IDs in epoch 14, only allowing the cross-calibration parameters to be free (although, as expected, we find them to be consistent with each other). Our procedure for finding a preferred set of phase-resolved spectral parameters is as follows:

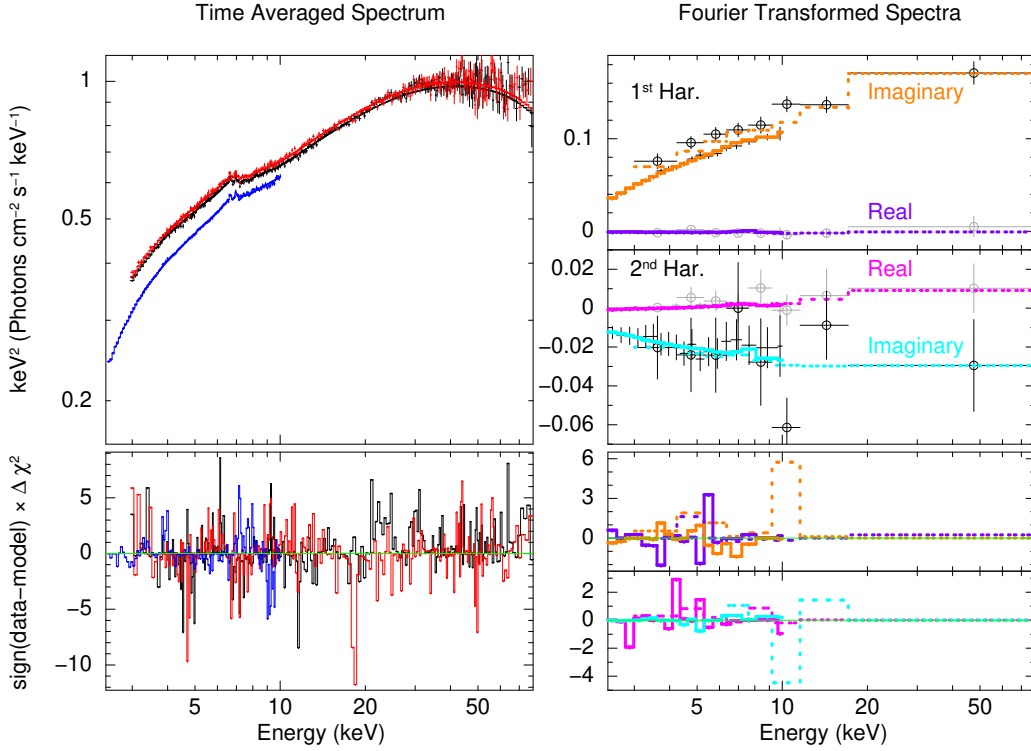


Figure 4.5: The unfolded data and best fitting model for epoch 16a, along with the χ^2 residuals. On the left is the phase-averaged spectrum, with *NuSTAR* FMPA and FMPB in red and black, and *XMM-Newton* EPIC-pn in blue. On the right are the Fourier transformed spectra, split into the first and second harmonics on the upper and lower panels. The real and imaginary components of the first harmonic of the model are in purple and orange respectively, whereas they are in pink and teal for the second harmonic. The lines are solid for the model of the *XMM-Newton* data, but dotted for the *NuSTAR* data. The markers are light grey for the real spectra, and black for the imaginary, with those of *NuSTAR* also having a circle.

- Begin by finding the global best fit spectra to the phase-averaged and QPO spectra, using the model such that all the parameters are allowed to be modulated with QPO phase: $N_c(\gamma)$, $f_R(\gamma)$, $\Gamma(\gamma)$, $T_e(\gamma)$, along with the asymmetry parameters A_1 and A_2 , and disc temperature modulation.
- Individually we switch the modulations in these parameters off for both harmonics, and find a new fit. Through the use of the statistical F-test we identify any modulations which are not required to the 1σ significance level. The F-test is required as by setting the amplitude of the modulations to 0, we

also remove the sensitivity of the model to the associated phase parameters.

- Of the parameters found to be $< 1\sigma$ significant, we turn off the least significant parameter modulation. We then repeat this until the remaining parameter modulations are all required to at least the 1σ significance level.

4.3.1.1 Results

After following the procedure described above, we find that the epochs require the following parameters to be modulated:

- Epoch 14: $f_R(\gamma)$ at 3.0σ significance; $N_c(\gamma)$ at $\gg 5\sigma$ significance; with χ^2 fit statistic 3351 / 3449 d.o.f.
- Epoch 16a: Asymmetry parameters at 1.1σ significance; $f_R(\gamma)$ at 4.2σ significance; T_i at 2.8σ significance; $N_c(\gamma)$ at $\gg 5\sigma$ significance; with χ^2 fit statistic 3261 / 3144 d.o.f.
- Epoch 16b: $\Gamma(\gamma)$ at 1.2σ significance; $N_c(\gamma)$ at $\gg 5\sigma$ significance; with χ^2 fit statistic 3176 / 3112 d.o.f.

Fig 4.5 demonstrates our fits by presenting the unfolded phase-averaged spectrum (left) and QPO harmonic spectra (right) for the example of epoch 16a. The best fitting parameters for all three epochs are presented in Table 4.2. Fig. 4.6 demonstrates the best-fitting parameter modulations with QPO phase for each epoch. They are presented as probability maps to provide an illustrative view of the statistical errors, although note that the errors from different modulation parameters were combined using a simplified prescription (see figure caption for details). When a parameter is held constant with QPO phase in the final fit, we simply plot the best fitting tied value of that parameter with associated uncertainty. We see that the continuum normalisation (bottom) is modulated for all three epochs, which is a trivial result given the very strong QPO present in each. The photon index (top) is modulated only for epoch 16b, but for a low significance. The reflection fraction (middle) is modulated in 14 and 16a with respectively 3.0σ and 4.0σ

significance. Ingram et al. (2017) also found the reflection fraction to be modulated (3.5σ significance) in a data set that included our epoch 14. The modulation we constrain here is compatible with that seen by Ingram et al. (2017), with a minimum at ~ 0.2 cycles and maxima at ~ 0.5 and ~ 1 QPO cycles. The mean value is different because here we employ a slightly different definition of reflection fraction (the observer's reflection fraction here and the system reflection fraction in the earlier study; see Ingram et al. 2019). We further explore the reflection fraction modulation by presenting 2D contour plots of the A_{1f} , A_{2f} plane for epochs 14 and 16a in Fig. 4.7 (left and right respectively).

The asymmetry parameters A_1 and A_2 are non-zero in our final fits only for epoch 16a, with our f-test procedure (described in the previous sub-section) preferring to freeze $A_1 = A_2 = 0$ in the final fits for epochs 14 and 16b. In contrast, Ingram et al. (2017) were able to rule out $A_1 = A_2 = 0$ with 2.4σ confidence for a larger data set that included our epoch 14. The difference here is that Ingram et al. (2017) considered more data than we do here (stacking the data from our epoch 14 with another *XMM-Newton* observation), plus they also employed a more restrictive model such that our more advanced model has more freedom to reproduce the data when the asymmetry parameters are held to zero. Even though our constraints are slightly worse, our results are compatible with those of Ingram et al. (2017). We test this by freeing up A_1 , A_2 , ϕ_1 and ϕ_2 in our best fitting epoch 14 model, finding a new fit and then calculating 1σ error bounds on the four newly free parameters. The new best fit, as expected, has a lower χ^2 of 3747, but for 4 extra degrees of freedom this is a $< 1\sigma$ improvement on our preferred fit. This procedure gives a 1σ confidence interval of $A_1 > 1$ (however the upper bound is unconstrained) and $\phi_1 = 0.56^{+0.11}_{-0.07}$ cycles. This A_1 value is compatible with the measurement from Ingram et al. (2017) of $A_1 = 0.93^{+1.3}_{-0.6}$, whereas our ϕ_1 range lies slightly above their 1σ range of $\phi_1 = 0.26^{+0.12}_{-0.13}$. We find that A_2 , and therefore ϕ_2 , are unconstrained so are therefore compatible with the previous $A_2 = 3.5^{+5.2}_{-2.2}$ of Ingram et al. (2017).

Variable	Epoch 14	Epoch 16a	Epoch 16b
FPMB / FMPA	1.044 ± 0.001	1.021 ± 0.001	1.0216 ± 0.001
XMM1 / FMPA	0.761 ± 0.002	0.785 ± 0.003	0.777 ± 0.004
XMM1 $\Delta\Gamma$	$0.100^{+0.002}_{-0.001}$	$0.081^{+0.001}_{-0.002}$	$0.087^{+0.002}_{-0.003}$
XMM1 ϕ_c / cyc	0.000 ± 0.002	-0.001 ± 0.001	0.001 ± 0.002
XMM2 / FMPA	$0.763^{+0.003}_{-0.001}$	—	—
XMM2 $\Delta\Gamma$	0.105 ± 0.002	—	—
XMM2 ϕ_c / cyc	-0.001 ± 0.002	—	—
N_{H} / 10^{22}	$3.7^{+0.1}_{-0.2}$	$2.9^{+0.3}_{-0.1}$	$3.2^{+0.2}_{-0.4}$
Incl. / deg	$79.7^{+0.2}_{-0.4}$	69 ± 2	$81.6^{+1.1}_{-0.3}$
r_{in} / R_{g}	8^{+3}_{-2}	27^{+4}_{-2}	17 ± 2
q_1	2^{+4}_{-2}	> 12.6	6 ± 1
q_2	$0.2^{+1.6}_{-0.2}$	$-1.3^{+0.8}_{-0.7}$	0 ± 2
$r_{\text{br},1}$ / r_{in}	2.4^*	2.4 ± 0.2	$3.5^{+0.6}_{-0.3}$
$r_{\text{br},2}$ / $r_{\text{br},1}$	1.92^*	< 14.5	> 3
kT_{e} / keV	60^{+2}_{-5}	37^{+2}_{-1}	39^{+3}_{-5}
$A_{1\Gamma}$	—	—	< 0.005
$A_{2\Gamma}$	—	—	0.02 ± 0.01
$\Gamma(\gamma)$ $\phi_{1\Gamma}$ / cyc	—	—	0^*
$\phi_{2\Gamma}$ / cyc	—	—	$0.23^{+0.01}_{-0.06}$
Γ	$1.689^{+0.004}_{-0.012}$	$1.750^{+0.006}_{-0.002}$	$1.729^{+0.006}_{-0.002}$
A_{1f}	0.08 ± 0.05	$0.17^{+0.17}_{-0.04}$	—
A_{2f}	0.20 ± 0.06	$0.22^{+0.13}_{-0.07}$	—
$f_{\text{R}}(\gamma)$ ϕ_{1f} / cyc	$0.43^{+0.05}_{-0.10}$	0.96 ± 0.02	—
ϕ_{2f} / cyc	$0.94^{+0.02}_{-0.03}$	0.97 ± 0.02	—
f_{R}	$1.16^{+0.07}_{-0.09}$	$0.30^{+0.03}_{-0.05}$	$0.8^{+0.3}_{-0.1}$
A_1	—	> 0.42	—
A_2	—	> 0.85	—
Asym. Φ_1 / cyc	—	0.53 ± 0.08	—
Φ_2 / cyc	—	$0.14^{+0.03}_{-0.04}$	—
$\log \xi$	$3.4^{+0.8}_{-1.7}$	2.8 ± 0.2	$4.2^{+0.3}_{-0.5}$
A_{Fe}	< 0.54	$0.75^{+0.05}_{-0.11}$	2.3 ± 0.4
$kT_{\text{v,max}}$ / keV	$0.484^{+0.032}_{-0.006}$	$0.54^{+0.04}_{-0.03}$	0.61 ± 0.04
kT_{i} / keV	—	$0.57^{+0.05}_{-0.07}$	—
$\Delta\gamma$ / cyc	—	$-0.01^{+0.01}_{-0.02}$	—
N_{d} / 10^{-3}	13 ± 5	$0.6^{+0.4}_{-0.2}$	3^{+3}_{-1}
A_{1N}	0.84 ± 0.1	$0.75^{+0.02}_{-0.04}$	0.76 ± 0.02
A_{2N}	0.326 ± 0.012	$0.32^{+0.06}_{-0.01}$	0.11 ± 0.06
$N(\gamma)$ ϕ_{1N} / cyc	$0.001^{+0.001}_{-0.002}$	$0.012^{+0.005}_{-0.002}$	-0.002 ± 0.002
ϕ_{2N} / cyc	0.280 ± 0.002	$0.244^{+0.004}_{-0.003}$	$0.28^{+0.03}_{-0.02}$
N_{c}	$3.65^{+0.07}_{-0.09}$	$3.17^{+0.06}_{-0.02}$	$2.49^{+0.08}_{-0.22}$
$\log \xi^{\dagger}$	$2.80^{+0.05}_{-0.07}$	$3.32^{+0.02}_{-0.03}$	$4.33^{+0.05}_{-0.07}$
norm^{\dagger} / 10^{-3}	$1.3^{+0.1}_{-0.2}$	$1.4^{+0.2}_{-0.1}$	$6.8^{+2.0}_{-0.5}$

Table 4.2: The best fitting and 1σ errors of the free parameters in the model. The first 7 rows show the cross-calibration parameters (with the normalisation of the different instruments given in reference to the normalisation of FMPA). The parameters with no entry were not applicable in the final fit. Some parameters only had one bound, which is shown as an equality for that bound. Others were unconstrained as indicated by the best fit with a *. The ionisation and normalisation of the distant reflection are indicated by \dagger .




Figure 4.6: Probability heat maps which demonstrate the QPO phase dependence on the photon index $\Gamma(\gamma)$, the reflection fraction $f_R(\gamma)$, and the coronal power law normalisation $N_c(\gamma)$; for the three epochs which we fit to the phase resolved model. These were created by producing a sample of each of the five parameters for each waveform (e.g. N_c , A_{1N} , P_{1N} , A_{2N} , P_{2N}), randomly selecting from a normal distribution based upon the parameter best fitting value, and the estimated error. From each these sample parameters a waveform was produced, to allow the heat map made. This is only for visualisation, and doesn't take into account for the actual posterior distributions being non-Gaussian, or any covariance between the parameters, which is actually likely. The model does not allow for $f_R(\gamma) < 0$, so the waveforms are cut-off there. Finally, for the parameters which are considered to be non-significantly modulated (Γ for epochs 14 and 16a, f_R for epoch 16b), the heatmap is produced for reference, representing the uncertainty on the time-averaged parameter only.




Figure 4.7: A contour plot from χ^2 space exploration of the reflection fraction modulation amplitudes A_{1f} and A_{2f} for epochs 14 (left) and 16a (right). The contours are shown corresponding to 1 – 5 σ and 1 – 6 σ confidence levels, for a χ^2 distribution with 2 degrees of freedom.

4.4 Discussion and Conclusions

We have analysed the type-C QPOs in long look *XMM-Newton* and *NuSTAR* observations of the BH XRB H1743-322 taken across four observational epochs during three outbursts. The range in QPO frequency ($\sim 0.2 - 1$) Hz provides the opportunity to study how the accretion flow parameters in this object evolve with QPO frequency. We find that the data are of high enough quality to conduct QPO phase-resolved spectroscopy for three of the four epochs. We use the phase-resolving method and apply the phase-resolved spectral model that are both described in chapter 2 of this thesis.

We find a statistically significant modulation of the reflection fraction in epochs 14 (3.0σ) and 16a (4.2σ). The former is in agreement with the results of Ingram et al. (2017), who analysed a larger data set containing our epoch 14, and the latter adds another detection of this effect to the original Ingram et al. (2017) detection and the modulation in GRS 1915+105 presented in chapter 2 of this thesis. These formative studies therefore imply that the reflection fraction is commonly modulated with QPO phase. This adds further to the mounting body of evidence that type-C QPOs are a geometric effect, particularly when considered in the context of the inclination dependence of QPO properties (Motta et al. 2015; van den Eijnden

et al. 2017). While we do not see a significant reflection fraction modulation in epoch 16b, we do see a modulation in the photon index $\Gamma(\gamma)$, albeit with low significance. When the distinctive features of the reflection spectrum are hard to detect in the QPO Fourier transformed spectra, a modulation in the photon index is able to produce a similar effect to the reflection fraction varying: the reflection spectrum is harder than the continuum, and so a varying power-law index produces similar spectral variability to a varying reflection fraction if the signal to noise is poor enough for variations in the iron line feature to be undetectable. Indeed, the modulations of $\Gamma(\gamma)$ in epoch 16b, as visualised in Fig. 4.6, are closely in anti-phase with the modulations in $f_{\text{R}}(\gamma)$ seen in the other epochs; suggesting that hardness variations with QPO phase in the epoch 16b data that may in reality be due to an increasing reflection fraction are instead modelled as a decreasing photon index. By considering the uncertainty in the phase-averaged reflection fraction of epoch 16b it is clear that $f_{\text{R}}(\gamma)$ modulations with a similar amplitude to those in the other epochs would be undetectable. We also note the same anti-phase between the reflection fraction $f_{\text{R}}(\gamma)$ and coronal normalisation $N_{\text{c}}(\gamma)$ in epochs 14 and 16a, as is seen for GRS 1915+105 in chapter 2 of this thesis.

Our model includes heating of the disc via irradiation with a phase lag parameter, $\Delta\gamma$, introduced to account for the time delay resulting from the time it takes irradiating photons to reach thermal equilibrium in the disc. For GRS 1915+105, the best-fitting value of this parameter implied an uncomfortably high thermalisation time of ~ 75 ms. Here, we are able to constrain such temperature variations with high statistical significance for epoch 16a only. The 1σ confidence range of the thermalisation phase lag is $\Delta\gamma = -0.033$ to 0.002 cycles, which corresponds to a time lag of -33 to 2 ms. A negative lag is incompatible with the interpretation of $\Delta\gamma$ as a thermalisation lag, but the 1σ upper bound of 2 ms is very reasonable in comparison to theory (Garcia et al in prep).

We are unable to rule out the null hypothesis of axisymmetric disc illumination ($A_1 = A_2 = 0$) with high statistical confidence for any of the epochs studied. The significance we find here for epoch 14 is lower than that found for a similar epoch

by Ingram et al. (2017). This is at least partly because Ingram et al. (2017) stacked more data together, and thus achieved a higher signal to noise, than we do here. Here we are also using a more complex and therefore more flexible model. We find that our weak constraints on the asymmetry parameters are compatible with the previous result within 1σ confidence. Since we are currently limited by data quality, in future we will consider epochs 16a and 16b (which correspond to consecutive *XMM-Newton* orbits around Earth) simultaneously to boost signal to noise as Ingram et al. (2017) did for the 2014 outburst.

The values we find for the disc inner radius r_{in} and the inclination are generally inconsistent with those found during Chapter 3, and are inconsistent with the general finding that r_{in} decreases as the QPO frequency increases (e.g. García et al. 2015). We also find that the best-fitting inclination angle is inconsistent with remaining constant across the epochs as would be expected. While these inconsistencies are seemingly worrying, there are potential modelling reasons which could cause this result. It may be that the inner radius and inclination are anti-correlated in our fits, and that either tying the inclination across all epochs or imposing a prior on it could yield more believable r_{in} values. Any correlations between inner radius and inclination can also be explored in future work using a Markov Chain Monte Carlo simulation, which would provide a deeper search of the parameter space and potentially a more accurate, and larger, statistical uncertainty. Conversely, these issues could be due to systemics from comparing similar yet different models. For example, the discrepancy in the r_{in} values measured here and in the previous chapter could be due to non-linear biases introduced by parameter modulations being averaged over the full QPO cycle. Such biases are accounted for in these phase-resolved spectral fits, but not in fits to the time-averaged spectrum only as were presented in chapter 3.

When these observations were taken, the intention was to probe how QPO properties varied for different QPO frequencies within the same source. However, this work shows that the signal requirements for advanced phase-resolved spectroscopy require stacking significant amounts of data, especially when using *XMM-Newton*

which cannot look at the brightest sources in the sky without experiencing issues due to pile-up. In future, we plan to apply our QPO phase-resolving technique to very high count rate observations taken by *NICER*, which is not constrained by pile-up. A highly luminous source such as MAXI J1820+070 (Thomas et al. 2022) or MAXI J1535-571 (Zhang et al. 2022) could finally provide the signal to noise necessary for a firm detection of asymmetric disc illumination.

*I've searched around the universe
Been down some black holes
There's nothing but space, man*

— Sam Ryder

5

Conclusions

In this thesis I have presented the work which I undertook over the last four years, focusing on spectral modelling of accreting black holes in X-ray binary systems, and in particular looking at the reflection signal. All of the work within this thesis used models built upon the state-of-the-art XILLVER reflection model of García & Kallman (2010).

In Chapter 3 I used the model introduced by Ingram et al. (2022) to make one of the first mass measurements of a black hole using only the X-ray spectrum. By using the shape of the observed reflection spectrum to infer the X-ray flux incident on (and therefore reflected by) the disc and comparing this to the observed flux, this technique is able to constrain mass-distance ratio of the BH. Therefore this method enables the measurement of black hole masses in systems where current mass measurement techniques cannot be applied (provided separate methods can provide a distance). This technique is able to probe the mass population of black holes in our Galaxy more fully, and therefore may resolve current observational biases towards lower black hole masses (Jonker et al. 2021). The work I presented was a proof-of-principle test of the method, and I found a mass consistent with mass ranges expected for black holes in X-ray binaries. However, I only considered a system which currently has no mass estimate, so any mass value could be reasonable, provided it is within an appropriate range. The truest test of this technique will be

to consider a system where the mass has already been dynamically measured! This will be the focus of future work, on X-ray binaries such as, for example, Cygnus X-1. Including information of the short timescale variability within the observation in the analysis will also provide a more stringent test of the model, since the timing properties can also be used to infer a mass that may or may not be compatible with the value I inferred from the spectrum.

A large part of this thesis involved not just studying the reflection spectrum from accreting black holes, but also probing how the reflection spectrum varies over the course of each cycle of quasi-periodic oscillations (QPOs) seen in the light curves, which is presented in Chapters 2 and 4. This work built upon the tomographic reflection model and the advanced timing techniques introduced by Ingram et al. (2016) and Ingram et al. (2017). I introduced a new QPO frequency tracking algorithm, based upon fitting the power spectra of short light curves using maximum likelihood estimation. I also introduced the bispectrum to contain the phase difference between QPO harmonics, which improves upon previous methods as it enables a proper Poisson noise correction. Together these innovations make the method more robust, and enable it to be applied to a greater number of observations. By understanding QPOs we can shed light on the innermost accretion flow in these strong gravity conditions, and the interplay between the disc and the corona. In Chapter 2 I analysed an observation of GRS 1915+105, and I was able to find strong evidence of the presence of a QPO phase dependent illumination cycle, along with a phase modulated reflection fraction, which adds to the growing body of evidence supporting geometrical models for QPOs. However, in Chapter 4 I had more limited success when studying observations of H1743-332, although I did again find good evidence of the reflection fraction phase modulation. My work highlights the need for high signal/noise observations for these techniques to be successful: while Ingram et al. (2017) had been able to constrain the phase dependent illumination profile, when using only a subset of the same data I was not able to constrain the illumination profile in the same way. This shows that while stacking the data can provide constraining results, it also restricts the ability to consider changes

in the QPO properties over this timescale (or the techniques could fail as it has to be assumed that the properties we're looking for are constant). One of the original goals of the observation campaign was to analyse how the QPO differed at different frequencies for the same source, so it is unfortunate that the data were not quite good enough to be used in this manner.

Reflection modelling of these highly relativistic systems is tricky to do, especially when working with the most advanced models. These models make assumptions about the nature of the reflection, such as the ionisation structure of the disc, or the illumination profile. While much work goes into understanding these choices to select the most appropriate option, these choices are always going to impart systematic effects on the parameters which are determined from the modelling. Further to this, the reflection model parameters often are at least partially degenerate, which makes exploring the parameter space tricky, and filled with many local minima. When close to the global best fit, Markov Chain Monte Carlo (MCMC) is a very good way to fully understand the parameter space and the statistical uncertainties. I used MCMC for the work in Chapters 2 and 3, however these algorithms take a long time to run. The sheer number of parameters in the reflection models means that many samples being required for the MCMC to converge. Since the advanced disc reflection models I have used are relatively computationally expensive, some of the MCMC analysis I have performed has taken months to complete. However, this could be improved in the future by moving away from more general routines (such as the MCMC functionality within XSPEC) to more specialised routines that account for the behaviour of the parameters in a manner that enables them to converge with fewer model evaluations.

During the production of this thesis, almost all the data I analysed came from simultaneous observations of *NuSTAR* alongside either *XMM-Newton* or *NICER*. These are instruments which compliment each other well, with *NuSTAR* providing coverage at high energy, while *XMM-Newton* and *NICER* cover the lower energies. However, despite having a very similar bandpass to *XMM-Newton*, *NICER* is able to look at brighter sources without being affected by pileup effects,

helping provide the signal requirements especially for timing analysis. I hope that as the catalogue of *NICER* observations grow, there will be increasingly more observations of QPOs suitable for phase resolving, therefore the dream of understanding their underlying mechanism will be soon realised. However, the future of X-ray astronomy is even brighter than just *NICER*. While the last year of my DPhil coincided with the curtailment of the upcoming *ATHENA* mission, and the interruption of *eROSITA* due to international politics, there are many missions being designed which will revolutionise the field of X-ray timing and spectroscopy, with many rich discoveries surely waiting.

One such mission has recently been launched, the Imaging X-ray Polarimetry Explorer (*IXPE*; Weisskopf et al. 2016), the first X-ray polarimeter to be launched in decades. As of the submission of this thesis, *IXPE* is already making very exciting discoveries, such as the higher than expected polarisation of Cygnus X-1 which gives insights into the geometry of the accretion flow (Krawczynski et al. 2022). With further, even more powerful X-ray polarimetry missions already on the horizon it is clear that the models and techniques I have used in this thesis will be able to be applied to X-ray polarisation (e.g. Ingram & Maccarone 2017), particularly in the case of joint observations between polarimeters and upcoming (and existing) pure spectrometry missions. Now my DPhil has concluded, I will be going forward to join the *NuSTAR* team at CalTech, where I will be working to expand the XILLVER reflection models to include polarisation, which I hope will be used in the analysis of future observations, and help push our understanding of black hole accretion into a new age.

Appendices

*There was always a logical explanation for everything,
even if you had to make it up.*

— Terry Pratchett, *Soul Music*

A

Appendix of Chapter 2

Contents

A.1 Phase difference from Bispectrum	147
A.1.1 Treatment of Poisson noise in the Bispectrum	147
A.1.2 Jellyfish Bootstrapping	149
A.2 Model Details	149
A.2.1 Ray tracing	149
A.2.2 Model Normalisation	150
A.3 Phase offsets with different reference bands	150
A.4 MCMC corner plots of modulated parameters	152

A.1 Phase difference from Bispectrum

A.1.1 Treatment of Poisson noise in the Bispectrum

The ‘jellyfish plots’ in Fig. 2.5 are corrected for Poisson noise and deadtime effects. Failing to correct for these effects introduces an instrumental component into the real part of the bispectrum, which presents itself as a general ‘drift’ of the random walks along the positive real axis. In turn, this reduces the measured bi-phase from the true value.

There exists a Poisson-noise free estimate of the bispectrum, as discussed in

Wirnitzer (1985), which we use in the form

$$\begin{aligned}
 B_m(\nu) &= R_m(\nu)R_m(\nu)R_m^*(2\nu) \\
 &\quad - 2|R_m(\nu)|^2 - |R_m(2\nu)|^2 \\
 &\quad + 2N,
 \end{aligned} \tag{A.1}$$

where $R_m(\nu)$ are the Fourier transforms of segments of the lightcurve, and $N = \langle R_m(0) \rangle$ is the average number of photons per light curve segment¹. Likewise, we want to use a Poisson-free estimate for the normalisation of $b^2(\nu)$. Considering the expectation of random variables X representing the Fourier transform of a series containing a signal with power P_s , but also noise P_n

$$\mathbb{E} [|X X|^2] = 2(P_s + P_n)^2, \tag{A.2}$$

and clearly in the noiseless case it is $2P_s^2$. While we could use a rearrangement of Eq. A.2, we instead opt to simply use a noise-subtracted estimate of P_s from the powerspectrum. Therefore, we also replace the denominator of Eq. 2.10 with

$$2 \left(\langle |R_m(\nu)|^2 \rangle - N \right)^2 \left(\langle |R_m(2\nu)|^2 \rangle - N \right). \tag{A.3}$$

However, the deadtime effects of *NuSTAR* are still more complex than the base Poisson noise. Therefore we take inspiration from Bachetti et al. (2015) and use the co-bispectrum between light curve segments from FPMA and FPMB ($R_{A,m}(\nu)$ and $R_{B,m}(\nu)$ respectively)

$$B_m(\nu) = R_{A,m}(\nu)X_{A,m}(\nu)R_{B,m}^*(2\nu), \tag{A.4}$$

and also normalise b^2 using the denominator

$$\mathbb{R} [\langle R_{A,m}(\nu)R_{B,m}(\nu) \rangle]^2 \mathbb{R} [\langle R_{A,m}(2\nu)R_{B,m}(2\nu) \rangle] \frac{\langle R_{A,m}(0) \rangle}{\langle R_{B,m}(0) \rangle}. \tag{A.5}$$

¹As here for normalisation consistency, we consider light curve segments measured in photon counts, as opposed to the count-rate.

A.1.2 Jellyfish Bootstrapping

We use a bootstrapping method to calculate the uncertainties in the bispectrum for our determination of the phase-difference between the QPO harmonics.

Each observation has already been split up into M segments, so we draw from this (with replacement) a sample of M segments, 1000 times². Therefore, each of these 1000 random draws is the same ‘length’ as the original observation. We then calculate the bispectrum for each of these samples. When performing the QPO tracking, we ensure to use the QPO frequency corresponding to the segment’s true time.

As b^2 is restricted to be positive³, we show the range corresponding to the $\pm 1\sigma$ and $\pm 3\sigma$ quantiles in Fig. 2.5.

To ensure we don’t fall foul of any issues relating to the cyclic nature of the biphasic, to calculate the error on the QPO phase-difference we take the measurement from each 1000-segment sample. For these values, we phase-wrap them so each lies within π rad of the true measurement and then report the standard deviation on these values as our uncertainty.

A.2 Model Details

A.2.1 Ray tracing

A disc patch at coordinate (r, ϕ) with radial and azimuthal extent dr and $d\phi$ will subtend a solid angle $d\Omega = d\alpha d\beta / D^2$ on the image plane, and will be centered at horizontal and vertical coordinates on the image plane α and β . Here, D is the distance from observer to source, and α and β are the impact parameters at infinity, where the singularity occupies the position $\alpha = \beta = 0$ on the image plane. The total QPO phase-dependent specific flux observed from the disc is (e.g. Ingram et al. 2019)

$$F_d(E, \gamma) = \int_0^{2\pi} \int_{r_{\text{in}}}^{r_{\text{out}}} g^3(r, \phi) \frac{I(E/g, r, \phi, \gamma)}{D^2} d\alpha d\beta, \quad (\text{A.6})$$

²As an example, if $M = 5$, with segments labelled ‘A’, ‘B’, ‘C’, ‘D’, and ‘E’ we would use 1000 random draws such as ‘ABAAD’, ‘EECDC’, and so on.

³However in the region of low signal/noise, the Poisson correction we make can push an estimate below 0, see Appendix A.1.2 for details.

where $g = E/E_d$ is the energy shift experienced by a photon travelling from disc coordinate (r, ϕ) to the observer (given by Equation 4 in Ingram et al. 2017) and E is photon energy in the observer restframe.

A.2.2 Model Normalisation

We normalise the XILLVERCP spectrum, $\mathcal{R}(E)$, such that the incident spectrum that goes into the calculation has an integral over all energies of unity. Due to the internal XILLVERCP normalisation (see Eq. 16 of Ingram et al. 2019), this gives

$$\int_0^{2\pi} \int_0^1 \int_0^\infty \mu_e \mathcal{R}(E) dE d\mu_e d\phi = 1, \quad (\text{A.7})$$

where μ_e is the cosine of the emission angle. We normalise the emissivity function such that

$$\int_{\alpha, \beta} g^4(r, \phi) \epsilon(r, \phi, \gamma) d\alpha d\beta = 1. \quad (\text{A.8})$$

These two conditions together ensure that the observed bolometric reflected flux (in the case of $N_d = 0$) would be $f_R(\gamma) N_c(\gamma)$ if the function $\mathcal{R}(E)$ was independent of emission angle.

A.3 Phase offsets with different reference bands

Recalling Eq. 2.3, the FT of the j^{th} QPO harmonic is

$$W_j(E) = \mu(E) \sigma_j(E) e^{i\Phi_j(E)}, \quad (\text{A.9})$$

where (recalling Eq. 2.2)

$$\begin{aligned} \Phi_1(E) &= \Phi_1 + \Delta_1(E) \\ \Phi_2(E) &= 2(\Phi_1 + \psi) + \Delta_2(E). \end{aligned} \quad (\text{A.10})$$

Here, denoting the subject band as S_j and reference band as R_j , the phase lags are (taking the argument of Eq. 2.5)

$$\Delta_j(E) = \arg [S_j(E) R_j^*]. \quad (\text{A.11})$$

When we measure the same QPO FT but with a different instrument, and therefore with a different reference band, we consider a different reference band T , where T_j lags behind R_j by a phase difference δ_j . The QPO FT we measure from our new instrument is therefore

$$Q_j^T(E) = \mu^T(E) \sigma_j(E) e^{\Phi_j^T(E)}, \quad (\text{A.12})$$

where $\Phi_j^T(E)$ is

$$\begin{aligned} \Phi_1^T(E) &= \Phi_1 + \Delta_1(E) - \delta_1 \\ \Phi_2^T(E) &= 2(\Phi_1 + \psi^T) + \Delta_2(E) - \delta_2, \end{aligned} \quad (\text{A.13})$$

and where ψ^T is the phase difference between harmonics in the T band, which in principle can be different from ψ as the reference bands come from instruments with different energy bands, and even have different responses within overlapping bands.

Combining Eqs. A.10 and A.13, we find that

$$\begin{aligned} \Phi_1^T(E) &= \Phi_1(E) - \delta_1 \\ \Phi_2^T(E) &= \Phi_2(E) + 2(\psi^T - \psi) - \delta_2, \end{aligned} \quad (\text{A.14})$$

Considering the phase difference between the harmonics in any energy band (Ingram et al. 2016)

$$\psi(E) = \psi - \Delta_1(E) + \frac{1}{2}\Delta_2(E). \quad (\text{A.15})$$

we can similarly construct that

$$\psi^T = \psi - \delta_1 + \frac{1}{2}\delta_2, \quad (\text{A.16})$$

and therefore

$$\begin{aligned} \Phi_1^T(E) &= \Phi_1(E) - \delta_1 \\ \Phi_2^T(E) &= \Phi_2(E) - 2\delta_1. \end{aligned} \quad (\text{A.17})$$

Putting this together, the relation between the QPO FT when measured with the different reference bands is

$$Q_j^T(E) = Q_j(E) \exp(-ij\delta_1), \quad (\text{A.18})$$

and so the phase offset of the j^{th} harmonic is

$$\phi_{c,j} = -j\delta_1, \quad (\text{A.19})$$

so finally we have

$$\begin{aligned} \phi_{c,1} &\equiv \phi_c \\ \phi_{c,2} &= 2\phi_c. \end{aligned} \quad (\text{A.20})$$

A.4 MCMC corner plots of modulated parameters

This appendix contains corner plots of some of the parameters from the MCMC. Fig. A.1 shows the phase average, and the 1st harmonic and 2nd harmonic amplitudes of the modulated parameters compared with each other. Fig. A.2 shows the 1st and 2nd harmonic amplitudes for each of the modulated parameters, from which we can see how significantly away from $(0,0)$ the parameter is, which would correspond to that parameter not being modulated.

For the majority of the parameters, the posterior is mostly symmetric and approximately Gaussian. However, this is not true for some of the radial emissivity parameters q_1 , q_2 , r_{in} , $r_{\text{br},1}$, and $r_{\text{br},2}$ likely due to their interdependence (see Eq. 2.15). For completeness they are shown as a corner plot in Fig. A.3; all parameters whose posterior distributions are not otherwise shown in this paper are included as histograms in Fig. A.4 with symbols matching those in Table 2.2.

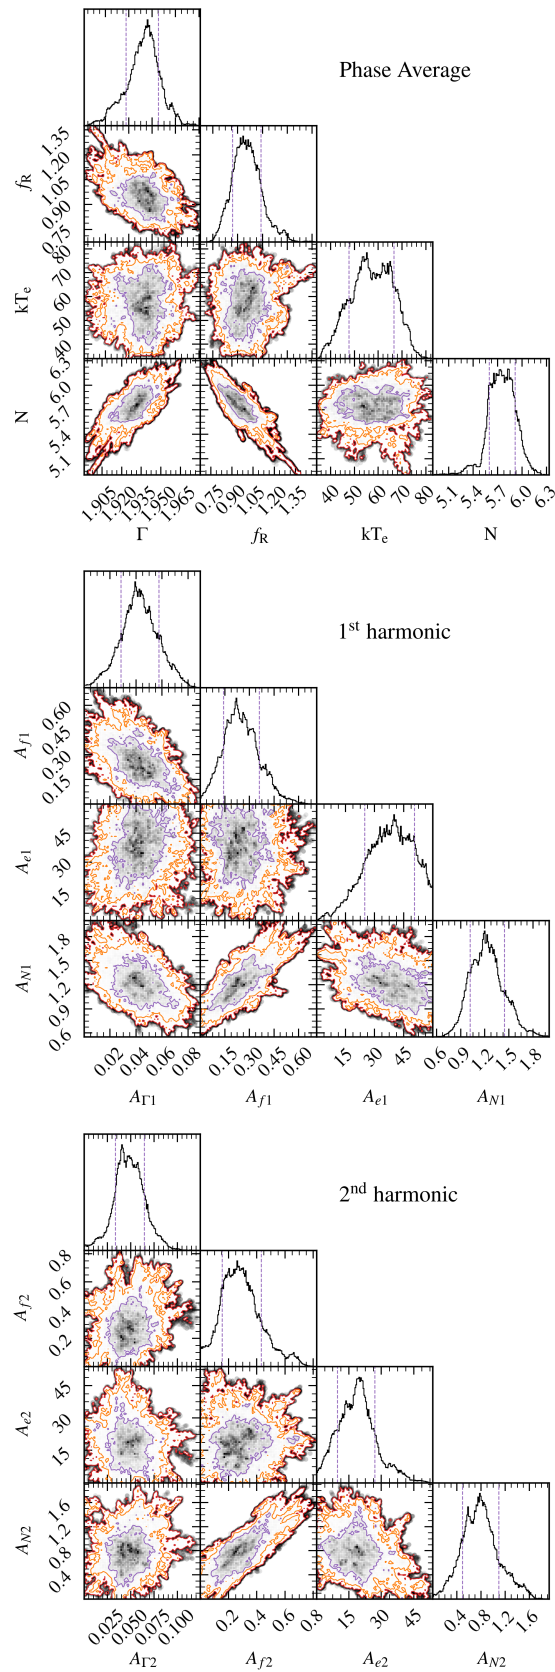


Figure A.1: Corner plots from the MCMC of the modulated parameters within the model. The purple, orange, and red lines show the 1, 2, 3 σ credible intervals respectively. The phase-average, and also the amplitudes of the first and second harmonics are shown separately.




Figure A.2: Corner plots from the MCMC of the amplitudes of the modulations within the model. The purple, orange, and red lines show the 1, 2, 3 σ credible intervals respectively. In all cases, the (0, 0) lies outside of the 3 σ contour showing that all four parameters are consistent with being modulated.




Figure A.3: Corner plots from the MCMC of the radial emissivity parameters. The purple, orange, and red lines show the 1, 2, 3 σ credible intervals respectively. It is important to note that the two break radii $r_{\text{br},1}$ and $r_{\text{br},2}$ have units of the disc truncation radius r_{in} and $r_{\text{br},1}$ respectively. As q_1 and q_2 are the power law indices within their relevant break radii, the model becomes insensitive to them when $r_{\text{br},1}$ and $r_{\text{br},2}$ approach 1 respectively.




Figure A.4: The posteriors of all the parameters not otherwise shown in this paper, with the $\pm 1\sigma$ credible interval shown by the dashed purple vertical lines. For visual simplicity the parameter values are not shown here, as the shape of the distribution is the key feature. The parameters marked with \dagger relate solely to the distant reflector component of the model.

You can use logic to justify almost anything. That's its power and its flaw.

— Captain Janeway, Star Trek: Voyager

B

Appendix of Chapter 4

Contents

B.1 Auto-bispectrum for epochs 14, 16b, and 18b	157
--	------------

B.1 Auto-bispectrum for epochs 14, 16b, and 18b

This appendix contains the ‘jellyfish’ plots for epochs 14 in Fig. B.1, 16b in Fig. B.2, and 18b in Fig. B.3. Those for epoch 16a were left in the main body as indicative example (Fig. 4.4). The bi-coherence corresponding to each of these is contained within Fig. 4.3, while the phase difference between the harmonics that these represent are all listed in Table 4.1. See section 4.2.3 for details explaining how they were calculated.

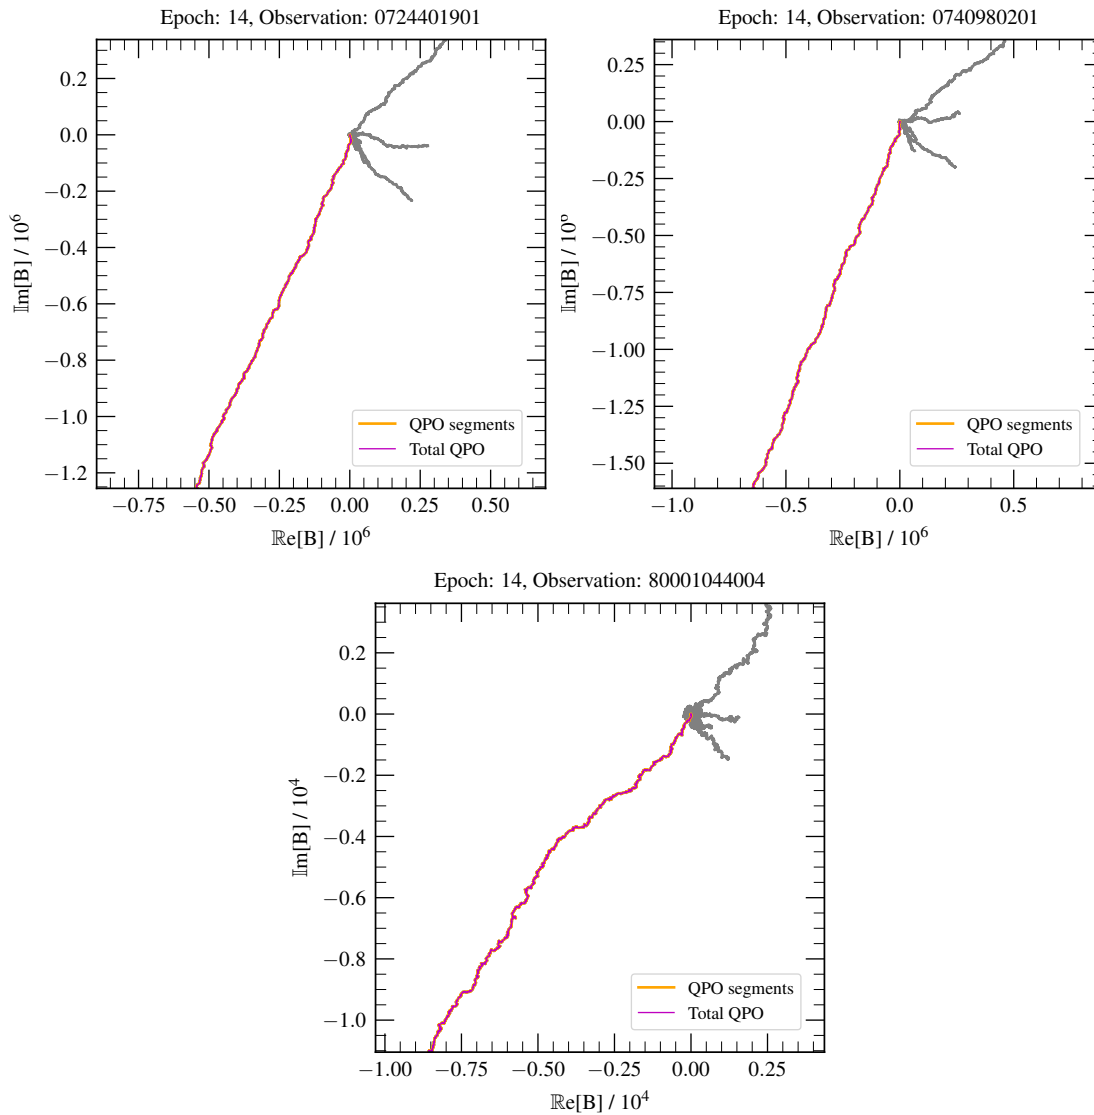


Figure B.1: The auto-bispectrum for epoch 14, with the *XMM-Newton* observations on the top, and the *NuSTAR* observation below. The grey lines show $B(\nu_k)$ for frequencies $\nu_k < 10$ Hz, with orange segments showing the frequency bins containing the QPO at any time, which are summed into the purple line (covering the orange segments).

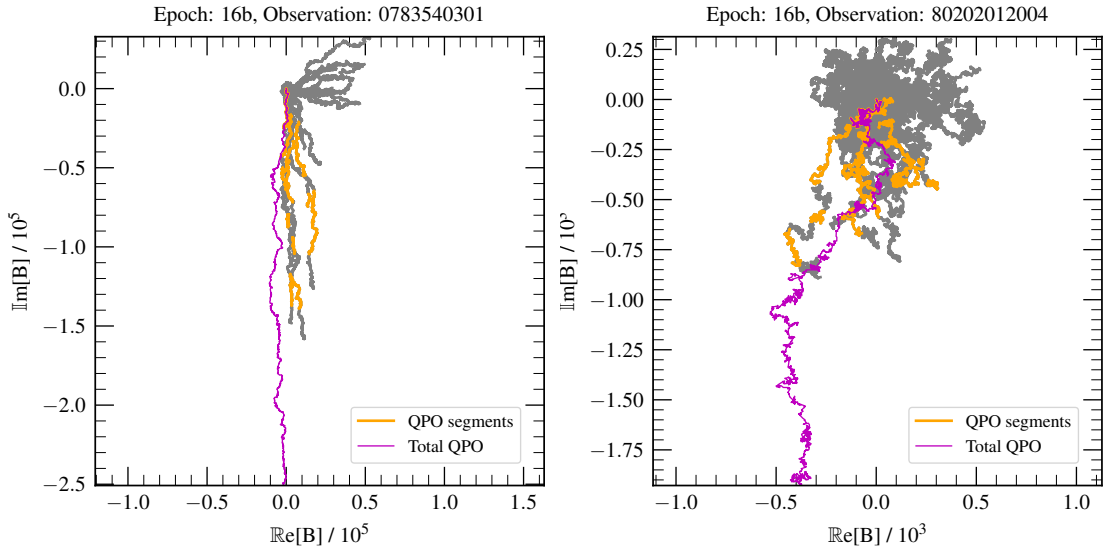


Figure B.2: The auto-bispectrum for epoch 16b, with the *XMM-Newton* observation on the left, and the *NuSTAR* observation on the right. The grey lines show $B(\nu_k)$ for frequencies $\nu_k < 10$ Hz, with orange segments showing the frequency bins containing the QPO at any time, which are summed into the purple line.

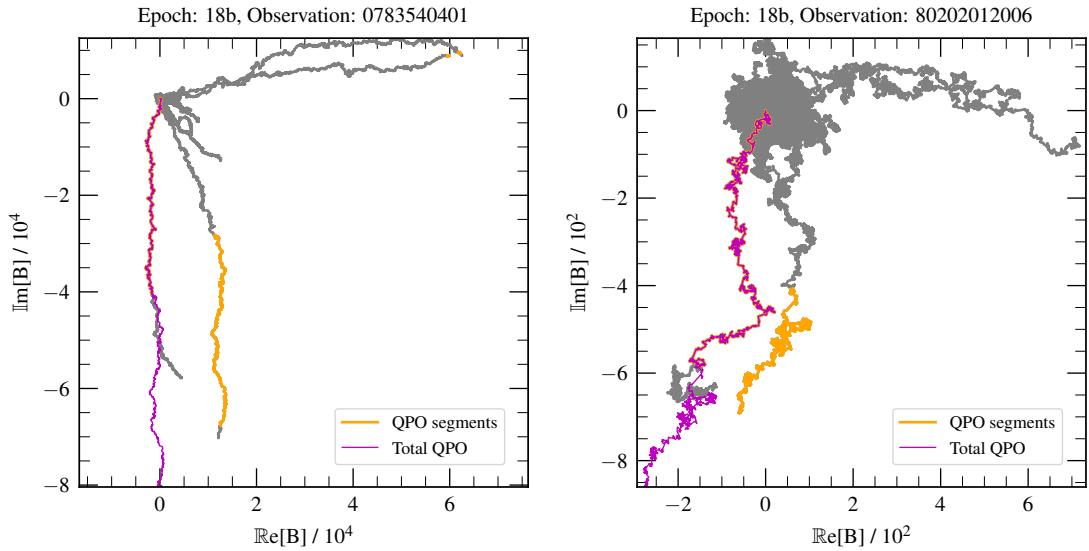


Figure B.3: The auto-bispectrum for epoch 18b, with the *XMM-Newton* observation on the left, and the *NuSTAR* observation on the right. The grey lines show $B(\nu_k)$ for frequencies $\nu_k < 10$ Hz, with orange segments showing the frequency bins containing the QPO at any time, which are summed into the purple line.

The relevant equation is: Knowledge = power = energy = matter = mass; a good bookshop is just a genteel Black Hole that knows how to read.

— Terry Pratchett, Guards! Guards!

Bibliography

- Abbott R., et al., 2021, ApJ, 913, L7
- Alabarta K., et al., 2021, MNRAS, 507, 5507
- Altamirano D., et al., 2011, ApJ, 742, L17
- Arévalo P., Uttley P., 2006, MNRAS, 367, 801
- Arur K., Maccarone T., 2019, Monthly Notices of the Royal Astronomical Society, 486, 3451
- Arur K., Maccarone T. J., 2022, MNRAS, 514, 1720
- Axelsson M., Done C., 2016, MNRAS, 458, 1778
- Axelsson M., Done C., Hjalmarsdotter L., 2014, MNRAS, 438, 657
- Bachetti M., et al., 2014, Nature, 514, 202
- Bachetti M., et al., 2015, ApJ, 800, 109
- Bambi C., 2017, X-Ray Reflection Spectroscopy. Springer Singapore, Singapore, pp 153–179, doi:10.1007/978-981-10-4524-0_8, https://doi.org/10.1007/978-981-10-4524-0_8
- Bardeen J. M., Petterson J. A., 1975, The Astrophysical Journal, 195, L65
- Bardeen J. M., Press W. H., Teukolsky S. A., 1972, ApJ, 178, 347
- Barret D., Vaughan S., 2012, ApJ, 746, 131
- Basak R., Zdziarski A. A., Parker M., Islam N., 2017, MNRAS, 472, 4220
- Belloni T. M., 2010, in Belloni T., ed., , Vol. 794, Lecture Notes in Physics, Berlin Springer Verlag. Springer Berlin Heidelberg, pp 53–84, doi:10.1007/978-3-540-76937-8_3
- Belloni T., Hasinger G., 1990, A&A, 230, 103
- Belloni T. M., Motta S. E., 2016, in Bambi C., ed., Astrophysics and Space Science Library Vol. 440, Astrophysics of Black Holes: From Fundamental Aspects to Latest Developments. p. 61 (arXiv:1603.07872), doi:10.1007/978-3-662-52859-4_2
- Belloni T., Klein-Wolt M., Méndez M., van der Klis M., van Paradijs J., 2000, A&A, 355, 271
- Blaauw A., 1961, Bull. Astron. Inst. Netherlands, 15, 265
- Bult P., 2017, ApJ, 837, 61

- Cabanac C., Henri G., Petrucci P. O., Malzac J., Ferreira J., Belloni T. M., 2010, *MNRAS*, 404, 738
- Cannizzo J. K., 1993, *ApJ*, 419, 318
- Cannizzo J. K., Reiff C. M., 1992, *ApJ*, 385, 87
- Casares J., Jonker P. G., 2014, *Space Sci. Rev.*, 183, 223
- Chainakun P., Young A. J., 2017, *MNRAS*, 465, 3965
- Chakrabarti S. K., 1997, *ApJ*, 484, 313
- Connors R. M. T., et al., 2021, *ApJ*, 909, 146
- Corral-Santana J., 2021, BlackCAT a catalog of stellar-mass black holes in X-ray binaries, <https://www.astro.puc.cl/BlackCAT/>
- Corral-Santana J. M., Casares J., Muñoz-Darias T., Bauer F. E., Martínez-Pais I. G., Russell D. M., 2016, *A&A*, 587, A61
- Dexter J., Agol E., 2009, *ApJ*, 696, 1616
- Done C., 2010, arXiv e-prints, p. arXiv:1008.2287
- Done C., Nayakshin S., 2007, *Monthly Notices of the Royal Astronomical Society: Letters*, 377, L59
- Done C., Gierliński M., Kubota A., 2007, *A&ARv*, 15, 1
- Dove J. B., Wilms J., Maisack M., Begelman M. C., 1997, *ApJ*, 487, 759
- Eardley D. M., Lightman A. P., Shapiro S. L., 1975, *The Astrophysical Journal*, 199, L153
- Fabian A. C., Rees M. J., Stella L., White N. E., 1989, *MNRAS*, 238, 729
- Fabian A. C., Iwasawa K., Reynolds C. S., Young A. J., 2000, *PASP*, 112, 1145
- Farr W. M., Sravan N., Cantrell A., Kreidberg L., Bailyn C. D., Mandel I., Kalogera V., 2011, *ApJ*, 741, 103
- Fender R., Muñoz-Darias T., 2016, in Haardt F., Gorini V., Moschella U., Treves A., Colpi M., eds, , Vol. 905, *Lecture Notes in Physics*. Springer International Publishing, p. 65, doi:10.1007/978-3-319-19416-5_3
- Fender R. P., Garrington S., McKay D., Muxlow T., Pooley G. G., Spencer R., Stirling A., Waltman E., 1999, *Monthly Notices of the Royal Astronomical Society*, 304, 865
- Fender R. P., Belloni T. M., Gallo E., 2004, *Monthly Notices of the Royal Astronomical Society*, 355, 1105
- Fender R. P., Homan J., Belloni T. M., 2009, *MNRAS*, 396, 1370
- Field G. B., Rogers R. D., 1993, *ApJ*, 403, 94
- Fragile P. C., Blaes O. M., Anninos P., Salmonson J. D., 2007, *The Astrophysical Journal*, 668, 417

- Fragile P. C., Straub O., Blaes O., 2016, *MNRAS*, 461, 1356
- Fragos T., Tremmel M., Rantsiou E., Belczynski K., 2010, *The Astrophysical Journal Letters*, 719, L79
- Frank J., King A., Raine D., et al., 2002, *Accretion power in astrophysics*. Cambridge university press
- Fryer C. L., Kalogera V., 2001, *ApJ*, 554, 548
- Fryer C. L., Belczynski K., Wiktorowicz G., Dominik M., Kalogera V., Holz D. E., 2012, *ApJ*, 749, 91
- Fürst F., et al., 2016, *ApJ*, 828, 34
- Gallo E., Fender R. P., Pooley G. G., 2003, *MNRAS*, 344, 60
- García J., Kallman T. R., 2010, *ApJ*, 718, 695
- García J., Dauser T., Reynolds C. S., Kallman T. R., McClintock J. E., Wilms J., Eikmann W., 2013a, *ApJ*, 768, 146
- García J., Elhoussieny E. E., Bautista M. A., Kallman T. R., 2013b, *ApJ*, 775, 8
- García J. A., Steiner J. F., McClintock J. E., Remillard R. A., Grinberg V., Dauser T., 2015, *ApJ*, 813, 84
- García J. A., Fabian A. C., Kallman T. R., Dauser T., Parker M. L., McClintock J. E., Steiner J. F., Wilms J., 2016, *MNRAS*, 462, 751
- García J. A., et al., 2019, *ApJ*, 871, 88
- Gendreau K. C., et al., 2016, in den Herder J.-W. A., Takahashi T., Bautz M., eds, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*. p. 99051H, doi:10.1117/12.2231304
- Geweke J., 1992, *Bayesian statistics*, 4, 641
- Greene J. E., Strader J., Ho L. C., 2020, *ARA&A*, 58, 257
- Guilbert P. W., Rees M. J., 1988, *MNRAS*, 233, 475
- HEASARC N., 2014, *Astrophysics Source Code Library ascl*, 1408
- Haardt F., Maraschi L., 1991, *ApJ*, 380, L51
- Haardt F., Maraschi L., 1993, *ApJ*, 413, 507
- Haardt F., Maraschi L., Ghisellini G., 1994, *ApJ*, 432, L95
- Harrison F. A., et al., 2013, *The Astrophysical Journal*, 770, 103
- Heil L. M., Vaughan S., Uttley P., 2011, *MNRAS*, 411, L66
- Heil L. M., Vaughan S., Uttley P., 2012, *MNRAS*, 422, 2620
- Heil L. M., Uttley P., Klein-Wolt M., 2015, *MNRAS*, 448, 3348
- Homan J., et al., 2020, *ApJ*, 891, L29

- Huang Y., et al., 2018, *ApJ*, 866, 122
- Huppenkothen D., Bachetti M., 2018, *The Astrophysical Journal Supplement Series*, 236, 13
- Huppenkothen D., Bachetti M., 2021, *MNRAS*,
- Huppenkothen D., Heil L. M., Hogg D. W., Mueller A., 2017, *Monthly Notices of the Royal Astronomical Society*, 466, 2364
- Hussein E. M., 2007, in Hussein E. M., ed., , *Radiation Mechanics*. Elsevier Science Ltd, Oxford, pp 1–65, doi:<https://doi.org/10.1016/B978-008045053-7/50002-1>, <https://www.sciencedirect.com/science/article/pii/B9780080450537500021>
- Ichimaru S., 1977, *ApJ*, 214, 840
- Ingram A., 2019, *MNRAS*,
- Ingram A., Done C., 2012, *Monthly Notices of the Royal Astronomical Society*, 427, 934
- Ingram A. R., Maccarone T. J., 2017, *MNRAS*, 471, 4206
- Ingram A., Motta S., 2014, *MNRAS*, 444, 2065
- Ingram A. R., Motta S. E., 2019, *New Astronomy Reviews*, 85, 101524
- Ingram A., van der Klis M., 2013, *MNRAS*, 434, 1476
- Ingram A., van der Klis M., 2015, *MNRAS*, 446, 3516
- Ingram A., Done C., Fragile P. C., 2009, *MNRAS*, 397, L101
- Ingram A., Maccarone T. J., Poutanen J., Krawczynski H., 2015, *The Astrophysical Journal*, 807, 53
- Ingram A., van der Klis M., Middleton M., Done C., Altamirano D., Heil L., Uttley P., Axelsson M., 2016, *MNRAS*, 461, 1967
- Ingram A., van der Klis M., Middleton M., Altamirano D., Uttley P., 2017, *MNRAS*, 464, 2979
- Ingram A., Mastroserio G., Dauser T., Hovenkamp P., van der Klis M., García J. A., 2019, *MNRAS*, 488, 324
- Ingram A., et al., 2022, *MNRAS*, 509, 619
- Jansen F., et al., 2001, *A&A*, 365, L1
- Jiang J., Fabian A. C., Wang J., Walton D. J., García J. A., Parker M. L., Steiner J. F., Tomsick J. A., 2019a, *MNRAS*, 484, 1972
- Jiang J., et al., 2019b, *MNRAS*, 489, 3436
- Jonker P. G., Kaur K., Stone N., Torres M. A. P., 2021, *ApJ*, 921, 131
- Kalamkar M., Reynolds M., van der Klis M., Altamirano D., Miller J., 2015, *The Astrophysical Journal*, 802, 23
- Kallman T., Bautista M., 2001, *ApJS*, 133, 221

- Kara E., et al., 2019, *Nature*, 565, 198
- Karpouzas K., Méndez M., Ribeiro E. M., Altamirano D., Blaes O., García F., 2020, *MNRAS*, 492, 1399
- Karpouzas K., Méndez M., García F., Zhang L., Altamirano D., Belloni T., Zhang Y., 2021, *MNRAS*, 503, 5522
- Kato S., ed. 1998, *Black-hole accretion disks*
- Kato S., 2001, *Publications of the Astronomical Society of Japan*, 53, 1
- Kato S., Fukue J., 1980, *PASJ*, 32, 377
- Kerr R. P., 1963, *Phys. Rev. Lett.*, 11, 237
- Kim Y. C., Powers E. J., 1979, *IEEE transactions on plasma science*, 7, 120
- King A., Nixon C., 2016, *MNRAS*, 462, 464
- King A. R., Ritter H., 1998, *MNRAS*, 293, L42
- Kolehmainen M., Done C., Díaz Trigo M., 2014, *MNRAS*, 437, 316
- Kovach C. K., Oya H., Kawasaki H., 2018, *Neuroimage*, 173, 518
- Krawczynski H., et al., 2022, *arXiv e-prints*, p. arXiv:2206.09972
- Kumar N., Misra R., 2014, *MNRAS*, 445, 2818
- Kylafis N. D., Belloni T. M., 2015, *A&A*, 574, A133
- Kylafis N. D., Reig P., Papadakis I., 2020, *A&A*, 640, L16
- Laor A., 1991, *ApJ*, 376, 90
- Lasota J.-P., 2001, *New Astron. Rev.*, 45, 449
- Lense J., Thirring H., 1918, *Physikalische Zeitschrift*, 19, 156
- Lewin W. H. G., van Paradijs J., van der Klis M., 1988, *Space Sci. Rev.*, 46, 273
- Lightman A. P., White T. R., 1988, *ApJ*, 335, 57
- Liska M., Hesp C., Tchekhovskoy A., Ingram A., van der Klis M., Markoff S., 2018, *MNRAS*, 474, L81
- Liska M., Tchekhovskoy A., Ingram A., van der Klis M., 2019, *MNRAS*, 487, 550
- Liska M., Hesp C., Tchekhovskoy A., Ingram A., van der Klis M., Markoff S. B., Van Moer M., 2021, *MNRAS*, 507, 983
- Longair M. S., 2011, *High Energy Astrophysics*. Cambridge University Press
- Lubow S. H., Ogilvie G. I., Pringle J. E., 2002, *MNRAS*, 337, 706
- Lyubarskii Y. E., 1997, *MNRAS*, 292, 679
- Maccarone T. J., 2013, *MNRAS*, 435, 3547

- Maccarone T. J., Uttley P., van der Klis M., Wijnands R. A. D., Coppi P. S., 2011, *MNRAS*, 413, 1819
- Machida M., Matsumoto R., 2008, *PASJ*, 60, 613
- Madsen K. K., Forster K., Grefenstette B. W., Harrison F. A., Stern D., 2017, *ApJ*, 841, 56
- Madsen K. K., Grefenstette B. W., Pike S., Miyasaka H., Brightman M., Forster K., Harrison F. A., 2020, arXiv preprint arXiv:2005.00569
- Mahmoud R. D., Done C., 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 4040
- Mahmoud R. D., Done C., De Marco B., 2019, *Monthly Notices of the Royal Astronomical Society*, 486, 2137
- Malzac J., Dumont A. M., Mouchet M., 2005, *A&A*, 430, 761
- Markoff S., Nowak M. A., Wilms J., 2005, *The Astrophysical Journal*, 635, 1203
- Martocchia A., Matt G., 1996, *MNRAS*, 282, L53
- Mastroserio G., Ingram A., van der Klis M., 2019, *MNRAS*, 488, 348
- Mastroserio G., et al., 2021, *MNRAS*, 507, 55
- Matt G., Perola G. C., Piro L., 1991, *A&A*, 247, 25
- Méndez M., et al., 1998, *ApJ*, 494, L65
- Middleton M., 2016, in Bambi C., ed., *Astrophysics and Space Science Library* Vol. 440, *Astrophysics of Black Holes: From Fundamental Aspects to Latest Developments*. p. 99 (arXiv:1507.06153), doi:10.1007/978-3-662-52859-4_3
- Middleton M. J., Sutton A. D., Roberts T. P., Jackson F. E., Done C., 2012, *MNRAS*, 420, 2969
- Miller-Jones J. C., et al., 2019, *Nature*, 569, 374
- Miller-Jones J. C. A., et al., 2021, *Science*, 371, 1046
- Miller J. M., et al., 2013, *ApJ*, 775, L45
- Mills B. S., Davis S. W., Middleton M. J., 2021, arXiv preprint arXiv:2101.11655
- Miyamoto S., Kitamoto S., 1991, *The Astrophysical Journal*, 374, 741
- Motta S. E., Belloni T. M., Stella L., Muñoz-Darias T., Fender R., 2014a, *MNRAS*, 437, 2554
- Motta S., Muñoz-Darias T., Sanna A., Fender R., Belloni T., Stella L., 2014b, *Monthly Notices of the Royal Astronomical Society: Letters*, 439, L65
- Motta S. E., Casella P., Henze M., Muñoz-Darias T., Sanna A., Fender R., Belloni T., 2015, *Monthly Notices of the Royal Astronomical Society*, 447, 2059
- Motta S., et al., 2021, arXiv preprint arXiv:2101.01187
- Mummery A., Balbus S. A., 2019a, *MNRAS*, 489, 132

- Mummery A., Balbus S. A., 2019b, *MNRAS*, 489, 143
- NASA 2021, NICER Data Analysis Threads - NICER Responses (ARFs and RMFs), https://heasarc.gsfc.nasa.gov/docs/nicer/analysis_threads/arf-rmf/
- Narayan R., Yi I., 1995, *ApJ*, 452, 710
- Nathan E., et al., 2022, *MNRAS*, 511, 255
- Newman E. T., Couch E., Chinnapared K., Exton A., Prakash A., Torrence R., 1965, *Journal of Mathematical Physics*, 6, 918
- Nixon C. J., King A. R., 2012, *Monthly Notices of the Royal Astronomical Society*, 421, 1201
- Novikov I. D., Thorne K. S., 1973, in *Black Holes (Les Astres Occlus)*. pp 343–450
- Nowak M., Lehr D., 1998, in *Abramowicz M. A., Björnsson G., Pringle J. E., eds, Theory of black hole accretion discs*. Cambridge University Press, Cambridge, Chapt. 15, pp 233–253
- O’Neill S. M., Reynolds C. S., Miller M. C., Sorathia K. A., 2011, *ApJ*, 736, 107
- Orosz J. A., McClintock J. E., Aufdenberg J. P., Remillard R. A., Reid M. J., Narayan R., Gou L., 2011, *ApJ*, 742, 84
- Özel F., Psaltis D., Narayan R., McClintock J. E., 2010, *ApJ*, 725, 1918
- Parker M. L., et al., 2016, *ApJ*, 821, L6
- Peterson B. M., et al., 2004, *ApJ*, 613, 682
- Plant D., Fender R., Ponti G., Muñoz-Darias T., Coriat M., 2014, *Monthly Notices of the Royal Astronomical Society*, 442, 1767
- Poutanen J., Coppi P. S., 1998, *Physica Scripta Volume T*, 77, 57
- Psaltis D., Norman C., 2000, arXiv e-prints, pp astro-ph/0001391
- Qu J., Lu F., Lu Y., Song L., Zhang S., Ding G., Wang J., 2010, *The Astrophysical Journal*, 710, 836
- Rees M. J., Begelman M. C., Blandford R. D., Phinney E. S., 1982, *Nature*, 295, 17
- Reid M., McClintock J. E., Steiner J., Steeghs D., Remillard R., Dhawan V., Narayan R., 2014, *The Astrophysical Journal*, 796, 2
- Reig P., Belloni T., van der Klis M., Méndez M., Kylafis N., Ford E., 2000, *The Astrophysical Journal*, 541, 883
- Reis R. C., Fabian A. C., Ross R. R., Miniutti G., Miller J. M., Reynolds C., 2008, *MNRAS*, 387, 1489
- Remillard R. A., et al., 2021, arXiv preprint arXiv:2105.09901
- Reynolds C. S., 2021, *ARA&A*, 59, 117
- Ross R. R., Fabian A. C., 1993, *MNRAS*, 261, 74

- Ross R., Fabian A., 2005, *Monthly Notices of the Royal Astronomical Society*, 358, 211
- Russell T. D., et al., 2019, *ApJ*, 883, 198
- Rybicki G. B., Lightman A. P., 1979, *Radiative processes in astrophysics*. WILEY-VCH
- Schnittman J. D., Homan J., Miller J. M., 2006, *ApJ*, 642, 420
- Schwarzschild K., 1916, in *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*. pp 424–434
- Shakura N. I., Sunyaev R. A., 1973, in Bradt H., Giacconi R., eds, *IAU Symposium Vol. 55, X- and Gamma-Ray Astronomy*. p. 155
- Shreeram S., Ingram A., 2020, *Monthly Notices of the Royal Astronomical Society*, 492, 405
- Sobolewska M. A., Życki P. T., 2006, *MNRAS*, 370, 405
- Steeghs D., McClintock J. E., Parsons S. G., Reid M., Littlefair S., Dhillon V. S., 2013, *The Astrophysical Journal*, 768, 185
- Steiner J. F., McClintock J. E., Reid M. J., 2012, *ApJ*, 745, L7
- Stella L., Vietri M., 1998, *ApJ*, 492, L59
- Stella L., Vietri M., Morsink S. M., 1999, *ApJ*, 524, L63
- Stern B. E., Poutanen J., Svensson R., Sikora M., Begelman M. C., 1995, *ApJ*, 449, L13
- Stevens A. L., Uttley P., 2016, *MNRAS*, 460, 2796
- Sunyaev R. A., Truemper J., 1979, *Nature*, 279, 506
- Tagger M., 2007, in *Revista Mexicana de Astronomia y Astrofisica*, vol. 27. pp 26–35
- Tagger M., Pellat R., 1999, *A&A*, 349, 1003
- Tagger M., Varnière P., Rodriguez J., Pellat R., 2004, *ApJ*, 607, 410
- Tarter C. B., Tucker W. H., Salpeter E. E., 1969, *ApJ*, 156, 943
- Thomas J. K., et al., 2022, *MNRAS*, 513, L35
- Tomsick J. A., et al., 2018, *ApJ*, 855, 3
- Tsang D., Butsky I., 2013, *MNRAS*, 435, 749
- Tsang D., Lai D., 2009, *MNRAS*, 393, 992
- Uttley P., McHardy I. M., 2001, *MNRAS*, 323, L26
- Uttley P., Wilkinson T., Cassatella P., Wilms J., Pottschmidt K., Hanke M., Böck M., 2011, *MNRAS*, 414, L60
- Uttley P., Cackett E., Fabian A., Kara E., Wilkins D., 2014, *The Astronomy and Astrophysics Review*, 22, 1

- Van Rossum G., Drake F. L., 2009, Python 3 Reference Manual. CreateSpace, Scotts Valley, CA
- Varnière P., Tagger M., 2002, *A&A*, 394, 329
- Veledina A., Poutanen J., Ingram A., 2013, *ApJ*, 778, 165
- Verbunt F., Igoshev A., Cator E., 2017, *A&A*, 608, A57
- Vignarca F., Migliari S., Belloni T., Psaltis D., van der Klis M., 2003, *A&A*, 397, 729
- Villanueva-Domingo P., Mena O., Palomares-Ruiz S., 2021, *Frontiers in Astronomy and Space Sciences*, 8, 87
- Wang Z.-Y., Huang C.-Y., Wang D.-X., Wang J.-Z., 2012, *Research in Astronomy and Astrophysics*, 12, 661
- Wang Y.-Z., Tang S.-P., Liang Y.-F., Han M.-Z., Li X., Jin Z.-P., Fan Y.-Z., Wei D.-M., 2021, *ApJ*, 913, 42
- Webster B. L., Murdin P., 1972, *Nature*, 235, 37
- Weisskopf M. C., et al., 2016, in den Herder J.-W. A., Takahashi T., Bautz M., eds, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*. p. 990517, doi:10.1117/12.2235240
- Wilkins D., Fabian A., 2011, *Monthly Notices of the Royal Astronomical Society*, 414, 1269
- Wilkins D., Fabian A., 2012, *Monthly Notices of the Royal Astronomical Society*, 424, 1284
- Wilkinson T., Uttley P., 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 666
- Wilms J., Allen A., McCray R., 2000, *ApJ*, 542, 914
- Wirnitzer B., 1985, *JOSA A*, 2, 14
- Yamada S., Makishima K., Done C., Torii S., Noda H., Sakurai S., 2013, *PASJ*, 65, 80
- Yan S.-P., Ding G.-Q., Wang N., Qu J.-L., Song L.-M., 2013, *Monthly Notices of the Royal Astronomical Society*, 434, 59
- Yang X., Wang J., 2013, *The Astrophysical Journal Supplement Series*, 207, 6
- You B., Życki P. T., Ingram A., Bursa M., Wang W., 2020, *The Astrophysical Journal*, 897, 27
- Young A. J., Fabian A. C., Ross R. R., Tanaka Y., 2001, *MNRAS*, 325, 1045
- Zdziarski A. A., Johnson W. N., Magdziarz P., 1996, *MNRAS*, 283, 193
- Zdziarski A. A., De Marco B., Szanecki M., Niedźwiecki A., Markowitz A., 2021a, *The Astrophysical Journal*, 906, 69
- Zdziarski A. A., Dzielak M. A., De Marco B., Szanecki M., Niedźwiecki A., 2021b, *ApJ*, 909, L9

- Zdziarski A. A., You B., Szanecki M., Li X.-B., Ge M., 2022, *ApJ*, 928, 11
- Zhang W., Jahoda K., Swank J. H., Morgan E. H., Giles A. B., 1995, *ApJ*, 449, 930
- Zhang Y., Abdikamalov A. B., Ayzenberg D., Bambi C., Dauser T., García J. A., Nampalliwar S., 2019, *The Astrophysical Journal*, 875, 41
- Zhang L., et al., 2020, *Monthly Notices of the Royal Astronomical Society*, 494, 1375
- Zhang Y., et al., 2022, *MNRAS*, 512, 2686
- de Avellar M. G. B., Méndez M., Sanna A., Horvath J. E., 2013, *MNRAS*, 433, 3453
- de Ruiter I., van den Eijnden J., Ingram A., Uttley P., 2019, *MNRAS*, 485, 3834
- van Straaten S., van der Klis M., di Salvo T., Belloni T., 2002, *ApJ*, 568, 912
- van den Eijnden J., Ingram A., Uttley P., 2016, *Monthly Notices of the Royal Astronomical Society*, 458, 3655
- van den Eijnden J., Ingram A., Uttley P., Motta S. E., Belloni T. M., Gardenier D. W., 2017, *MNRAS*, 464, 2643
- van der Klis M., 1989, in Ögelman H., van den Heuvel E. P. J., eds, *NATO Advanced Science Institutes (ASI) Series C Vol. 262*, NATO Advanced Science Institutes (ASI) Series C. p. 27
- van der Klis M., Hasinger G., Stella L., Langmeier A., van Paradijs J., Lewin W. H. G., 1987, *ApJ*, 319, L13
- van der Sluys M., 2006, Roche Potential, <https://commons.wikimedia.org/wiki/File:RochePotential.jpg>