New techniques for understanding rapid X-ray variability from compact objects

Stevens, A.L.

New techniques for understanding rapid X-ray variability from compact objects

Abigail L. Stevens

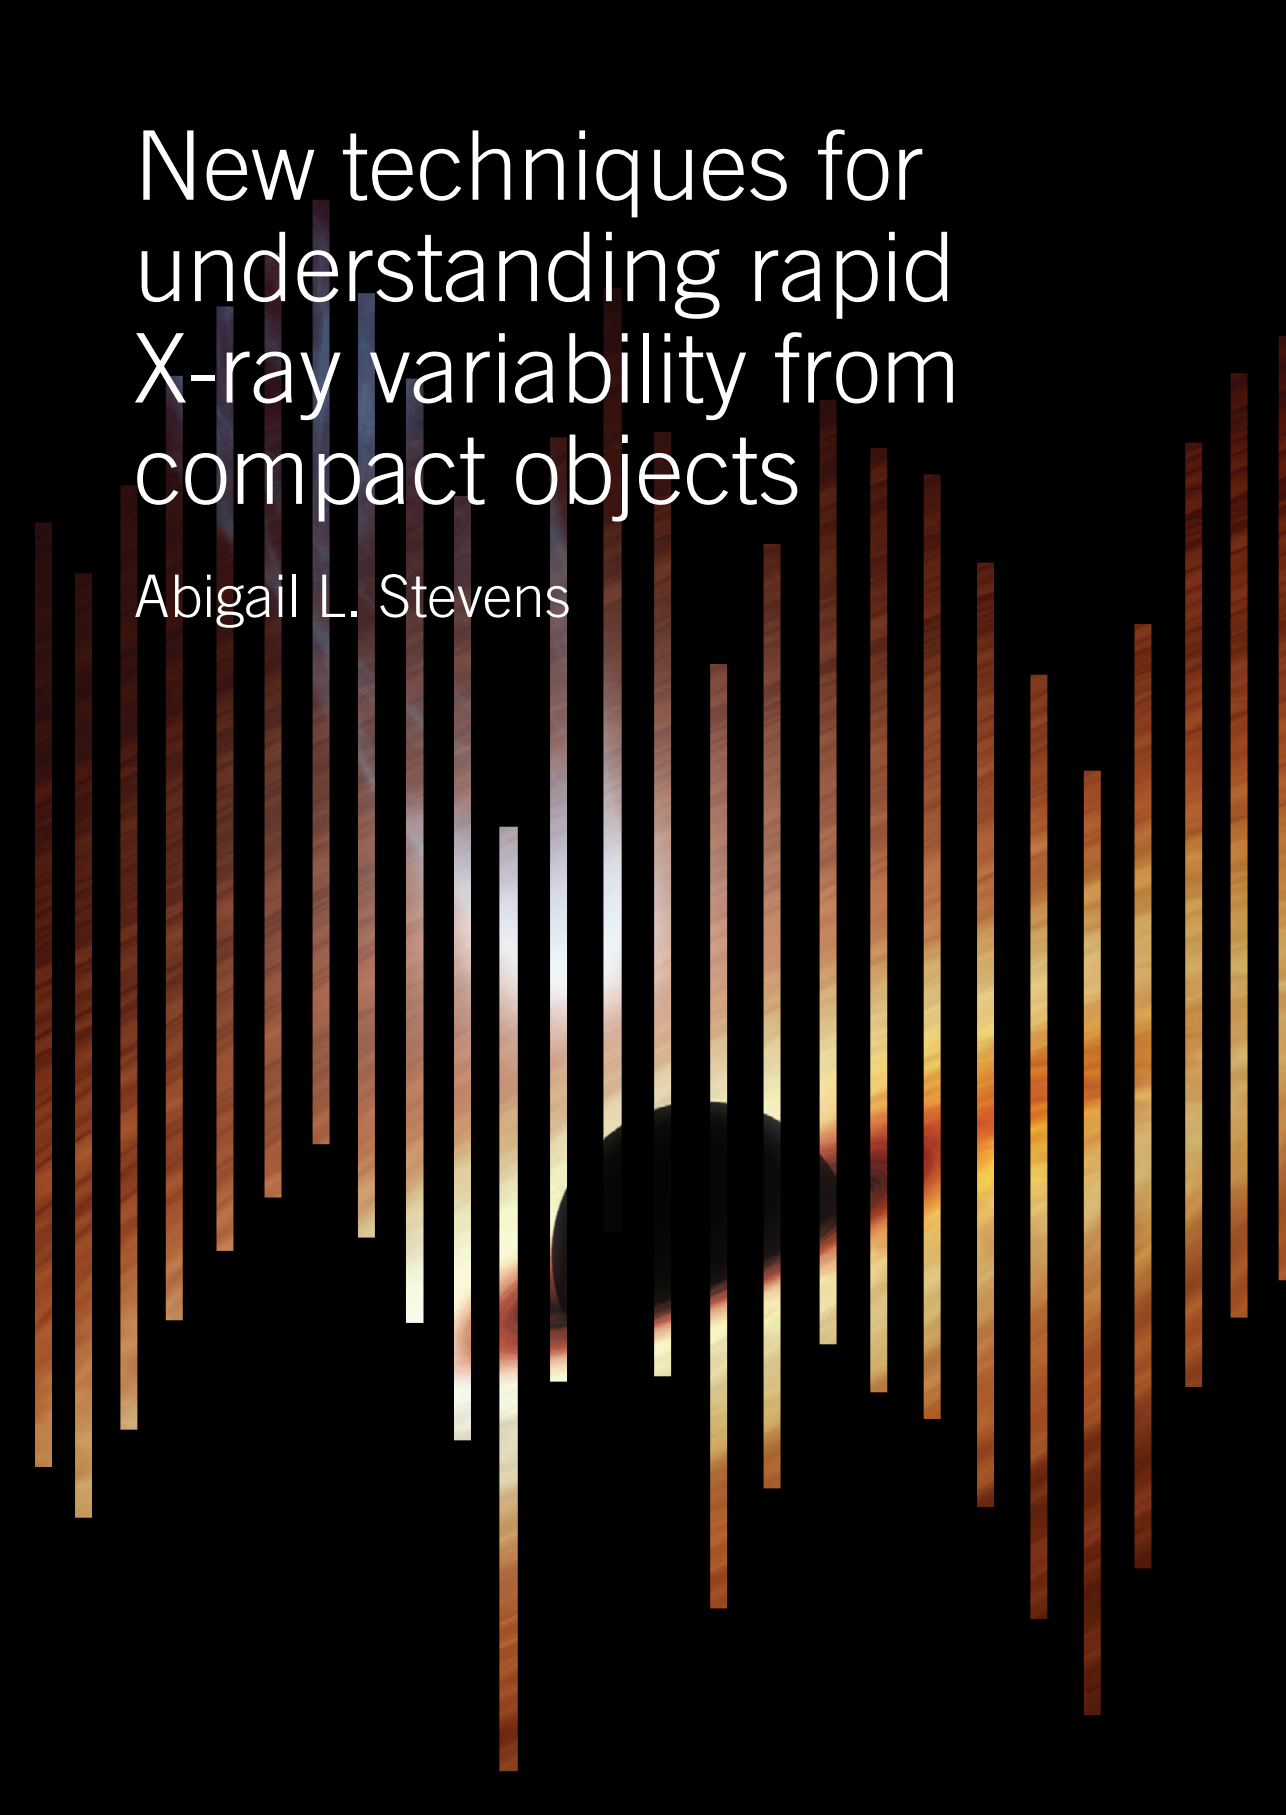

New techniques for understanding rapid X-ray variability from compact objects

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K. I. J. Maex
ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen in de Agnietenkapel
op donderdag 19 april 2018, te 12:00

door

Abigail Lauren Stevens

geboren te New York, Verenigde Staten van Amerika

Promotiecommissie:

Promotores:	dr. P. Uttley	Universiteit van Amsterdam
	prof. dr. M. B. M. van der Klis	Universiteit van Amsterdam
Overige leden:	prof. dr. C. Done	Durham University
	prof. dr. E. M. Cackett	Wayne State University
	prof. dr. S. B. Markoff	Universiteit van Amsterdam
	prof. dr. R. A. D. Wijnands	Universiteit van Amsterdam
	dr. A. L. Watts	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

The research described in this thesis was supported by the Nederlandse Onderzoekschool Voor Astronomie (NOVA) and was primarily carried out at the Anton Pannekoek Institute for Astronomy (API), Universiteit van Amsterdam. Some of the research described in Chapter 3 was carried out at the Department of Physics, University of Alberta in Edmonton, Alberta, Canada. Some travel support was provided by the Leids Kerkhoven-Bosscha Fonds (LKBF).

Cover: An accreting black hole with a jet, seen through quasi-periodic variability in a light curve. Design by Libby Taggart Singh. Image by NASA/JPL-Caltech.

This is my timey-wimey detector.
It goes 'ding' when there's stuff.
The Doctor, *Doctor Who*

To my past self, who didn't know I could do this.

Contents

1

Introduction

1.1 Accretion in Strong Gravity

Gravity is one of the four fundamental forces of nature. We can understand gravity in the weak-field limit, that is, in flat and weakly-curved spacetime, with Newtonian and post-Newtonian approximations. However, strong-field gravity is far more difficult to examine, since it is not replicable within our solar system. In the general relativistic framework, compact objects like black holes and neutron stars are dense enough to significantly bend spacetime in their vicinity. To measure relative sizes for comparing phenomena around neutron stars and black holes, we use the gravitational radius ($\mathrm{r}_g = GM/c^2$), a length unit that scales with the mass of the compact object. For a $10\,\mathrm{M}_\odot$ black hole, $1\,\mathrm{r}_g = 14.7\,\mathrm{km}$; for a $1.4\,\mathrm{M}_\odot$ neutron star, $1\,\mathrm{r}_g = 2.1\,\mathrm{km}$. The strong-field gravitational regime is invoked when an r_g is within two orders of magnitude of the size of the massive object (see Baker et al. (2015) for a comparison of tests in different gravitational regimes).

Accretion is one of most efficient mechanisms with which to generate energy in the universe. Low-mass X-ray binaries (LMXBs; Figure 1.1) are unparallelled laboratories to probe accretion in strong gravity. In an LMXB, the compact object is in a binary system with a low-mass stellar companion, and the companion overflows its Roche lobe. The overflowing matter is gravitationally attracted towards the compact object, but it has angular momentum. Since it cannot directly accrete, it forms a disk around the compact object, which removes both energy and angular momentum

and thus enables the matter to accrete. Most X-ray binaries are transients, and they go into outburst, or periods of high activity and increased emission (by many orders of magnitude). Outbursts tend to last weeks to months, and recurrence times are on the order of years to decades. In outburst, there is a high mass accretion rate, and viscous heating in the accretion disk causes the accreting plasma to glow brightly in the X-rays as it falls down the potential well of the compact object. The X-rays are dominated by emission processes from deep within the strong-field regime.

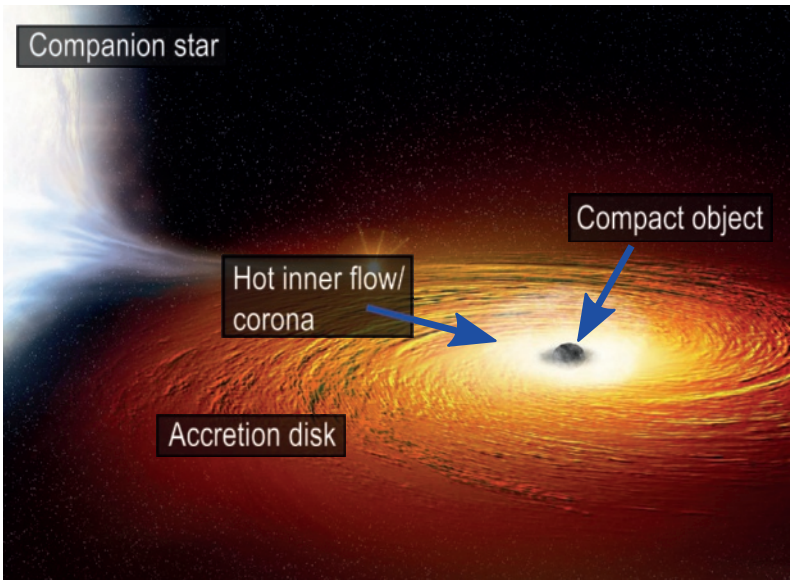

Figure 1.1: An artist's illustration of an LMXB, with key components labelled. The low-mass companion star is comparatively not luminous in X-rays, so we can assume that the observed X-ray emission is coming from the accretion disk, hot inner flow/corona, and compact object (if it is a neutron star). Image adapted from NASA/CXC/M. Weiss.

To all previous, existing, and planned X-ray telescopes, LMXBs appear as point sources, since they are both very small and very far away. The inner part of the accretion disk is less than or equal to nanoarcseconds in angular size, so spatial resolution of the black hole event horizon or neutron star surface would require even finer resolution.[1] Since we cannot directly image the inner regions of LMXBs to study them closely, we turn to spectroscopy and timing to understand them better.

[1] For a $10\,M_\odot$ black hole, its event horizon ($2\,r_g$) is ~ 30 km, which at a distance of 6 kpc, would have an angular diameter of ~ 1 picoarcsecond as seen from Earth. A canonical neutron star with a diameter of 24 km at the same distance would have an angular diameter of just less than a picoarcsecond. For comparison, the *Event Horizon Telescope* will have an angular resolution of 25–300 microarcseconds (Fish et al. 2016), which is ~ 10 million times larger.

In the following sections, we will discuss phenomena and analysis techniques specific to stellar-mass black holes and neutron stars.

1.2 Black Holes

Black holes are the most compact objects in the universe. Due to this compactness, they significantly warp spacetime and play host to extremes of physics that are not found elsewhere in the universe.

According to the No-Hair Theorem, the only observable properties an astrophysical black hole can have are mass and spin (since electric charge would dissipate in an astrophysical setting; Israel 1968; Misner et al. 1973). The event horizon of a black hole is a mathematical surface, the size of which scales linearly with the mass of the black hole for a given spin. Barring exotic theories like Hawking radiation (Hawking 1975), non-accreting black holes are not visible via electromagnetic radiation. In order to study the motion of matter in the strong gravitational regime, we therefore must observe black holes when they are actively accreting. As studied in astronomy and astrophysics, black holes generally belong to one of two families, stellar-mass and supermassive. In this thesis we consider stellar-mass black holes, which are the end product of core-collapse supernovae of the most massive stars (initially $\gtrsim 20\,\mathrm{M}_{\odot}$). A comparison of the known masses of stellar-mass black holes (and neutron stars) is shown in Figure 1.2.

In the following sections we will review the spectral and timing properties of accreting stellar-mass black holes.

1.2.1 Spectral properties

X-ray spectroscopy, in which we analyze the distribution of energies of the incoming photons, has long been a tool to deduce the time-averaged (over an exposure) physical properties of the system. Using X-ray spectroscopy, we can determine the physical components in an LMXB (see Figure 1.1; for reviews, see Done et al. 2007 and Gilfanov 2010): i) a thermally-emitting, optically thick, multi-colour blackbody accretion disk; ii) a hot Comptonizing region of optically thin thermal or non-thermal electrons that up-scatter a source of seed photons; and iii) reflection of high-energy photons emitted by the Comptonizing region that have been reprocessed by the accretion disk before reaching the observer. See Figure 1.3 for an example spectrum with the three components labelled. Depending on the configuration of the Comptonizing region, it is sometimes referred to as a *hot inner flow* if it is in the plane of the accretion disk between the inner edge of the disk and the black hole, or as a *corona* if it is halo-like above the disk and black hole. The most prominent spectral signature of this reflection is the fluorescent iron line at 6.4–6.7 keV in the rest frame (e.g., Ross & Fabian 2005; Miller 2007; Dauser et al. 2013). Due to classical Doppler

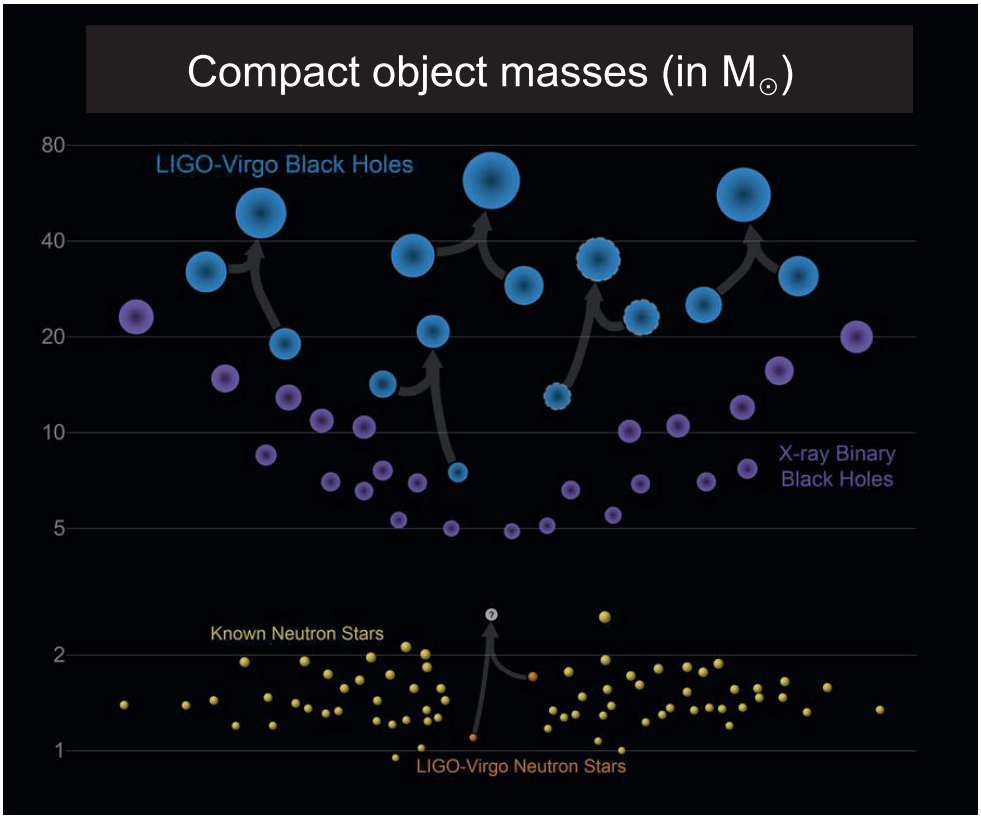

Figure 1.2: An illustration of the known masses of stellar-mass compact objects. The merged neutron star product is labeled with a question mark, since it is not presently known whether a neutron star or black hole was formed. Image adapted from LIGO-Virgo/F. Elavsky/Northwestern University.

boosting of the emission from azimuthal sections of the accretion disk, the iron line has a double-horned profile in the spectrum. This profile is then redshifted by special and general relativistic time-dilation effects and unevenly amplified in intensity by relativistic beaming (Figure 1.4; see Reynolds & Nowak 2003 for a review). The contribution from multiple radii in the disk smears the summed profile. The precise shape of the iron line profile imparts information like the size and location of the Comptonizing region, the velocity, ionization, and gravitational redshift of the material in the accretion disk, and the spin of the black hole.

General relativity predicts that the inner edge of accretion disk lies at the innermost stable circular orbit (ISCO), which is $1\,r_g$ in the Kerr metric for a prograde maximally rotating black hole, and $6\,r_g$ in the Schwarzschild metric for a non-rotating black hole, and $9\,r_g$ in the Kerr metric for a retrograde maximally rotating black hole

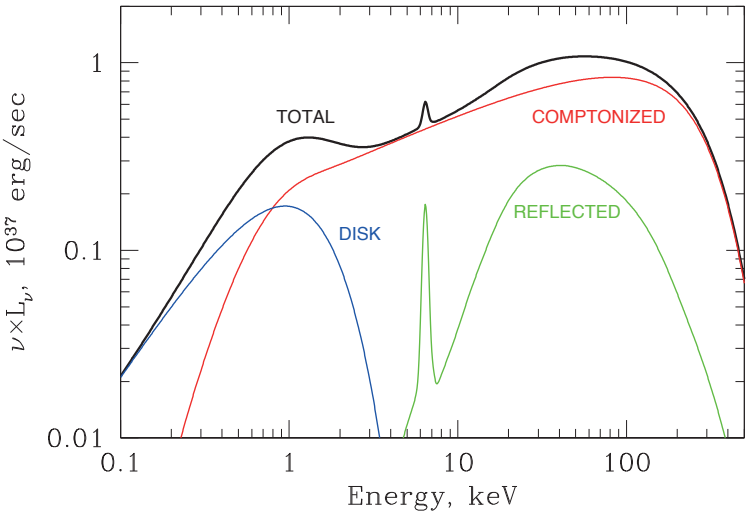

Figure 1.3: An example (hard-state) energy spectrum of an accreting black hole LMXB in outburst. The accretion disk multi-colour blackbody component, Comptonized component, and reflection component are labelled. Image from Gilfanov (2010).

(Misner et al. 1973). In the systems we observe, the ISCO is likely somewhere in between, depending on the spin of the black hole.

Over the course of an outburst, the energy spectrum of a black hole LMXB changes (see, e.g., Miyamoto et al. 1991; Nowak 1995; Belloni 2004; Done et al. 2007; Belloni 2010). These different spectral states are classed as hard (when the spectrum dominated by hard X-ray Comptonized emission), intermediate (both Comptonized and soft blackbody emission), and soft (when the spectrum is dominated by soft X-ray blackbody emission). As a further identifier, 'high' and 'low' are sometimes used, referring to X-ray flux or count rate.

The hardness evolution is typically plotted in a hardness-intensity diagram (Figure 1.5), in which the hardness ratio (a ratio of the flux in a hard energy band to the flux in a soft energy band) is plotted versus the X-ray flux or count rate in a broad energy band. The outburst starts when the source rises out of quiescence into the low-hard state, labelled 'A' in Figure 1.5. In the low-hard state, there is a low photon count rate and the spectrum is dominated by hard Compton up-scattered photons. The systems then evolve upward to the hard state (technically the high-hard state, labelled 'B') , where spectral shape remains the same but the flux increases. Then the system maintains roughly the same count rate, but the distribution of photons changes dramatically through the hard- and soft-intermediate states ('C' and 'D', respectively). This transition through the intermediate states to the soft state happens

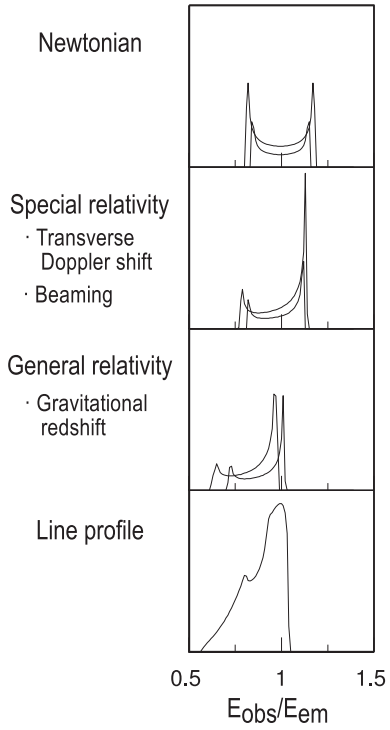

Figure 1.4: A diagram of the different effects on the iron line profile, shown versus the ratio of the observed energy to the emitted energy. The two lines in the top three panels represent two different disk radii. Image adapted from Fabian (2013).

on the order of days to weeks, and the transition is often not smooth, as evidenced by the zig-zags between 'C' and 'D' in Figure 1.5. In the soft state ('E'), the soft X-ray emission from the accretion disk dominates the spectrum. Finally, while in the soft state, the count rate decreases, and then the source moves back through the intermediate states to the low-hard state at a lower overall count rate than in the rise. Most outbursting black hole LMXBs follow similar tracks in the hardness-intensity diagram.

The spectral evolution through the hardness-intensity diagram tracks the changing accretion flow. In quiescence, the hydrogen in the disk is not ionized, and there is a low mass accretion rate. In a truncated disk scenario, the accretion disk would extend inwards until $\sim 10^2$–$10^4\,\mathrm{r}_g$, and the region between the accretion disk and compact object would have a geometrically thick, optically thin, hot inner accretion flow (Esin et al. 1997; Frank et al. 2002; Done et al. 2007). Then to spur the source into outburst, a disk instability mechanism occurs where cascading ionization of the disk hydrogen increases viscosity, which thus drastically increases the mass accretion

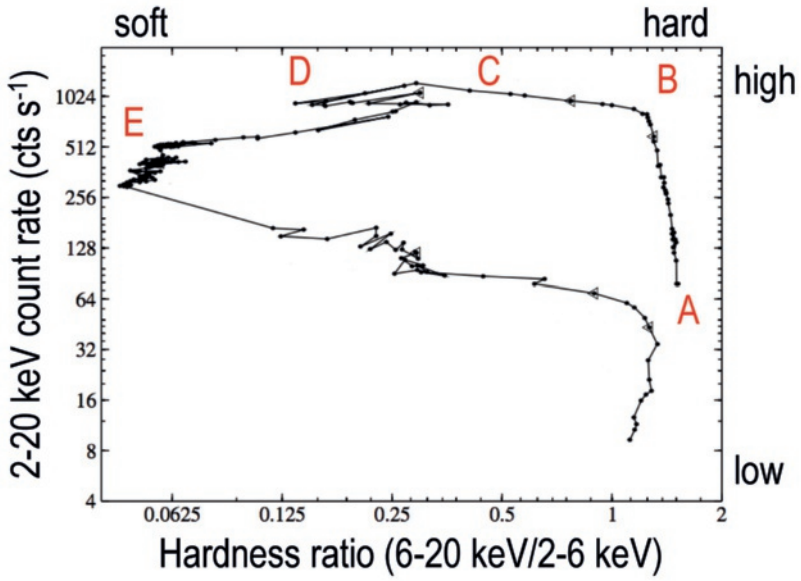

Figure 1.5: A hardness-intensity diagram of the black hole GX 339–4 during its 2012 outburst, sometimes also called a 'turtle head' diagram. This data was taken using the *RXTE* PCA (see Section 1.4.1 for more on *RXTE*). Five spectral states are labelled A-E, corresponding to the low-hard ('A'), high-hard or hard ('B'), hard-intermediate ('C'), soft-intermediate ('D'), and high-soft ('E') states. Observations from the soft-intermediate state of this outburst, near 'D', are used in Chapter 2. Image adapted from Nandi et al. (2012).

rate and luminosity (Lasota 2001). While the system is in the hard state, the Comptonization strongly dominates the spectrum, with little-to-no disk blackbody (Figure 1.6; e.g., Gierliński et al. 1999; García et al. 2015). Additionally, a reflection component is often seen in the spectrum with a reflection fraction of ~ 0.2–0.3 (García et al. 2015) to ~ 0.4–0.5 (Gilfanov et al. 2000). As the mass accretion rate increases, the truncated disk scenario has the disk inner radius moving inwards towards the ISCO (Done et al. 2007). In the intermediate states, the Comptonizing region dissipates as the disk emission increases (e.g., Gierliński et al. 1999). In the soft state, the disk dominates the X-ray emission (Figure 1.6) and the inner edge of the disk is at the ISCO. In the decline of the outburst, the disk emission and mass accretion rate die down (and the inner edge of the disk possibly truncates), and the Comptonizing region either re-forms or takes over again (Nowak et al. 2002). In the final low-hard state, the Comptonizing region again dominates the X-ray emission.

Some black hole LMXBs have radio jets, which are collimated beams of relativistic particles commonly observed via synchrotron emission in the radio to sub-mm bands (see Fender et al. 2004 and Homan & Belloni 2005 for reviews). Black hole LMXBs with jets are sometimes called 'microquasars'. The jets are not constant, and their

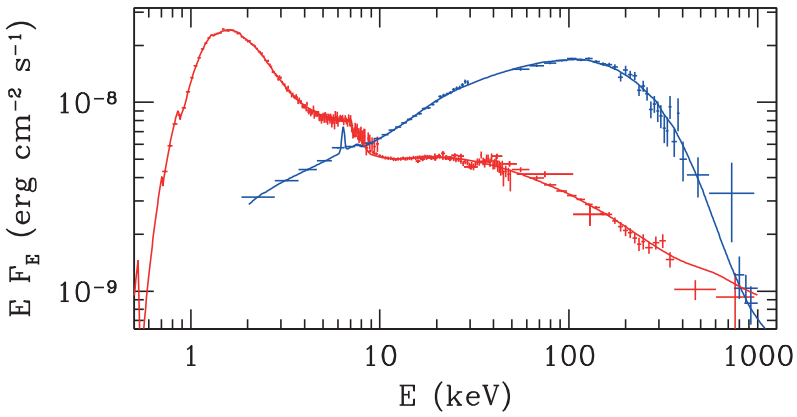

Figure 1.6: Example spectra in the hard state (blue) and soft state (red) from Cygnux X-1. Image from Done et al. (2007), adapted from Gierliński et al. (1999).

properties correlate with the spectral state of the black hole. In the hard state, the jet is 'on' and the emission is persistent (e.g., Gallo et al. 2003). In the hard-intermediate state, the jet enters the 'ballistic' regime, in which the emission is not constant and the material in the jet is clumpy (see Fender et al. 2004 and references therein). Then, in the soft-intermediate state, the jet turns off and no radio emission is detected (e.g., Fender et al. 2009). In the soft state, the jet is not present (e.g., Corbel et al. 2001). Then in decline of outburst, the jet turns back on in the intermediate states, and there is a strong jet again in hard state.

Since the jet occurs in both the rise and decline of the outburst but only in the hard-intermediate and hard states, we infer that the radio jet does not correlate with the count rate or luminosity (a proxy for the accretion rate), just the spectral state (a proxy for the physical geometry of the system). Furthermore, the presence of the jet correlates with the presence of stronger Comptonization in the spectral states, so it is thought that the two could be physically connected. For example, the Comptonizing region could be the base of the jet, as suggested by Markoff et al. (2005).

So, spectroscopy has informed a significant amount of what we know about the time-averaged properties and physics processes of LMXBs. However, since traditional X-ray spectroscopy only measures the time-averaged energy distribution of the observed photons, it cannot probe rapid changes in these systems. As we will discuss in the next section, there is rapid variability in the X-ray emission from black hole LMXBs, and the variability also changes with outburst evolution.

1.2.2 Timing properties

When actively accreting, the X-ray emission from black hole LMXBs can be highly variable on a wide range of timescales (for reviews, see McClintock & Remillard 2006; van der Klis 2006b). Timing analysis naturally gives the timescale of the variable emission, but it does not directly give the physical processes responsible for it. The timing features in black hole LMXBs are low- and high-frequency quasi-periodic oscillations, in addition to noise processes like observed Poisson noise and intrinsic broadband noise. We use power spectra, the modulus squared of the Fourier transform of a light curve, to assess the timing properties of a light curve. A power spectrum tells us the amount of variability in a light curve at each Fourier frequency (see van der Klis 1989 for a detailed explanation of power spectral analysis). While technically the power spectrum (or power density spectrum) represents the mean amplitude of variability of the underlying process and the periodogram is the frequency distribution of the variability that is measured with timing analysis, astronomers typically refer to the measured frequency distribution itself as a power spectrum.

For an X-ray light curve $x(t)$ expressed in counts per time bin, it has a Fourier transform $X(\nu)$ that can be computed by the discrete Fourier transform,

$$X_n = \sum_{k=0}^{N-1} x_k \exp\left(2\pi i n k/N\right) \tag{1.1}$$

for each frequency bin n and each time bin k, where the frequency in each frequency bin $f_n = n/(N\,dt)$, N is the number of consecutive bins in the light curve, dt is the sampling time step of the light curve, and $n = 1, 2, 3, .., N/2$. The power spectrum $P(\nu)$ is defined as

$$P(\nu) = \left|X(\nu)^2\right| . \tag{1.2}$$

The three normalizations for power spectra are *absolute rms-squared*, *fractional rms-squared*, and *Leahy* (Leahy et al. 1983), respectively computed by

$$P(\nu)_{\rm abs} = P(\nu) \times \frac{2\,dt}{N} \tag{1.3}$$

$$P(\nu)_{\rm frac} = P(\nu) \times \frac{2\,dt}{N\langle x\rangle^2} \tag{1.4}$$

$$P(\nu)_{\rm Leahy} = P(\nu) \times \frac{2\,dt}{N\langle x\rangle} \tag{1.5}$$

where $\langle x\rangle$ is the mean count rate of the light curve. Power spectra are typically averaged over many sequential segments of a light curve to reduce the statistical noise, and they can also be re-binned in frequency. The integral of the fractional

rms-squared power spectrum (either over the whole set of positive Fourier frequencies or over a smaller range of consecutive Fourier frequencies) is the variance. The square root of the variance is the fractional rms ('root mean square'), which measures the overall amount of variability in a power spectrum in a given frequency range. For more technical details of the mathematics of the power spectrum, we refer to van der Klis (1989) and Uttley et al. (2014).

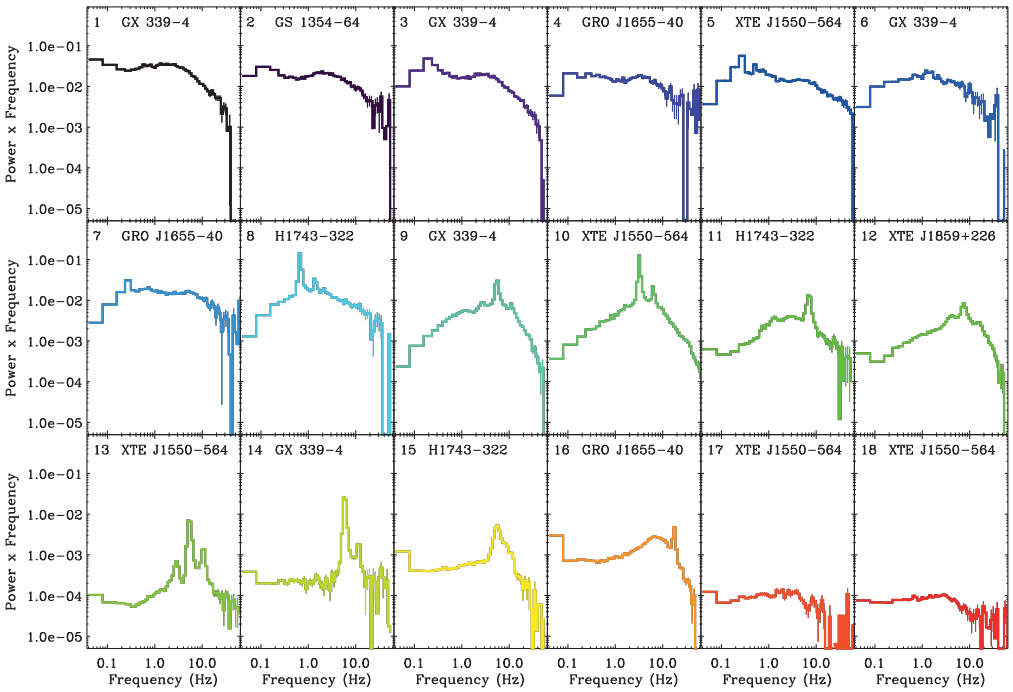

Figure 1.7: Power spectra from different sources showing typical timing properties of a black hole LMXB in outburst as it transitions through the spectral states in the first half of the outburst, progressing left to right and top to bottom. As determined by the shape and rms of the noise in Heil et al. (2015b), panels 1–7 are in the hard state, panels 8–11 are in the hard-intermediate state, panels 12–15 are in the soft-intermediate state, and panels 16-18 are in the the soft state. There are Type C QPOs in panels 8–11, and Type B QPOs in panels 13–15. Image from Heil et al. (2015b).

Quasi-periodic oscillations (QPOs) have been seen in the X-ray light curves of accreting black holes and neutron stars since the mid 1980s (van der Klis et al. 1985). QPOs from black hole binaries are broadly classed in two categories, low-frequency ($\sim$0.1–30 Hz) and high-frequency ($\gtrsim$100 Hz). Low-frequency QPOs in black hole LMXBs are further subdivided into three categories, based on their spectral and timing properties: Type A, Type B, and Type C (Remillard et al. 2002; Casella et al. 2005). Type A QPOs appear in the soft-intermediate and soft states with weak broadband noise and with a very weak, broad QPO peak at around 8 Hz. Type B

QPOs appear fleetingly in the soft-intermediate state, and have weak broadband noise with a very strong QPO at around 6 Hz and relatively weak QPO harmonic. Type C QPOs appear in the hard and hard-intermediate states, sit on top of a significant broadband noise component, and have a QPO with a relatively strong QPO harmonic. The Type C QPO centroid frequency tends to vary significantly over the frequency range ~ 0.1–15 Hz, with the higher frequencies typically found in softer spectral states. A diagram of characteristic power spectra throughout the spectral states in the first half of an outburst is shown in Figure 1.7. Type Cs are the most common type of low-frequency QPO detected, followed by some Type Bs and very few Type As (Motta et al. 2015). Type B and C QPOs tend to be mostly comprised of medium-to-hard X-ray emission ($\gtrsim 5$ keV). Type B QPOs are analyzed in Chapters 2 and 5.

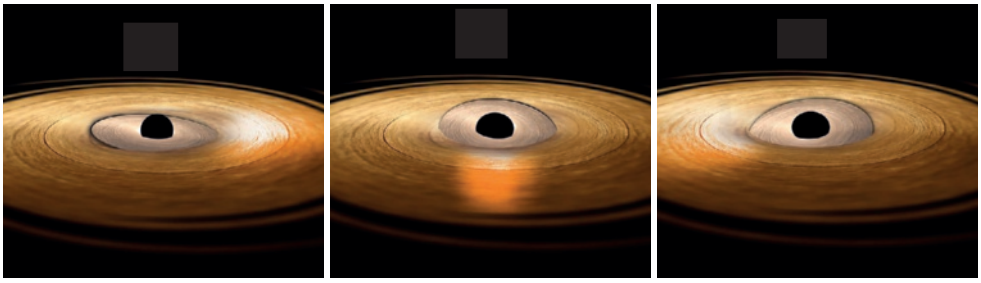

Figure 1.8: An artistic depiction of an accreting black hole with a precessing hot inner flow. The configuration of the hot inner flow relative to the inner accretion disk is shown at 0, 0.25, and 0.5 of a QPO cycle (respectively, left to right). We see that the emission from the hot inner flow sweeps around the inner accretion disk illuminating and heating it. Image adapted from ESA/ATG medialab/A. Ingram.

High-frequency QPOs, which are found at frequencies above about 100 Hz, occur at frequencies consistent with Keplerian motion at the inner edge of the accretion disk (e.g., Remillard et al. 1999a,b; Strohmayer 2001; Kluzniak & Abramowicz 2001). In black holes, high-frequency QPOs are not commonly observed and have historically had rather poor signal-to-noise. The high-frequency QPOs that have been observed are seen when the source is in the soft state (Belloni et al. 2012). Black hole high-frequency QPOs are somewhat analogous to kHz QPOs in accreting neutron stars (though not entirely; see, e.g., Motta et al. 2017), the latter of which are frequently seen with high signal-to-noise in a variety of sources (discussed in more detail in Section 1.3.2).

The most ubiquitous feature in a power spectrum of an accreting compact object is flat Poisson noise. Poisson noise is due to the counting statistics of X-ray timing observations. For a power spectrum with absolute rms-squared normalization, the expected Poisson noise level is $2\langle x \rangle$, for fractional rms-squared, the expected Poisson noise level is $2/\langle x \rangle$, and for Leahy normalization, the expected Poisson noise level is

2. When plotted as power × frequency,[2] the Poisson noise has a power-law index of +1. Poisson noise dominates the power spectrum at high frequencies.

Broadband noise, or band-limited noise, is weakly-peaked low-frequency intrinsic noise ($\lesssim 10$ Hz) that appears in the hard and hard-intermediate states. It was first observed in the early 1970s and is often modelled with a broad Lorentzian (or combination of broad Lorentzians) and/or a broken power-law (Terrell 1972). Broadband noise is most prominent in the hard and hard-intermediate states (see panels 1–12 in Figure 1.7) and is thought to be due to accretion fluctuations propagating inward through the accretion disk (Lyubarskii 1997; Arévalo & Uttley 2006; Uttley et al. 2011; Rapisarda et al. 2016). There is also general low-frequency noise in the soft spectral states that tends to be modelled with a power-law.

Low-frequency QPO models

The physical origin of low-frequency QPOs is unclear. Models of low-frequency QPOs generally fall under one of two categories: intrinsic variability and geometric variability.

Intrinsic models invoke an intrinsically quasi-periodic process causing overall luminosity variability in one or more spectral components. Examples of intrinsic models include trapped oscillations in the inner accretion disk (Nowak & Wagoner 1991), standing shockwaves in the accretion flow (Chakrabarti 1996), Rayleigh-Taylor gravity waves in the inner accretion flow (Titarchuk 2003), or Comptonized oscillations in the base of a jet (Giannios et al. 2004). *Geometric* models produce variations in the observed luminosity due to geometric variations in the system. Two examples of geometric models are general relativistic Lense-Thirring precession (Stella & Vietri 1998) and an oscillating shock causing a moving corona (Suková et al. 2017). Often, theoretical geometric models have a periodic geometric variation that is made to be quasi-periodic in appearance due to accretion rate fluctuations.

Evidence points towards a geometric origin due to the dependence of QPO properties on geometric parameters of the system. Schnittman et al. (2006) found a correlation between QPO amplitude and the inclination of the binary orbit with respect to our line of sight for a sample of 7 sources with Type C QPOs. Larger studies of this nature were carried out simultaneously by Motta et al. (2015) and Heil et al. (2015a) and found a clear correlation between the QPO rms and the inclination of the binary orbit with respect to our line of sight. Furthermore, van den Eijnden et al. (2017) found a correlation between the sign of the energy-dependent QPO phase lags and the binary orbit inclination. Interestingly, Type B and Type C low-frequency QPOs show opposite trends in the QPO rms-inclination correlation, indicating that they have different origins.

[2] Power × frequency is often plotted against frequency to make it easier to determine the *relative* contribution of the variability at each frequency to the total variability in the light curve.

One geometric model gaining traction is general relativistic Lense-Thirring precession of the entire Comptonizing region (Figure 1.8; Stella & Vietri 1998; Ingram et al. 2009). Lense-Thirring precession, also known as frame dragging, is a consequence of the Kerr metric, which describes the spacetime outside a rotating black hole (see Misner et al. 1973). It arises when the orbital plane of a test particle orbit is misaligned with respect to the spin axis of the black hole, and the orbital plane precesses around the spin axis of the black hole. The frequencies for standard Lense-Thirring precession of test particles in orbit at the inner edge of the accretion disk around a stellar-mass black hole are marginally higher than the low-frequency QPOs that are observed. However, in this model, the entire hot inner flow precesses like a rigid body in and out of the plane of the accretion disk (Figure 1.8), which causes the precession to occur at a lower frequency, matching the range of observed frequencies of low-frequency QPOs. Hydrodynamical simulations have shown that this is possible if the orbital axis of the inner accretion disk is misaligned with the spin axis of the black hole Fragile & Anninos 2005; Fragile et al. 2016).

The most recent evidence of Lense-Thirring precession is a modulating iron line energy, from a precessing hot inner flow reflecting alternately off the blue-shifted and red-shifted sides of the accretion disk (Fragile et al. 2005; Ingram & Done 2012). Ingram et al. (2016, 2017) applied this model to Type C QPOs, which have strong harmonics, where the harmonic comes from the underside of the precessing flow illuminating the disk as well. In addition, Sobolewska & Życki (2006), Axelsson et al. (2014), and Axelsson & Done (2016) found via rms spectra that the Type C QPO and its harmonic were originating from the Comptonized component (in Axelsson & Done 2016, they found that the QPO and harmonic could be produced from different parts of an inhomogeneous Comptonizing region). Finally, new simulations from Liska et al. (2017) show that it is possible to get a precessing corona-like base of a radio jet to give low-frequency QPOs in the X-ray light curve. This work, published after Stevens & Uttley (2016), supports the physical interpretation in Chapter 2 that a wobbling tall jet-like Comptonizing region could be responsible for the Type B QPO we analyzed.

On the other hand, it might not be possible to have sufficient misalignment between the black hole spin axis and inner disk axis for Lense-Thirring precession to occur with strong-enough modulation to generate QPOs in the X-ray light curves. Much of the iron line spectroscopy and continuum fitting literature assumes that the disk is not truncated in the hard state and extends inward to the ISCO, and that the orbital plane of the disk is fully perpendicular to the black hole spin axis (e.g., Steiner et al. 2011). Likewise, Maccarone (2002, 2015) found that if there were a misalignment when the LMXB formed, it would not last long-enough in the LMXB's lifetime to explain the number of observed sources with low-frequency QPOs.

Each model for the origin of low-frequency QPOs can predict QPO-phase-dependent changes to the energy spectrum. In order to break degeneracies between QPO mod-

els, we want to measure the spectral-timing behaviour of the QPO using a growing arsenal of techniques.

1.2.3 Spectral-timing techniques

Spectral-timing analysis is a new pathway to understanding X-ray data that simultaneously examines its spectral and temporal properties to deduce underlying causal relationships between different emitting components. The foundation of spectral-timing analysis is in the understanding that photon arrival time and photon energy should not be analyzed separately, since they are intrinsically connected by the underlying physics. Since spectral analysis gives the physical process and timing analysis gives the timescale, the combination can link the physical processes to timescales to better understand the emission region. For a full description of spectral-timing techniques with recipes and examples, see Uttley et al. (2014). Here we give a brief overview.

The *cross-spectrum* is a type of Fourier analysis that compares two light curves and measures the components that are correlated between the two. On an intuitive level, it finds the timescale(s) where two independent quantities x and y vary together. For two independent X-ray light curves $x(t)$ and $y(t)$ that have the Fourier transforms $X(\nu)$ and $Y(\nu)$, the cross-spectrum $C(\nu)$ is defined as

$$C(\nu) = X(\nu)\, Y^*(\nu)\,, \tag{1.6}$$

where * denotes the complex conjugate. The cross spectrum is complex, and thus it has an amplitude and a phase at each Fourier frequency. The cross-spectrum amplitude is the cross-spectral version of the power spectrum. It is computed as

$$A(\nu) = \sqrt{R(\nu)^2 + I(\nu)^2}\,, \tag{1.7}$$

where $R(\nu)$ is the amplitude of the real component and $I(\nu)$ is the amplitude of the imaginary component of the cross spectrum $C(\nu)$. It measures the amount of correlated variability at each Fourier frequency. The cross spectrum is typically averaged over segments, and can be binned in frequency (just as with the power spectrum). The same normalizations from equations 1.3–1.5 can be applied to the cross spectrum amplitude using $\sqrt{\langle x\rangle\langle y\rangle}$ as the mean count rate, where $\langle x\rangle$ and $\langle y\rangle$ are the mean count rates of light curves $x(t)$ and $y(t)$, respectively. We note that the mean of the Poisson noise in the cross spectrum is zero, since Poisson noise is by definition not correlated between independent light curves, but the scatter on the Poisson noise does contribute to the cross spectrum by a minor amount (Vaughan & Nowak 1997). For an example of the cross spectrum amplitude, see Figure 4.3 in Chapter 4.

Lags, computed from the cross-spectrum phase, measure the frequency- and energy-dependent phase shifts in the variability (Vaughan et al. 1994). *Lag-frequency spectra* are measured between a soft and hard light curve. It measures the variability in the hard light curve that leads or lags the variability in the soft light curve at each Fourier

frequency. *Lag-energy spectra* are measured between a 'reference' light curve from a broad energy band and 'channel of interest' light curves in many different narrow energy bands. It measures if the variability in each channel of interest leads or lags the variability in the reference band, and has a spectral resolution corresponding to the width of the channels of interest. The lag-energy spectrum gives indications of causality with respect to the coarse energy information and the amplitude of the time lag. Non-monotonic structure in the lag-energy spectrum indicates that there is likely interesting spectral-timing behaviour, such as spectral pivoting or phase shifts between spectral components (as opposed to, e.g., a modulation in the overall flux normalization). For example lag-energy spectra, see Figures 2.2, 4.4, and 5.4 in Chapters 2, 4, and 5, respectively.

The *rms spectrum* measures the amount of variability in a light curve in narrow energy bands (Revnivtsev et al. 1999). It indicates at which energies the variability is strongest, and is best applied to signals that are thought to be from a changing normalization of a single or unrelated spectral component(s). For examples of rms spectra applied to low-frequency QPOs, see Sobolewska & Życki (2006); Axelsson et al. (2014); Axelsson & Done (2016). The *covariance spectrum* is the cross-spectral counterpart to the rms spectrum (Wilkinson & Uttley 2009). It measures the amplitude of correlated variability between a light curve from a narrow energy band with a light curve from a broad 'reference' energy band. We note that the idea to use a broad reference band for cross-spectra and cross-correlation (as used in Chapters 2, 4, and 5) was first developed for covariance spectra. Yet more spectral-timing methods include the Hilbert-Huang Transform to decompose a non-stationary time signal (e.g., Su et al. 2015) and spectral decomposition with principle component analysis (e.g., Koljonen 2015).

The culmination of spectral-timing analysis for periodic and quasi-periodic signals is phase-resolved spectroscopy, in which the energy spectrum can be measured and fitted at multiple phases across one period or cycle of variability. Phase-resolved spectra can be more straightforward to physically interpret than the rms spectrum or covariance spectrum, and makes use of both the amplitude and phase information. For periodic signals coming from pulsars, phase-resolved spectroscopy can be done by folding light curves on the period of the pulsation to build up an average pulse shape. Such analyses, when applied to Fermi observations of gamma-ray pulsars, broke degeneracies between models for pulsar emission and constrained the structure of the magnetic field (Harding 2013). However, phase-resolved spectroscopy of QPOs (*quasi*-periodic signals by definition) had been an unreachable goal, since they are not coherent enough to phase-fold in the time domain. In Chapter 2, we introduce a new technique for phase-resolved spectroscopy that we apply to low-frequency and kHz QPOs in Chapters 2, 4, and 5. Our phase-resolved spectroscopy technique averages the QPO signals in the cross-spectrum, so that the phases of the signals add while the phases of the noise cancels. Two other techniques for phase-resolved spectroscopy

of low-frequency QPOs are from Miller & Homan (2005) and Ingram & van der Klis (2015). Miller & Homan (2005) combined spectra at the well-defined maxima and minima of the QPO in the light curve. Ingram & van der Klis (2015) measured the complex amplitude and phase of the QPO and harmonic in the Fourier transform of the light curves in each energy channel and used this information to reconstruct the energy-dependent QPO signal.

1.3 Neutron Stars

Neutron star LMXBs have many types of emission and variability that are similar to those seen from black holes. However, there is an extra layer of complexity with neutron stars, in that they have a surface.

Underneath the surface, neutron stars contain the densest form of matter known to exist. Their average density is that of an atomic nucleus ($\sim 2.8 \times 10^{14}\,\mathrm{g\,cm^{-3}}$), and given that they have a density gradient, their inner cores are up to an order of magnitude denser than an atomic nucleus (Glendenning 2000; Camenzind 2007; Watts et al. 2016). The regions inside a neutron star, moving radially inward, are the outer crust comprised of an iron-rich lattice, an inner crust of neutron-rich atomic nuclei, and a core likely comprised of superfluid neutrons with some superconducting protons and relativistic electrons (Figure 1.9; Watts et al. 2016). The inner core composition is unknown, with theories ranging from quark matter to various particle condensates. Furthermore, neutron star matter is considered to be cold, since the thermal energy of the particles is less than the Fermi energy ($\sim 30\,\mathrm{MeV}$; Glendenning 2000). While particle physics experiments like ALICE at CERN can probe some regimes of hot dense matter, the cold supra-nuclear-density conditions inside neutron stars cannot be re-created on Earth (Watts et al. 2016).

The holy grail of neutron star astrophysics research is to determine the equation of state of the neutron star inner core. The equation of state of the cold matter in neutron stars is expressed as a relation of pressure and density. Different neutron star equations of state are classified as 'soft' or 'stiff', referring to whether they give lower or higher pressure at a certain density, which then gives smaller or larger neutron star masses, respectively, for a given radius. Since the pressure-density relation can be derived by combining many mass and radius measurements of neutron stars, obtaining constraints on the masses and radii will place limits on which theoretical equations of state are plausible (see Lattimer & Prakash 2016 for a recent overview of realistic neutron star equations of state).

To study the unique phenomena associated with neutron stars, we can analyze the unique variability in X-ray emission from neutron stars, such as *pulsations* from accretion onto the surface at the magnetic poles (see van der Klis 2000 for a review), *thermonuclear burst oscillations* from a thermonuclear explosion of accreted material

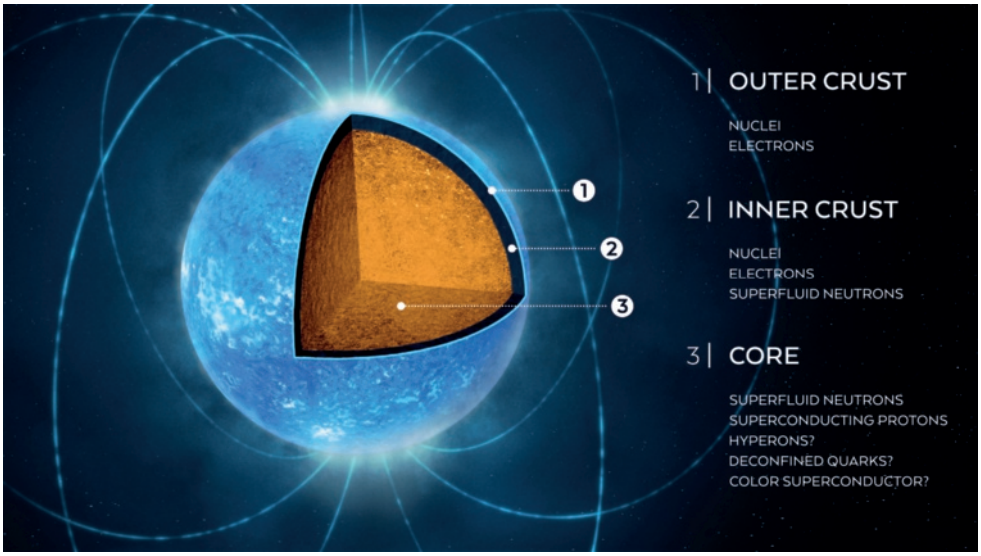

Figure 1.9: An artistic illustration of the interior of a neutron star, with the outer crust, inner crust, and core labelled. Image from Watts et al. (2016).

on the surface (see Galloway et al. 2008 for a review), *millihertz QPOs* from marginally stable nuclear burning on the surface (see Bult et al. 2017 for a recent example), and *kilohertz QPOs* from the 'boundary layer' accretion flow just above the surface or inhomogeneities in the inner accretion disk (colloquially, 'hot blobs') orbiting the neutron star (e.g., Wijnands et al. 2003).[3] Thermonuclear burst oscillations and kilohertz QPOs are explored in this thesis, and descriptions of each are provided in the following sections.

1.3.1 Thermonuclear burst oscillations

When a neutron star is actively accreting, matter can accumulate on its surface. Once the matter reaches a critical temperature and density, thermonuclear burning (i.e., nuclear fusion) is triggered and the accreted surface layers ignite in an unstable thermonuclear runaway (see Watts 2012 for a review). The thermonuclear burning releases a significant burst of X-rays, and this scenario is called a Type I X-ray burst.

Some Type I X-ray bursts show oscillations in the light curve, the frequency of which strongly corresponds to the known spin frequency of the neutron star, during the rise and decay of the burst. A common theory (and the one assumed in this thesis) is that the burst oscillations come from a 'hotspot' on the surface of the neutron star

[3]There are also low-frequency QPOs seen in neutron star LMXBs: horizontal branch oscillations, normal branch oscillations, and flaring branch oscillations. There are many similarities between these and black hole low-frequency QPOs; see, e.g., Motta et al. (2017) for a recent comparison.

(e.g., Artigue et al. 2013). This hotspot is misaligned with the spin axis of the neutron star, and the observed luminosity changes as the neutron star spins, which gives a periodic light curve. This 'lighthouse effect' is also part of the mechanism behind accreting X-ray pulsars. Since the X-rays emitted during the thermonuclear burst are from the surface layers of the neutron star, the geometric parameters of the system, including the neutron star's mass and radius, are encoded in the burst photons (Watts 2012).

Burst oscillations can be modelled with ray tracing, in which the path of the burst photons is computed through the curved spacetime around the neutron star at sequential phases in the neutron star's rapid rotation. Due to gravitational light-bending, an observer at infinity is able to see roughly 3/4 of the neutron star surface at once, so the hotspot is typically visible for the entire rotation. It is possible to have an antipodal hotspot in the southern hemisphere of the neutron star, though in Chapter 3 we only model one hotspot in the northern hemisphere. The shape of the burst oscillation light curve is referred to as the 'pulse profile'.

The first ray-tracing formalism for point-like emission from a Schwarzschild-metric (i.e., non-rotating) neutron star is from Pechenick et al. (1983). This was built upon by Miller & Lamb (1998) for a Schwarzschild metric with special-relativistic Doppler effects, which arise from the rapid rotation of these neutron stars. Poutanen & Gierliński (2003) then used the 'Schwarzschild+Doppler' formalism with the Beloborodov (2002) approximation for computing the photon trajectory, which improves computation time, and expanded the algorithm to include non-trivially-sized hotspots. Since rapidly rotating neutron stars can become oblate in shape due to centrifugal forces, and the effects of spin frequencies $\gtrsim 300$ Hz are evident in the shape of the pulse profile, Cadeau et al. (2007) and Morsink et al. (2007) explored ray-tracing in an 'oblate Schwarzschild+Doppler' framework. There are also some examples in the literature of ray-tracing emission from neutron stars in non-Schwarzschild metrics, such as a Kerr metric (Psaltis & Johannsen 2012) and a Hartle-Thorne metric (Bauböck et al. 2013). For more background on pulse profile simulations, see the introduction of Stevens (2013). The Schwarzschild+Doppler and oblate Schwarzschild+Doppler ray tracing formalisms are used in Chapter 3.

1.3.2 Kilohertz QPOs

Kilohertz (kHz) QPOs occur at frequencies ~ 400–1200 Hz (van der Klis et al. 1996b; Strohmayer et al. 1996b; for reviews, see van der Klis 1998, 2000, 2006b). These frequencies are the most rapid type of variability that has been observed in an LMXB. kHz QPOs are most often seen in the soft and soft-intermediate spectral states (Motta et al. 2017). In Fourier space, kHz QPOs are often double-peaked (e.g., van der Klis et al. 1997; Méndez et al. 1998a), with the 'upper' and 'lower' kHz QPO exhibiting different spectral-timing properties (e.g., Barret 2013; de Avellar et al. 2013; Peille

et al. 2015; Cackett 2016). The frequencies of kHz QPOs are consistent with the Keplerian velocity in the inner part of the accretion disk (Stella & Vietri 1999). This region may be where the disk meets the neutron star surface in neutron stars with weaker magnetic fields ($\lesssim 10^{10}$ G), or where the magnetic field truncates the accretion disk in neutron stars with stronger magnetic fields ($\gtrsim 10^{10}$ G). As such, kHz QPOs are a diagnostic tool for dynamical interactions between strong-gravity, strong magnetic fields, a surface, and accretion flows.

Neutron stars rotate rapidly, but slower than the Keplerian velocity of the inner disk edge. The boundary layer, where the inner accretion disk may meet the neutron star surface, then corresponds to the region where matter slows down and hits the surface of the neutron star. The emission from the boundary layer is spectrally modelled as a thermal Comptonized component with a low-energy cutoff temperature tied to an accretion-induced hot spot on the neutron star surface (not the temperature of the inner disk) and a high-energy cutoff temperature at ~ 3 keV (for a recent example, see Armas Padilla et al. 2017). The lower kHz QPO may originate from the boundary layer, as explored in Chapter 4.

1.4 X-ray Observatories for Spectral-Timing

The efficacy of the methods we have described to study this plethora of exotic phenomena in black holes and neutron stars is heavily dependent on the quality of the data. Since these sources are very luminous when in outburst, background emission (i.e., non-source emission in the field of view) is largely not an issue compared to other fields in X-ray astronomy like quiescent compact objects or distant galaxies. However, the more photons detected, the better the statistics are, so the best X-ray instruments have a large effective area with which to collect photons. In order to temporally resolve the most rapid variability, very fast time resolution of the detector is necessary. In terms of spectral range, it is necessary to cover both the soft blackbody emission and the hard Comptonized emission, preferably with fine-enough energy resolution around the iron line to characterize the shape of the line profile. Finally, some outbursting sources get exceptionally bright, so the ideal detector will be able to accommodate thousands of photons per second without losing spectral or temporal resolution. In the following sections we highlight the best past mission, best current mission, and best two future missions for spectral-timing observations of X-ray binaries.

1.4.1 *RXTE*

The *Rossi X-ray Timing Explorer* (*RXTE*) was launched by NASA in December 1995 (see Bradt et al. 1993). The most notable instrument for our research is the Proportional Counter Array, or PCA (see Jahoda et al. 1996). The PCA was comprised of

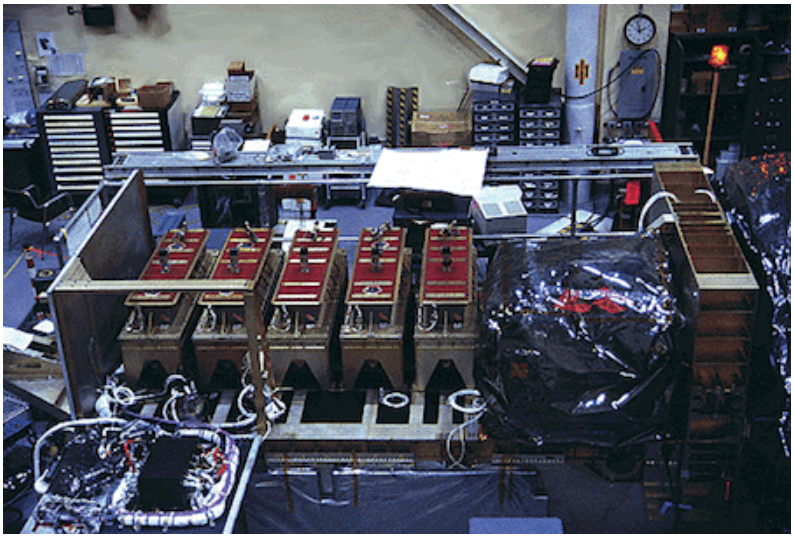

Figure 1.10: A photograph of the *RXTE* PCA as it was being assembled. Image from NASA/GSFC.

five identical proportional counter units (PCUs) that were filled with xenon gas. The five PCUs are visible in Figure 1.10 as the red vertical strips of instrumentation. The PCA detected X-rays in the 2–60 keV energy range with a best energy resolution of 1.08 keV at 6 keV. Since it was built as a timing instrument, its best possible time resolution was 1 μs, though the *RXTE* PCA data used in this thesis has a time resolution of 122 μs. The PCA had a total effective area of 0.65 m^2, and held the title of largest effective area in the 2–60 keV bandpass until the launch of *AstroSat* in September 2015 (for more on *AstroSat*, see Singh et al. 2014 and Yadav et al. 2016). Since *RXTE* was used for spectroscopy as well as precise timing measurements, there were two types of data mode: binned mode, which was used for spectra (i.e., Standard-2 data) and for timing of some very bright sources, and event mode, which was used for timing.

RXTE has an extensive legacy in X-ray timing of both black holes and neutron stars. Among the more than 2,100 papers from *RXTE* data during its lifetime alone,[4] it enabled the discovery of kHz QPOs (van der Klis et al. 1996a; Strohmayer et al. 1996a; van der Klis et al. 1996b; Strohmayer et al. 1996b), accretion-powered millisecond pulsars (Wijnands & van der Klis 1998a; Chakrabarty & Morgan 1998a; Wijnands & van der Klis 1998b; Chakrabarty & Morgan 1998b), magnetars (highly magnetic neutron stars; Ibrahim et al. 2002), and many stellar-mass black holes (McClintock & Remillard 2006). After a very successful 16-year mission lifetime, *RXTE* was de-

[4] https://heasarc.gsfc.nasa.gov/docs/xte/whatsnew/xte_refereed_all.html

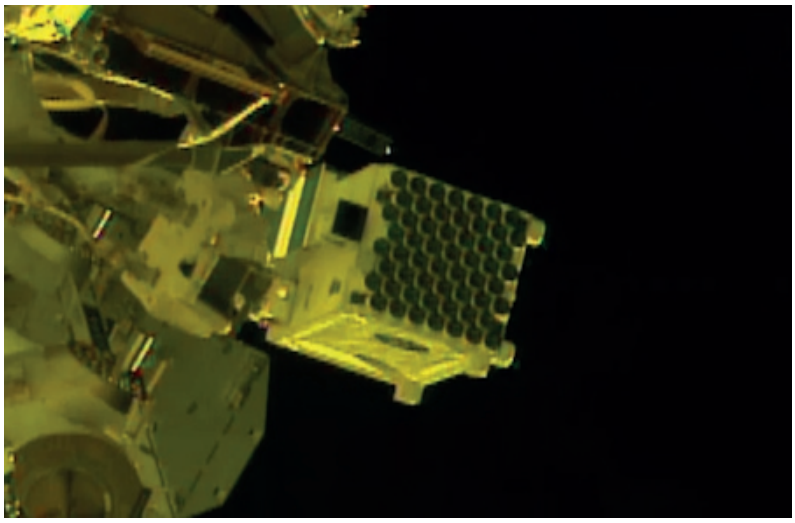

Figure 1.11: A picture of *NICER* in action onboard the International Space Station. Image from video by NASA/GSFC/K. Gendreau.

commissioned in January 2012. Chapters 2 and 4 use *RXTE* PCA data, and Chapter 3 includes simulations that represent *RXTE*-like data.

1.4.2 *NICER*

The *Neutron star Interior Composition Explorer* (*NICER*) is a soft X-ray mission launched by NASA in June 2017, and is attached to an external module on the International Space Station (see Gendreau et al. 2012, Arzoumanian et al. 2014). The *NICER* X-ray Timing Instrument (XTI) has 56 concentrator optics (visible in Figure 1.11) that direct the incoming X-rays onto the CCD detector. The XTI is capable of detecting X-rays in the 0.2–12 keV band, with a minimum energy resolution of 180 eV at 6 keV and 80 eV at 1 keV, from the 1,500 separate energy channels of the CCD detector. The time resolution is an unprecedented 40 ns, and the effective area is 0.23 m^2 (about the same as two *RXTE* PCUs) at 2 keV. Luckily, telemetry is not a concern, so there are no separate data modes for the observations; the photon count event lists have 40 ns time resolution and CCD energy resolution simultaneously. Furthermore, *NICER* does not have pile-up issues, so it can observe bright sources, such as MAXI J1535–571 (see Chapter 5) at 17,000 counts per second (Gendreau et al. 2017).

NICER is a dedicated mission for studying neutron stars. One of its primary science goals is to precisely measure the pulse profiles of persistent X-ray pulsars, and use this to constrain the neutron star matter equation of state. This is very

similar to what was described in Section 1.3.1, except for a difference in emission mechanism. The other component of the mission, SEXTANT, is devoted to precisely clocking X-ray pulsars for use in interplanetary navigation. In addition to the primary science, there is a wide array of observatory science possible with *NICER*, made possible by *NICER*'s energy range sampling the peaks of the accretion disk multicolour blackbody, neutron star surface blackbody and boundary layer, fine energy resolution around the fluorescent iron line, all coupled with its truly remarkable time resolution. Chapter 5 uses data from *NICER* taken in the autumn of 2017 to study a luminous new black hole transient.

1.4.3 Future missions

Continuing with the trend of *RXTE* and *NICER*, the instrumentation goal of X-ray telescopes for spectral-timing analysis is to have a larger effective area and higher spectral resolution, without losing information due to pile-up distortions or throughput limitations. One major proposed mission that ultimately did not go forward is *LOFT*, the *Large Observatory For X-ray Timing*, which was an ESA M3 and M4 candidate (Feroci et al. 2012). *LOFT* brought together the study of neutron star dense matter and accretion in strong gravity as the two primary science cases for an 8.5–10 m^2 X-ray observatory. While it did not pass the final down-select, much of the instrumentation and analysis expertise from *LOFT* has laid the groundwork for the currently proposed missions *eXTP* and *STROBE-X*. A common feature of *eXTP* and *STROBE-X* are the large-area panels of silicon drift detectors for observing hard X-rays that were pioneered for astrophysical use by the *LOFT* team. Chapter 3 includes simulations that represent *eXTP*- or *STROBE-X*-like hard X-ray data.

eXTP

The *Enhanced X-ray Timing Polarimeter* (*eXTP*) is a proposed mission to the Chinese Academy of Sciences (Zhang et al. 2016). It grew out of partnering the previously proposed Chinese mission *XTP* with some of the large-area panels and the wide-field monitor from *LOFT*. The two instruments ideal for our applications are the Large-Area Detector (LAD) and the Spectroscopic Focussing Array (SFA). The LAD would be a hard X-ray detector in the range 2–30 keV (possibly up to 80 keV for bright transients) with an energy resolution of 250 eV at 6 keV. It would have an effective area of 3.4 m^2 at 6 keV, and a time resolution of 1 μs. The SFA would be a soft X-ray detector in the range 0.5–20 keV, with 180 eV energy resolution at 6 keV. Its effective area would be 0.6 m^2 at 6 keV and 0.9 m^2 at 1–2 keV, with 10 μs time resolution. Since the instruments would be pointing together, we would get simultaneous coverage from both instruments for each pointing. The primary science case for *eXTP* heavily emphasizes variability from accreting stellar-mass black holes and neutron stars.

In addition, *eXTP* has the Polarimetry Focusing Array (PFA), a gas pixel-based

detector for linear polarization in the 2–10 keV energy band. The time resolution of the PFA is 500 μs, and the effective area is 0.025 m^2 at 2 keV. While the research in this thesis does not cover X-ray polarization, having spectral-timing-polarization measurements opens up an entirely new axis of analysis.

STROBE-X

The *Spectral- and Time-Resolved Observatory for Broadband Energy X-rays* (*STROBE-X*) is a proposed mission to NASA (Wilson-Hodge et al. 2017). It is currently undergoing concept study in preparation for the NASA Astro2020 Decadal Survey. The science case and instrumentation setup for *STROBE-X* comes from combining hard X-ray *LOFT* panels with a 'super *NICER*' soft X-ray detector scheme.

Of the three instruments in the nominal configuration for *STROBE-X*, the Large-Area Detector (LAD) and XRCA (X-Ray Concentrator Array) are best suited to our research. The LAD would have 200–240 eV energy resolution across its 2–30 keV bandpass. It would have an effective area of 7.6 m^2 at 10 keV, and 10 μs time resolution. The XRCA would have an energy range of 0.2–12 keV with 85–130 eV energy resolution, 100 ns time resolution, and 3.4 m^2 of effective area at 1.5 keV. As with *eXTP*, a large part of *STROBE-X*'s primary science case focuses on variability from accreting stellar compact objects.

1.5 Outline of this Thesis

In Chapter Two, we introduce our new technique for phase-resolved spectroscopy of QPOs and apply it to the Type B low-frequency QPO from the black hole GX 339–4. This research shows how QPO-phase-resolved spectroscopy can provide a physical interpretation for previously intriguing but unaccountable spectral-timing behaviour in a very well-studied source, and start to break degeneracies between physical models for low-frequency QPOs. The data for this chapter comes from *RXTE* observations in 2010.

In Chapter Three, we simulate thermonuclear burst oscillation pulse profiles from a neutron star and fit them with an evolutionary optimization algorithm for the first time in the literature. This work shows how new-to-the-field optimization methods are able to efficiently provide constraints in highly degenerate parameter spaces. The simulations in this chapter represent *LOFT*/*STROBE-X*- and *RXTE*-quality data.

In Chapter Four, we apply our phase-resolved spectroscopy technique to the lower kHz QPO in the neutron star 4U 1608–52. This proof-of-principle chapter shows that kHz QPOs can be easily phase-resolved with our technique from Chapter 2. The data for this chapter comes from *RXTE* observations in 1996.

In Chapter Five, we carry out spectral-timing analysis of the Type B low-frequency QPO in the new transient black hole MAXI J1535–571. With this work, we showcase

how momentous *NICER*'s capabilities are for spectral-timing of compact objects. The data for this chapter comes from *NICER* observations in 2017.

2

Phase-Resolved Spectroscopy of Type B Quasi-Periodic Oscillations in GX 339–4

A. L. Stevens and P. Uttley

Monthly Notices of the Royal Astronomical Society, 2016, 460(3), 2796-2810

Abstract

We present a new spectral-timing technique for phase-resolved spectroscopy and apply it to the low-frequency Type B quasi-periodic oscillation (QPO) from the black hole X-ray binary GX 339–4. We show that on the QPO time-scale the spectrum changes not only in normalisation, but also in spectral shape. Using several different spectral models which parameterise the blackbody and power-law components seen in the time-averaged spectrum, we find that both components are required to vary, although the fractional rms amplitude of blackbody emission is small, ~ 1.4 per cent compared to ~ 25 per cent for the power-law emission. However the blackbody variation leads the power-law variation by ~ 0.3 in relative phase (~ 110 degrees), giving a significant break in the Fourier lag-energy spectrum that our phase-resolved spectral models are able to reproduce. Our results support a geometric interpretation for the QPO variations where the blackbody variation and its phase relation to the power-law are explained by quasi-periodic heating of the approaching and receding sides of the disk by a precessing Comptonising region. The small amplitude of blackbody variations

suggests that the Comptonising region producing the QPO has a relatively large scale-height, and may be linked to the base of the jet, as has previously been suggested to explain the binary orbit inclination-dependence of Type B QPO amplitudes.

2.1 Introduction

The origin of quasi-periodic oscillations (QPOs) in X-ray binaries is still an enigma. They appear as Gaussian- or Lorentzian-shaped features in the averaged power spectrum of a light curve. QPOs in black hole X-ray binaries have been observed at $0.01 - 450\,\mathrm{Hz}$, and are broadly classified into two types, low-frequency (LF; ~ 0.1 to tens of Hz) and high-frequency (HF; ~ 100 Hz or more) (van der Klis 2005; Remillard & McClintock 2006).

Due to the different timescales of variability, it is difficult to define a single mechanism responsible for the whole observed frequency range of QPOs. HFQPO frequencies correspond to the dynamical timescale in the inner accretion disk of the X-ray binary, which suggests that they are related to the Keplerian velocity of the accretion flow (Strohmayer 2001; Kluzniak & Abramowicz 2001). LFQPOs correspond to longer timescales, and have garnered explanations based on either a varying intrinsic luminosity or a varying emission geometry. Examples of LFQPO physical models include seismic oscillations in the accretion disk (Nowak & Wagoner 1991), shocks in the accretion flow (Chakrabarti 1996), a precessing inner hot accretion flow (Stella & Vietri 1999; Ingram et al. 2009), nodal modulation of dynamical disk fluctuations (Psaltis & Norman 2000), gravity waves in the accretion disk (Titarchuk 2003), and intrinsic variability in the base of a jet (Giannios et al. 2004).

In the past 10 years there has been increasing evidence suggesting that LFQPOs have a geometric origin. Schnittman et al. (2006) found that for a sample of 10 sources there is a clear correlation between binary orbit inclination and QPO amplitude, as predicted by their precessing ring model (which was motivated by the first sub-QPO-cycle spectroscopy by Miller & Homan 2005). This result was later expanded upon for a much larger sample of observations by Motta et al. (2015), who determined that the observed LFQPO amplitude has a statistically significant dependence on the orbital inclination (see also Heil et al. 2015a).

Studying power spectra alone has not provided the ability to distinguish between theoretical models. By combining energy spectral and timing information simultaneously so that spectroscopy can probe the QPO variability timescale, we are able to look at the causal relationships between different spectral components and consider whether the QPO is caused by accretion rate fluctuations or geometric changes. For example, the precessing inner accretion flow model would give an observable phase relationship between the blackbody and power law emission due to varying illumination and heating of the accretion disk by the hot inner flow.

Previous spectral-timing methods used to study QPOs such as rms spectra (e.g., Sobolewska & Życki 2006, Gao et al. 2014, Axelsson & Done 2016) found that the power law component was likely responsible for the QPO emission. These methods give the amplitude of variability as a function of energy, but do not incorporate phase information about different energy spectral components. For broadband noise

components, conventional cross-spectral lag studies are used to determine the phase relationship as a function of energy and frequency (Uttley et al. 2011). However, this type of signal is not appropriate for time-domain phase-resolved study, since broadband noise is not coherent on any timescale. For periodic and quasi-periodic variability, it is useful to attempt phase-resolved spectroscopy.

In this paper we introduce a novel spectral-timing technique for phase-resolved spectroscopy of short-timescale variability, specifically applied here to Type B LFQPOs from the black hole binary GX 339–4. Type B QPOs are characterised by a low level of broadband noise (correlated noise in the lowest frequencies of the power spectrum), so most of their variability power is contained in the QPO. This makes Type B QPOs the cleanest QPO signal for our analysis. Our technique uses a combination of the cross-correlation function, energy spectroscopy, and simulated lag-energy spectrum to find a model that can explain the data in both the energy and Fourier domains. In the discussion we interpret our QPO-phase-resolved spectroscopy results through the lens of a geometric precession LFQPO model, but we emphasise that this model-independent technique can be used to test any physical model that predicts spectral changes on the variability timescale. We conclude with a discussion of the results and their interpretation, with a forward look to how combining our new technique with data from new missions will be able to further constrain models for LFQPOs, which can constrain how matter behaves in strong gravitational fields.

2.2 Data and Basic Spectral-Timing Behaviour

GX 339–4 is a black hole candidate X-ray binary (Hynes et al. 2003) with a black hole of lower mass limit $\sim 7\,M_\odot$ (Muñoz-Darias et al. 2008) and which may possess near-maximal spin (Ludlam et al. 2015). The system also likely has a low binary orbit inclination ($\sim 40°$; Muñoz-Darias et al. 2013). The source exhibits X-ray variability on a variety of timescales. In the 2010 outburst (Yamaoka et al. 2010), GX 339–4 was observed for some time in the soft intermediate spectral state, which showed strong Type B QPOs (Motta et al. 2011).

2.2.1 Data

For this analysis we used data from NASA's High Energy Astrophysics Science Archive Research Center (HEASARC) taken with the Proportional Counter Array (PCA) onboard the Rossi X-ray Timing Explorer (RXTE) in 64-channel event mode with $122\,\mu$s time resolution (`E_125us_64M_0_1s`).

The following filtering criteria were used to obtain Good Time Intervals (GTIs) for analysis: Proportional Counter Unit (PCU) 2 is on, two or more PCUs are on, elevation angle $> 10°$, and target offset $< 0.02°$. Time since the South Atlantic Anomaly passage was not filtered on. The nine observation IDs with events fitting these criteria

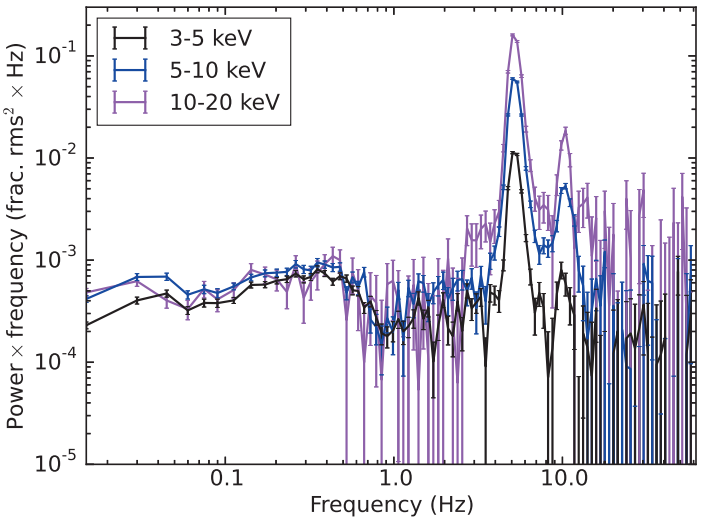

Figure 2.1: Power spectra in different energy bands, averaged over all available PCUs in the filtered event-mode data. The energy bands used are $3-5$ keV (rms = 7.7%; mean rate = $1627.4\,\mathrm{cts\,s^{-1}}$), $5-10$ keV (rms = 13.9%; mean rate = $872.3\,\mathrm{cts\,s^{-1}}$), and $10-20$ keV (rms = 23.1%; mean rate = $183.9\,\mathrm{cts\,s^{-1}}$). The power spectra were geometrically re-binned in frequency with a binning factor of 1.06. There is a Type B QPO with a centroid frequency $\nu_{\mathrm{centroid}} \simeq 5.2$ Hz and a harmonic just above 10 Hz. Individual observations have been shifted in frequency to correct for QPO centroid variations, as explained in the text.

are: 95335-01-01-05, 95335-01-01-07, 95335-01-01-00, 95335-01-01-01, 95409-01-18-00, 95409-01-17-05, 95409-01-18-05, 95409-01-17-06, and 95409-01-15-06.

The data were initially binned to 7.8125 ms time resolution ($64\times$ the intrinsic resolution of the data) and divided into 64 s segments, giving 8192 time bins per segment. After rejecting the segments with a negative integrated noise-subtracted power (6 segments in total), there were 198 good segments of data over the nine observations, or an exposure time of 12.672 ks. Note that the Type B QPO appears to sharply switch on and off on timescales of less than a few minutes (e.g., Belloni et al. 2005 figures A.3 and A.4), so is not expected to always contribute to the light curve, hence the segments with no detectable signal (and negative integrated power) are discarded.

2.2.2 Power spectra

We computed the average power spectrum in three energy bands (Figure 2.1) using our own code.[1] The power spectra were geometrically re-binned in frequency with a binning factor of 1.06. A power-law model of the noise level was fitted to the higher frequency power of each spectrum (> 25 Hz, by a chi-squared minimisation with a Levenberg-Marquardt algorithm) then subtracted. There is a strong Type B QPO with a centroid frequency $\nu_{\rm centroid} = 5.20$ Hz (as determined by a chi-squared minimisation of a Gaussian profile to the average power spectrum over all energy channels) and a quality factor $Q = 6.6$ ($Q = \nu_{\rm centroid}/{\rm FWHM}$), and weak broadband noise below 1 Hz. In the literature it is common to fit QPOs with a Lorentzian profile, but since these Type B QPOs have a relatively smooth peak, they are better fit with a Gaussian.

The QPO centroid frequency shifts slightly between observations ($\nu_{\rm centroid}$ ranges from 4.87 Hz to 5.65 Hz). The frequency shift smears out the QPO signal averaged over all observations, and thus reduces our signal to noise. To combat this, we artificially shifted the QPO centroid frequencies obtained from each *RXTE* ObsID to line up at the centroid of the unaltered average power spectrum (5.20 Hz). To line up the centroid frequencies, we kept 8192 bins for each segment of data but adjusted the segment lengths, so that the width of the time bins dt changed. The segments were adjusted by the same amount per observation. After adjusting, the average dt over all observations is 8.153 ms, which gives a Nyquist frequency of 61.33 Hz and a Fourier-space frequency resolution of 0.015 Hz, and changes the exposure time to 13.224 ks. This adjustment has the same effect as the shift-and-add technique commonly used when averaging multiple power spectra of neutron star kHz QPOs (Méndez et al. 1998a). The adjustment also corrects our phase-resolved analysis (in Section 2.3) to ensure that we compare relative phases, allowing for more accurate phase-resolved spectroscopy by forcing the same number of QPO cycles in one segment of data.

2.2.3 Lag-energy spectrum

We also compute a lag-energy spectrum for the QPO (by averaging the cross-spectrum in the $4-7$ Hz range), using our own lag spectrum code [1] following the outline in Uttley et al. (2014). Using the same approach described in Section 2.3.2, we measure the lags for event mode energy channels from PCU2, relative to a reference band which includes the counts in the 3–20 keV range from all other available PCUs (excluding PCU2). The lag-energy spectrum is plotted in Figure 2.2. It crosses the zero lag mark (dashed line) at the average energy of the reference band variations. The negative time lags denote the energy-dependent variations that lead the reference band variations, and positive time lags denote the variations that lag the reference band.

[1] See Appendix 2A for URL.

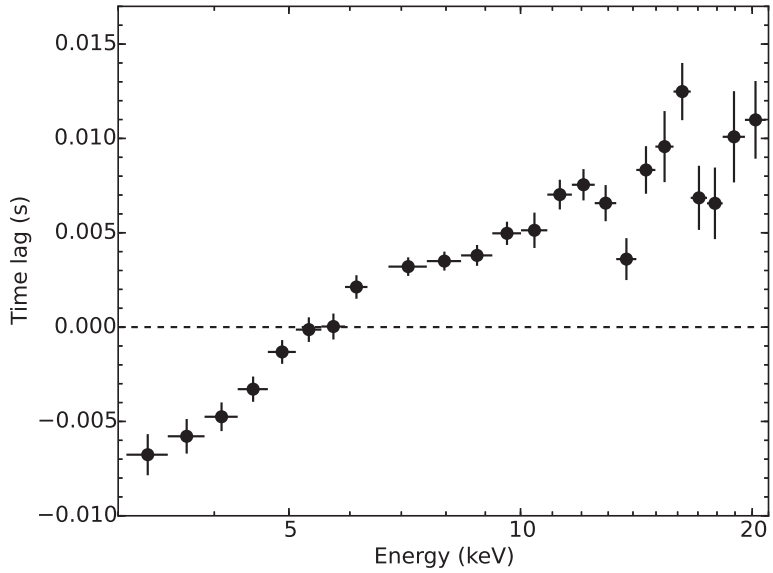

Figure 2.2: Time lag obtained from the cross-spectrum averaged over $4 - 7$ Hz, plotted versus energy. The dashed line indicates no time lag with respect to the variability in the reference band. Positive lags indicate energies which lag the reference band while negative lags lead the reference band. There is a clear break in the slope of the lag-energy spectrum at around 6 keV. The PCA has zero counts in detector channel 11 when in event-mode, hence the gap in the data at ~ 6.5 keV.

A flat or smooth lag-energy spectrum would indicate a simple evolution of the continuum or one spectral component over the averaged timescale. However, in Figure 2.2 there is a bump or break at ~ 6 keV. This indicates a more complex spectral behaviour, such as a causal relationship between separate spectral components (Uttley et al. 2011). Investigating the cause of this lag behaviour requires a more comprehensive analysis of individual energy spectra at different phases of the relative QPO cycle, or phase-resolved spectroscopy.

2.3 Phase-Resolved Spectroscopy Technique

One of the ways to probe the physics and geometry of black hole binaries is via X-ray spectroscopy. Traditionally this is done by analysing the mean spectrum of a source over a set of observations, or per state of a transient outburst. However, even looking at spectra per observation only shows the overall spectral trends, as these timescales are much larger than those of the variability processes. A short segment of data would be too noisy, due to insufficient counts, for drawing detailed conclusions (see however Skipper & McHardy 2016 for a more general approach using short-

term spectral variability). The most desirable solution is to carry out sub-variability-timescale spectroscopy, or effectively phase-resolved spectroscopy for quasi-periodic signals.

For periodic signals, one can do phase-resolved spectroscopy by phase-folding the variations to stack the light curves (e.g., Wilkinson et al. 2011). However, this approach requires exactly periodic signals from sources with known ephemerides, which is not appropriate for QPOs (*quasi*-periodic phenomena by definition).

Previous approaches for spectral-timing of QPOs have been carried out in the time domain. Miller & Homan (2005) used a bright source to extract spectra at maxima and minima of a QPO waveform in the broadband light curve. Recently, Ingram & van der Klis (2015) developed a method to reconstruct a QPO waveform for relatively-high-count-rate light curves per narrow energy band, then selecting different times in a QPO cycle for all energy bands to create and compare energy spectra. These methods yielded new discoveries, but are only applicable to bright sources with a well-defined QPO waveform. Here, we develop a more general approach that gives some of the benefits of phase-resolved spectroscopy without the requirements of a periodic signal or high count rate, using the cross-correlation function.

2.3.1 Approximating phase-resolved spectra

Consider correlated light curves defined at every time k with a sinusoidal modulation at the same angular frequency ω. A narrow energy band light curve x is determined for every discrete energy E_i with an energy-dependent amplitude $a(E_i)$ and phase $\psi(E_i)$, while a reference band light curve y for a much broader energy range has amplitude $a_{\rm ref}$ and phase $\psi_{\rm ref}$. Both have noise components $n(E_i)$ and $n_{\rm ref}$, respectively. The light curves are then expressed as:

$$x_k(E_i) = a(E_i)\,\sin\left(\omega\, t_k + \psi(E_i)\right) + n_k(E_i) \tag{2.1}$$

$$y_{{\rm ref},k} = a_{\rm ref}\,\sin\left(\omega\, t_k + \psi_{\rm ref}\right) + n_{{\rm ref},k} \tag{2.2}$$

At the Fourier frequency corresponding to ω, the cross spectrum C with the reference band for each energy channel is:

$$C_{E_i,{\rm ref}} = A(E_i)\,A_{\rm ref}\,\exp\left(i\,[\psi(E_i) - \psi_{\rm ref}]\right) + C_{\rm noise} \tag{2.3}$$

If the absolute rms-squared normalisation is used (e.g. see Uttley et al. 2014) the Fourier amplitudes $A(E_i)$ and $A_{\rm ref}$ scale with the rms values of the sine wave in the energy channel and reference band. Note that a similar term is produced at the corresponding negative frequency in the Fourier transform, which must also be kept for the next step. The noise term is produced by the products of the Fourier transforms of noise components with each other and with those of the sinusoidal signal.

The cross-spectrum is a Fourier pair with the cross-correlation function, or CCF, so that if we inverse Fourier transform the cross-spectrum we obtain the CCF per

time-delay bin τ_l for each energy channel E_i:

$$CCF_{E_i,\mathrm{ref}}(\tau_l) = \mathcal{F}^{-1}\left(C_{E_i,\mathrm{ref}}\right) \tag{2.4}$$

$$= a(E_i)\, a_{\mathrm{ref}}\, \sin\left(\omega\, \tau_l + [\psi(E_i) - \psi_{\mathrm{ref}}]\right) \tag{2.5}$$

$$+ CCF_{\mathrm{noise}}(\tau_l) \tag{2.6}$$

After calculating the CCF, we normalise by the noise-subtracted integrated rms of the reference band, to divide out the reference band normalisation a_{ref}. Thus, the CCF of the signal component is of the same form as the energy-dependent and sinusoidal original light curve of interest x, but with the phase defined as being relative to the reference band. In this way, we can recover the *phase-resolved* energy dependent signal.

The CCF of the noise component is statistically independent between time segments used to calculate the CCF, and has a mean of zero. Thus, averaging many CCFs obtained from identical-length light curve segments can be used to suppress the noise. However, it is important to bear in mind that the noise in the CCF is not in general statistically independent between CCF time-delay bins, which complicates the interpretation of the errors using our approach, as discussed in the next section.

The above demonstration is shown for sinusoidal signals, whereas the Fourier pair of a QPO signal more closely resembles a damped sinusoid, so some caution needs to be applied in interpreting the energy-resolved CCF in terms of phase-resolved spectra. We therefore use this approach mainly as a guide to determine the best way to model the energy-dependent QPO signal in terms of variable spectral components. However, in Section 2.4.5 we show that we can independently re-create the conventional lag-energy spectrum of the data by sinusoidally varying energy spectral parameters as inferred from our CCF method, which suggests that our approach is valid for deriving QPO phase-resolved spectra.

2.3.2 Cross correlation

Using the CCF we can isolate the characteristic QPO features in the time domain by correlating with a higher count-rate light curve from a broad reference energy band. By comparing different energy channels ("channels of interest") with the same reference band, we can connect the QPO's relative flux changes using energy channel resolution. The CCF gives the average QPO signal of a channel of interest correlated with and relative to the reference band.

Central to this method is establishing narrow energy channels of interest and a broad reference energy band, as outlined for cross-spectral analysis in Uttley et al. (2014). For our *RXTE* PCA data, the channels of interest are taken from PCU2, as it has a well-established calibration and is switched on for the greatest number of observations. Since the data are filtered to have at least two PCUs on, the broad reference band comes from any other PCUs that are on (i.e. excluding PCU2, so that

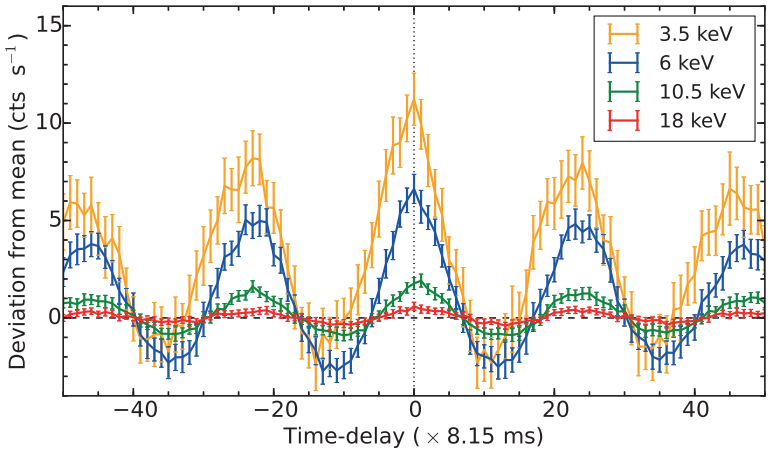

Figure 2.3: The CCF in four energy channels (corresponding to 3.5, 6, 10.5, and 18 keV), normalised to have units of count rate. The CCF has a larger amplitude in the lower energy channels, but the lower energy channels also have higher mean count rates.

the reference band and channel of interest light curves are statistically independent). The reference band corresponds to the energy range $3 - 20$ keV (absolute channels $6 - 48$, inclusive); this range ensures plenty of photons for optimal signal-to-noise, and fully covers the energy range used in spectral fitting in Section 2.4.3.

The time binning of the data is first adjusted per observation as explained in Section 2.2.2, and the cross spectrum is computed per energy channel per segment of data. Applying an inverse DFT to each cross spectrum yields the CCF in each channel of interest per segment of data. The CCFs are then averaged together in the time domain per energy channel over all segments. We then normalise the averaged CCF by $2/(K\,\sigma_{\rm ref})$, where $\sigma_{\rm ref}$ is the absolute-normalised integrated rms of the averaged power spectrum of the reference band, and K is the number of time bins per segment (as in Uttley et al. 2014). This gives the CCF units of count rate as deviations from the mean count rate per channel of interest, and corrects for the mixed PCUs in the reference band. The error is calculated for each time-delay and energy bin from the standard error on the mean CCF in that bin.

The CCF in four energy channels is shown in Figure 2.3. Since cross-correlation phase-locks the features of the signal in the channels of interest to the signal in the reference band, it produces a relatively strong signal modulation, even with a relatively low count rate in each energy channel.

This CCF phase-resolved spectroscopy technique does not make assumptions about an underlying QPO waveform, and is applicable independent of the QPO emission mechanism. Thus it is a powerful method to test physical QPO models that predict specific spectral-timing behaviour. Since the reference band and channels of interest

overlap in energy space and are taken simultaneously from independent detectors, we assume that if a signal is present in the light curves, there will be some degree of correlation between the channels of interest and the reference band. The Poisson noise components are uncorrelated, so the resulting CCF contains the (correlated) QPO signal, with less noise than the input light curves. Our technique yields a count rate per energy channel in each time-delay bin, as deviations from the mean count rate per energy channel. After adding the CCF to the mean count rate per energy channel, these effectively-phase-resolved spectra can be loaded into a spectral fitting program like XSPEC.

As already briefly mentioned in the previous section, statistical fluctuations between CCF time-delay bins for a given energy channel are correlated with one another. This effect arises from the component of the noise resulting from the multiplication of the Poisson noise by the intrinsic signal, resulting in 'structure' in the noise across phase bins, the strength of which depends on the shape of the intrinsic signal power spectrum (see Bartlett 1955; Box & Jenkins 1976 for a detailed description of the errors in the CCF associated with autocorrelated processes). Thus the phase-resolved spectra produced by our method are not strictly independent in time, although they are independent in energy (since different energy bins have independent Poisson noise terms). Therefore one cannot easily interpret the results of conventional fitting techniques, such as χ^2, which are used to simultaneously fit the spectra from multiple phase bins. To compensate, we will implement bootstrapping (Section 2.4.4) and lag-energy spectral comparison (Section 2.4.5) to determine accurate errors on our description of the modulation of spectral parameter values and check the appropriateness of the spectral model.

2.4 Results

2.4.1 Energy-dependent CCF

The CCF can be plotted in a two-dimensional colour plot (Figure 2.4), to show the simultaneous time variability and energy spectra. The colour map shows the ratio of the CCF amplitude to the mean count rate of each energy channel, because the lower energy channels have larger CCF amplitudes but also higher mean count rates. The energy range shown along the y-axis, $3 - 20$ keV, is the same range used for the reference band in Section 2.3 and for spectral fitting in Section 2.4.3. Although the mean count rates are higher at lower energies, the amplitude of the QPO signal in the CCF is largest at $\sim$10 keV. The CCF of the QPO has the approximate shape of a damped sinusoid in each energy channel, so henceforth we will refer to a QPO phase in the same way that one can consider a sinusoidal phase. As noted in Section 2.2.2, we correct for variations in the QPO centroid frequency between observations by artificially adjusting the length of the light curve segments, such that the same

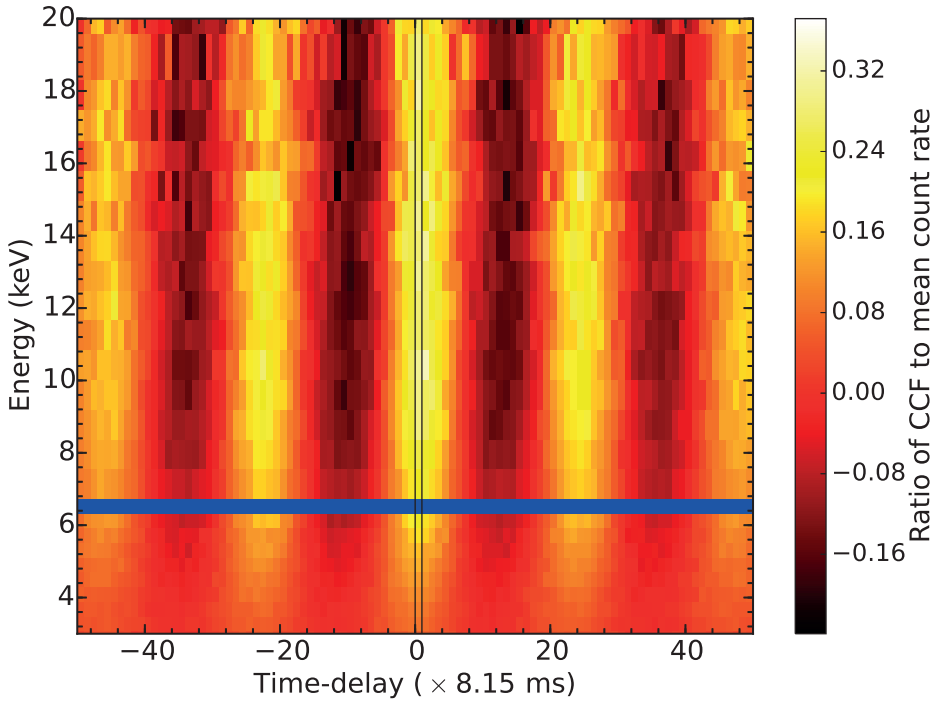

Figure 2.4: Cross-correlation function (CCF) plotted in two dimensions. The colour map shows the ratio of the CCF amplitude in each bin to the mean count rate of the energy channel. The zero time-delay bin is indicated by the solid black outline. The PCA has zero counts in detector channel 11 when in event-mode, hence the gap in the data at ∼6.5 keV as indicated by the blue bar.

number of QPO cycles are obtained in each data segment used to calculate the CCF. In doing this and then interpreting the results as a phase-resolved spectrum, we implicitly assume that time-delays between energy bands are constant in phase and not constant in time, so that the intrinsic phase relationships between energy bands are preserved and not blurred out by the adjustment process. This approach will be borne out by the success of our phase-resolved spectral model in reproducing the shape of the lag-energy spectrum, whether it is adjusted for QPO frequency changes or not (see Section 2.4.5).

2.4.2 Time-averaged energy spectrum

The time-averaged spectrum obtained from the *RXTE* PCA Standard-2 data mode, shown in Figure 2.5, was used for identifying the general energy spectral features and considering which spectral models to fit for the phase-resolved data. This data mode has finer energy binning than the high time resolution event-mode data used for the

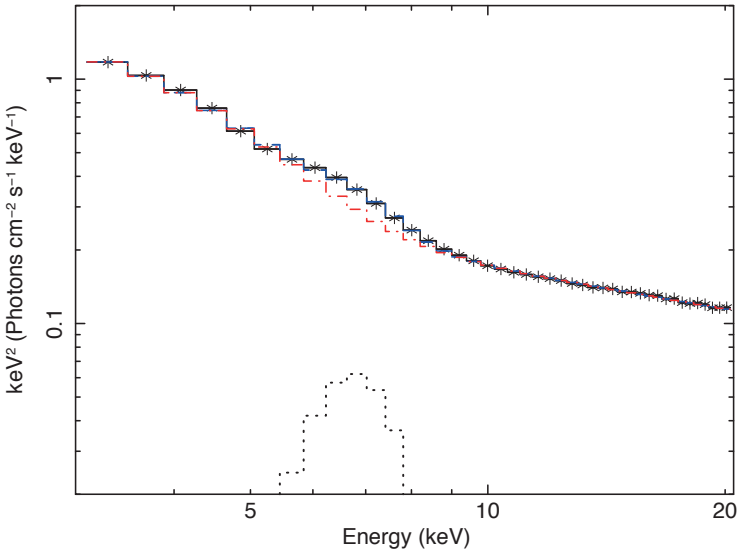

Figure 2.5: The time-averaged spectrum of the data in Standard-2 data mode, unfolded through the instrument response matrix and fit with the model PHABS × (SIMPL * DISKBB + GAUSS). The solid black line with star points represents the data, the dashed blue line represents the total model, the dash-dotted red line is SIMPL convolved with DISKBB (since it is a convolution model, there is only one output curve), and the dotted line is GAUSS representing the iron Kα emission line. This is a simple phenomenological spectral fit, with $\chi^2/\,\mathrm{d.o.f.} = 229.2/37$.

rest of the analysis.

The spectrum shows a soft blackbody-like component, a hard power law component, an iron Kα emission line, and neutral hydrogen absorption. We fit this with PHABS × (SIMPL * DISKBB + GAUSS), a simple phenomenological model containing the key components. While the presence of the iron line indicates that a broad reflection continuum should also be included, we expect that there would be degeneracies between an upscattered power law and the reflection components, particularly for the coarser energy binning of event-mode spectra. Since our main interest was the possible variability of the soft and hard spectral components and any phase delays between them, any reflection variability was accounted for with the power law and Gaussian components. As later noted, we find that the Gaussian iron line is not required by our fits to vary on the QPO timescale so we do not expect variable reflection to contaminate the inferred variation of other spectral components.

2.4.3 CCF phase-resolved spectroscopy

To carry out phase-resolved spectroscopy using the energy-dependent CCF, we select energy spectra at each time-delay bin of the CCF covering one QPO cycle. Since we

focus on relative phases instead of absolute phases for simplicity, we define the QPO phase to be zero when the CCF in channel 15 is at a minimum; this is in time-delay bin -10. By selecting until the next minimum in channel 15 at time-delay bin +13, we have 24 time-delay bins in total for one QPO cycle. This time range of data covers the largest maximum and minimum in CCF amplitude so that we can see the most prominent differences in energy spectra. The instrument response matrix uses the calibration of PCU 2, since the channels of interest are taken from that detector.

Figure 2.6a shows the deviations from the mean spectrum at four QPO phases, plotted in $E\,F(E)$ units. Figure 2.6b shows the total energy spectra at the same four QPO phases, which are the deviations shown in the upper panel added to the mean spectrum. For visualisation here, the energy spectra have been "fluxed", i.e., divided by the effective area of the instrument response (as in Vaughan et al. 2011). In this figure, the QPO phase of $\sim 0°$is at time-delay bin -10, $\sim 90°$is at time-delay bin -5, $\sim 180°$is at time-delay bin +1, and $\sim 270°$is at time-delay bin +6. Time-delay bins within one QPO cycle will be referred to as phase bins. In the remainder of this paper, we only focus on the spectral-parameterisation of variations during the central, highest-amplitude cycle seen in the CCF. However, we note here for completeness that an analysis of additional cycles yields consistent parameter variations (with similar phase relations to the strongest cycle), albeit with decaying amplitudes of variability, as expected due to the damped sinusoidal shape of the CCF.

For the full phase-resolved spectroscopy, we simultaneously fit energy spectra from time-delay bins -10 to +13 (inclusive, and including 0), giving 24 spectra in total; the 5.2 Hz centroid frequency corresponds to a timescale of 0.192 s, so the 8.153 ms time binning gives 23.5 phase bins per QPO cycle (rounded to 24, to give complete coverage of a QPO cycle). Our aim is to obtain the simplest possible parameterisation of the QPO spectral variability in terms of a spectral model with variable components. We therefore systematically untied and froze different combinations of parameters to find an optimal fit without over-fitting, such that the resulting model makes physical sense. We tested spectral models which all include constant neutral absorption, a soft blackbody component, a hard power law component and an iron Kα line. These components are clearly required by the time-averaged spectrum (Figure 2.5).

The power law is likely to be produced by Compton-upscattered disk photons. To model this component, we use SIMPL. SIMPL is a convolution model intended as a simple parameterisation of Comptonisation from Compton-thin plasmas by scattering a fraction ($f_{\rm scatt}$) of the input seed spectrum to make a power law of photon index Γ (Steiner et al. 2009). For a given, constant seed spectrum, $f_{\rm scatt}$ is thus a proxy for the normalisation of the power law component. In this spectral model, the power law lower-energy cutoff follows the low energy shape of the seed spectrum, which in our case is a multi-colour disk blackbody. SIMPL outputs the sum of the scattered power law and the unscattered seed photon spectrum, so it contains a built-in anticorrelation between the normalisation of the power law and the normalisation of the observed

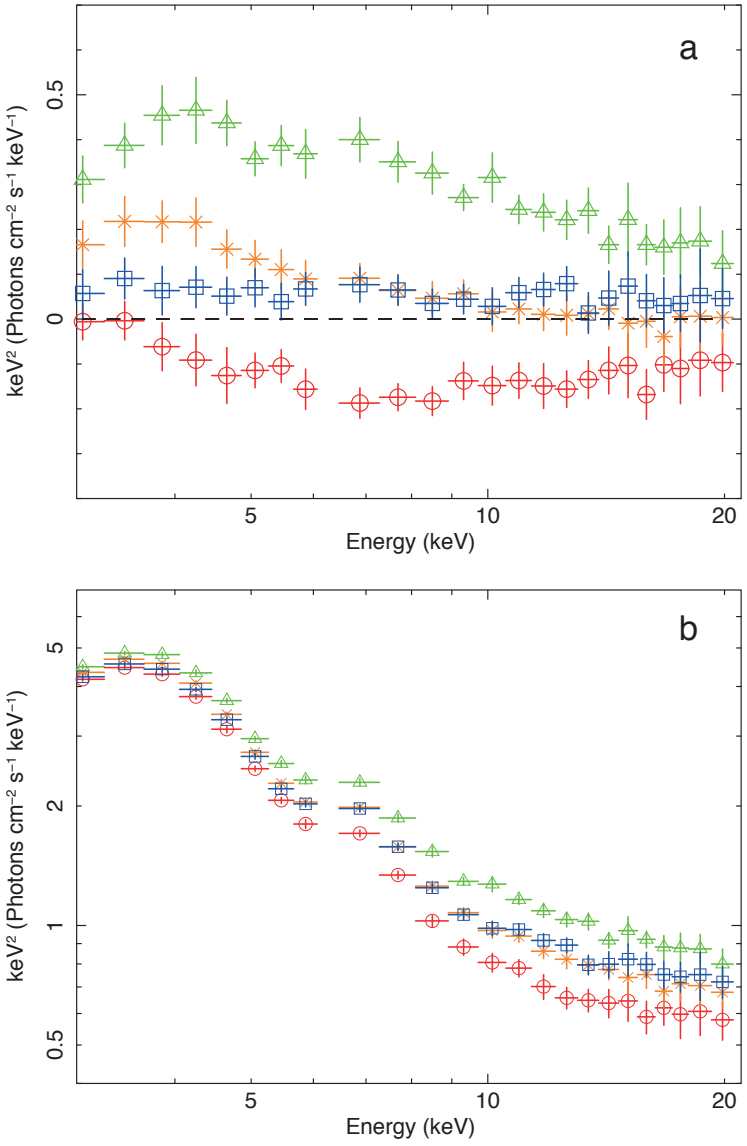

Figure 2.6: a: Deviations from the mean spectrum at four QPO phases discussed in the text. The black dashed line denotes zero deviation from the mean spectrum. The red circles are for a QPO phase of 0°, orange X's for 90°, green triangles for 180°, and blue squares for 270°. The PCA has zero counts in detector channel 11 when in event-mode, hence the gap in the data at ~ 6.5 keV. It is evident that the shape of the spectrum, in addition to the normalisation, changes with QPO phase. **b:** Energy spectra at the same four QPO phases, generated by adding the above deviations to the mean spectrum. Data points have the same meaning as in the top panel. For visualisation in both plots, the spectra have been divided by the instrument's effective area as mentioned in the text.

part of the seed spectrum. To remove this effect, we created a new model SIMPLER where the output seed spectrum is the same as the input. This would correspond to the case where the Compton scatterer does not obscure the observer's view of the seed spectrum, as would be expected for an inner hot flow that is physically separate from the accretion disk. As is the default for SIMPL, only upscattering is accounted for (indicated by freezing the model scattering switch $U = 1$ in Tables 2.2–2.4).

Using XSPEC version 12.8.2 (Arnaud 1996) we fitted the phase-resolved energy spectra in the energy range $3 - 20\,\mathrm{keV}$, ignoring detector channel 11 (which has zero counts in *RXTE* PCA event-mode). The solar abundance table was set to `vern` (Verner et al. 1996) and the photoionization cross-section table was set to `wilm` (Wilms et al. 2000). The systematic error was set to 0.5 per cent.

The models we tested are PHABS × (SIMPLER * DISKBB + GAUSS) ("model 1"), PHABS × (SIMPLER * DISKBB + BBODYRAD + GAUSS) ("model 2"), and PHABS × (SIMPLER * DISKPBB + GAUSS) ("model 3"); they are summarised with their best parameterisations in Table 2.1. These models represent the simplest possible system geometries: a power law-emitting region and a blackbody-emitting accretion disk; a power-law emitting region, a blackbody-emitting accretion disk, and a secondary blackbody from the power-law heating a smaller, single-temperature region; and a power law-emitting region and a blackbody-emitting accretion disk with a different (compared to the standard disk) and possibly varying radial dependence of temperature. Details about the fits and motivations are given in subsections 2.4.3, 2.4.3, and 2.4.3, respectively. The neutral hydrogen column density was frozen to $N_H = 6 \times 10^{21}\,\mathrm{cm}^{-2}$ (Reynolds & Miller 2013).

For carrying out the QPO phase-resolved spectroscopy, we began with every spectral parameter tied across QPO phase. We then systematically stepped through the parameters, untying each individually so that it can vary across phase, and assessing the fit. If untying a parameter gave a lower χ^2 and physically reasonable parameter values, it remained untied and we investigated whether a second parameter should be untied. This continued until there was no significant improvement in χ^2 and/or the parameter values became physically unreasonable.

As mentioned in Section 2.3.2, although the count rates in energy bins within a single phase-resolved spectrum are statistically independent, the phase-resolved spectra are not statistically independent from each other, in that there is some correlation in errors between different phases of the same energy bin. Therefore the χ^2 fit statistic returned by fitting the phase-resolved spectra together does not have the same meaning as a conventional χ^2 statistic. The correlation between phase bins effectively means that ratio between the χ^2 and degrees of freedom is reduced, so we use the following procedure only as a guide to find the best parameterisation of each model, rather than using an explicit goodness-of-fit test. We found during our fits that the plausible best-fitting models converged to $\chi^2/\mathrm{d.o.f.} \sim 130/500$. Assuming that this χ^2 corresponds to a reduced $\chi^2 \sim 1$, a significant improvement in the fit

at the 3-σ level should correspond to a $\Delta\chi^2$ of $46 \times 130/500 \simeq 12$, where 46 is the formal $\Delta\chi^2$ improvement corresponding to a 3-σ significance improvement in a nested model for 23 additional free parameters, i.e. equivalent to freeing a spectral model parameter in our simultaneous fit of 24 spectra. In Sections 2.4.3, 2.4.3, and 2.4.3 we use this $\Delta\chi^2 > 12$ criterion for a significant improvement to investigate which spectral parameters vary with QPO phase. However, in view of the fact that our approach is merely a guide to search for a spectral model parameterisation of the phase-dependent QPO variability, and is not statistically rigorous, in Sections 2.4.4 and 2.4.5 we assess the statistical significance of the parameter variations and their phase lags with bootstrapping and direct comparison of our best-fitting QPO spectral models with the lag-energy spectrum.

Spectral model 1

The first energy spectral model tested was PHABS × (SIMPLER * DISKBB + GAUSS). This is the simplest possible model to explain the features in Figures 2.5 and 2.6, where the blackbody component is intended to capture all possible soft blackbody emission. To test which parameterisation best fits the 24 sequential energy spectra from one QPO cycle, we began with all parameters tied (and $N_{\rm H}$ frozen and the upscattering-only flag set), which gave $\chi^2 = 7114.5$ for 569 d.o.f. For all fits, the power-law photon index Γ was allowed values in the range $1.0 - 3.1$, disk maximum temperature $T_{\rm disc}$ was allowed values in the range $0.5 - 1.0\,$keV, the Gaussian line centroid energy $E_{\rm line}$ was allowed $5.5 - 7.0\,$keV, and all other parameters were allowed the full range of values (unless frozen).

We then untied each parameter one at a time to see which was required to vary given our simple $\Delta\chi^2$ criterion. Untying the scattering fraction $f_{\rm scatt}$ gave a significant improvement to the fit with $\chi^2 = 393.0$ for 546 d.o.f. This is also physically motivated: untying $T_{\rm disc}$, the disk blackbody normalisation $N_{\rm disc}$, or any Gaussian line parameters did not account for the large hard continuum variations clearly seen in Figures 2.4 and 2.6. By next untying Γ, we saw a large improvement in the fit to $\chi^2 = 176.5$ for 523 d.o.f. This changing power law slope is also visibly evident in Figure 2.6a when comparing phases 0° and 180° (the red and green points) above $\sim 6.5\,$keV. A third parameter that significantly improved the fit when allowed to vary with phase is $T_{\rm disc}$, which when untied gave $\chi^2 = 141.4$ for 500 d.o.f. Variations in the blackbody component are required to account for the variations of the lower-energy 'soft excess' feature, seemingly out of phase with the harder power-law continuum, which are seen in Figure 2.6a. Untying $N_{\rm bb}$, $E_{\rm line}$, or $\sigma_{\rm line}$ as the third varying parameter returned worse fits. An untied Gaussian normalisation $N_{\rm line}$ yielded a smaller $\chi^2\,/\,$d.o.f. (132.8/500) than an untied $T_{\rm disc}$, although the improvement was not formally accepted according to our significance criterion. However, examining the returned parameter values and resultant spectra showed that the broadened iron line

#	Energy Spectral Model	Untied	χ^2 / d.o.f.	$\phi_{FS} - \phi_\Gamma$	$\phi_{FS} - \phi_{BB}$	Lag χ^2 / d.o.f.
1	PHABS $\times$ (SIMPLER * DISKBB + GAUSS)	$f_{\rm scatt}$, Γ, $T_{\rm disc}$	143.9/502	0.01 ± 0.02	0.32 ± 0.02	37.9/24
2	PHABS $\times$ (SIMPLER * DISKBB + BBODYRAD + GAUSS)	$f_{\rm scatt}$, Γ, $T_{\rm bb}$	128.4/500	0.019 ± 0.009	0.28 ± 0.02	45.5/24
3	PHABS $\times$ (SIMPLER * DISKPBB + GAUSS)	$f_{\rm scatt}$, Γ, p	126.8/501	0.02 ± 0.01	0.10 ± 0.04	47.2/24

Table 2.1: A summary of the best parameterisation for each spectral model tested. The first column gives the number of each energy spectral model, for reference throughout the paper, and the second column states the models. The "Untied" column indicates which model parameters are untied between phase-resolved spectra to give the best parameterisation (see Sections 2.4.3-2.4.3). The next column gives the χ^2 fit statistic for the energy spectral fits of each parameterisation (assuming the data are statistically independent, which they are not; see Section 2.4.4). The columns "$\phi_{FS} - \phi_\Gamma$" and "$\phi_{FS} - \phi_{BB}$" give the normalised phase difference in the parameter variations between an untied $f_{\rm scatt}$ and untied Γ, and an untied $f_{\rm scatt}$ and an untied blackbody parameter (see Section 2.4.4). The column "Lag χ^2 / d.o.f." gives the fit statistics for lag-energy spectrum simulated from the energy spectral model and fitted to the data (see Section 2.4.5 and Figure 2.11).

component was subsuming the role of the higher temperature end of the blackbody emission. The fit pushed $E_{\rm line}$ to the lowest value of the allowed range (5.5 keV) and gave the line width $\sigma_{\rm line} > 1$ keV, which we believe is unphysical. To mitigate this degeneracy, we froze $\sigma_{\rm line} = 0.97$ in the final parameterisation. The value of $\sigma_{\rm line}$ was found using the XSPEC `steppar` command on the mean event-mode spectrum; the best-fitting value was $\sigma_{\rm line} = 1.0$, so to push the value of $\sigma_{\rm line}$ away from the boundary we decided on the lowest value while keeping within the 1-σ fit region, $\sigma_{\rm line} = 0.97$. In general, the smaller the width of the iron line, the more physically consistent the model fit is. Untying a fourth parameter was not able to both significantly improve the fit and return physically reasonable parameter values.

While the fits somewhat preferred untying the disk blackbody temperature to the disk blackbody normalisation, the two parameters were highly degenerate. To disentangle this degeneracy and ensure that any temperature variations in the bootstrapping were not due to a different normalisation, $N_{\rm disc}$ was frozen at 2505.72 (the best-fitting value).

Component	Parameter	Value	Notes
PHABS	$N_{\rm H}$ ($\times 10^{21}$ cm^{-2})	6.0	Frozen
SIMPLER	Γ	2.52 †	Untied
SIMPLER	$f_{\rm scatt}$	0.156 †	Untied
SIMPLER	U	1	Frozen
DISKBB	$T_{\rm disc}$(keV)	0.8307 †	Untied
DISKBB	$N_{\rm disc}$	2505.72	Frozen
GAUSS	$E_{\rm line}$(keV)	6.28 ± 0.03	Tied
GAUSS	$\sigma_{\rm line}$ (keV)	0.97	Frozen
GAUSS	$N_{\rm line}$	$2.28 \pm 0.08 \times 10^{-2}$	Tied

Table 2.2: Values of the best-fitting parameterisation for model 1 as explained in Section 2.4.3. Mean values are listed for the untied parameters, indicated by †. Errors on the tied parameters were computed for the 90% confidence interval with XSPEC's MCMC error routine. XSPEC gives a fit statistic of $\chi^2 = 143.9$ for 502 d.o.f.

The best parameterisation for model 1 has $f_{\rm scatt}$, Γ, and $T_{\rm disc}$ untied, and $N_{\rm disc}$ and $\sigma_{\rm line}$ frozen. The value of each parameter is listed in Table 2.2, with the mean value quoted for the untied parameters. Note that the χ^2 / d.o.f. in the table caption corresponds to the fit with $N_{\rm disc}$ and $\sigma_{\rm line}$ frozen. The same four spectra shown in Figure 2.6 are also shown in Figure 2.7, unfolded through the instrument response matrix and through models with the best-fitting spectral parameter values. The value of each untied parameter per phase bin is shown in Figure 2.8. The errors on the values of the untied spectral parameters in each bin, and on the phase of their variations, were determined via bootstrapping the data (see Section 2.4.4). The phase

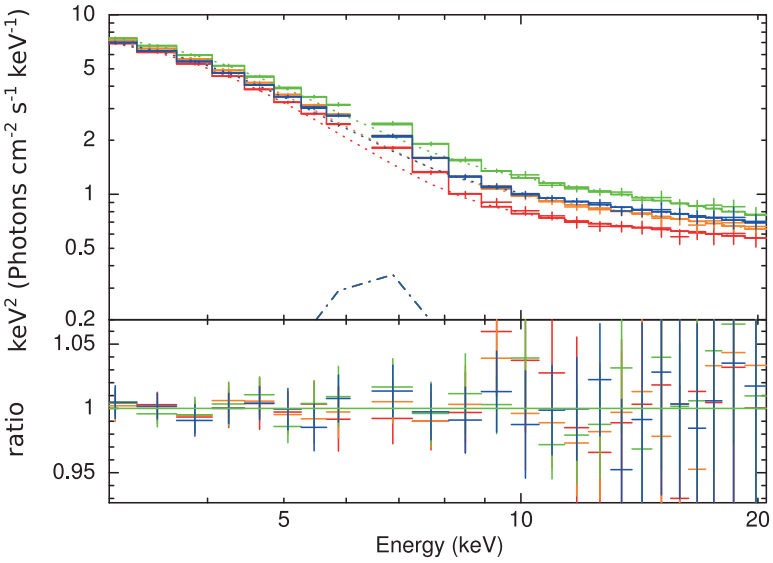

Figure 2.7: Phase-resolved energy spectra with model fits at four relative QPO phases, unfolded with the instrument response matrix and plotted in $EF(E)$ units. The colours have the same meaning as in Figure 2.6. The ratio of the data to the model is shown below the spectra. Since the iron line values are frozen or tied, the iron line components of each spectrum stack perfectly. The PCA has zero counts in detector channel 11 when in event-mode, hence the gap in the data at ∼6.5 keV.

differences between the untied parameter variations are discussed in Section 2.4.4. The equivalent width of the iron line varies with the opposite phase of the power-law continuum during a QPO cycle in the range $0.35 - 0.48$ keV. Note that since the iron line flux is constant, these equivalent width variations are due to the observed power-law continuum variations (which as we later show are likely due to geometric, not intrinsic, changes). Therefore, in this interpretation, the varying equivalent width does not represent a changing intrinsic disk reflection, but we report it here for the sake of completeness.

As an extension of model 1, we considered a variant on the DISKBB model with a free spectral hardening factor (Shimura & Takahara 1995) that changes the apparent peak temperature of the disk without changing the integrated disk flux, to represent variable scattering in the disk atmosphere. Our fits showed that the data also required a variation of DISKBB normalisation and/or temperature, i.e. the disk flux is required to change and hardening factor variations alone cannot explain the modulation of the blackbody emission. Effectively, hardening subsumes only part of the effect of an intrinsic temperature change (the change in spectral shape), without producing variations of disk flux which are also required by the data. As these results did not contribute any additional interpretation to the data beyond that of model 1, they are not reported further here.

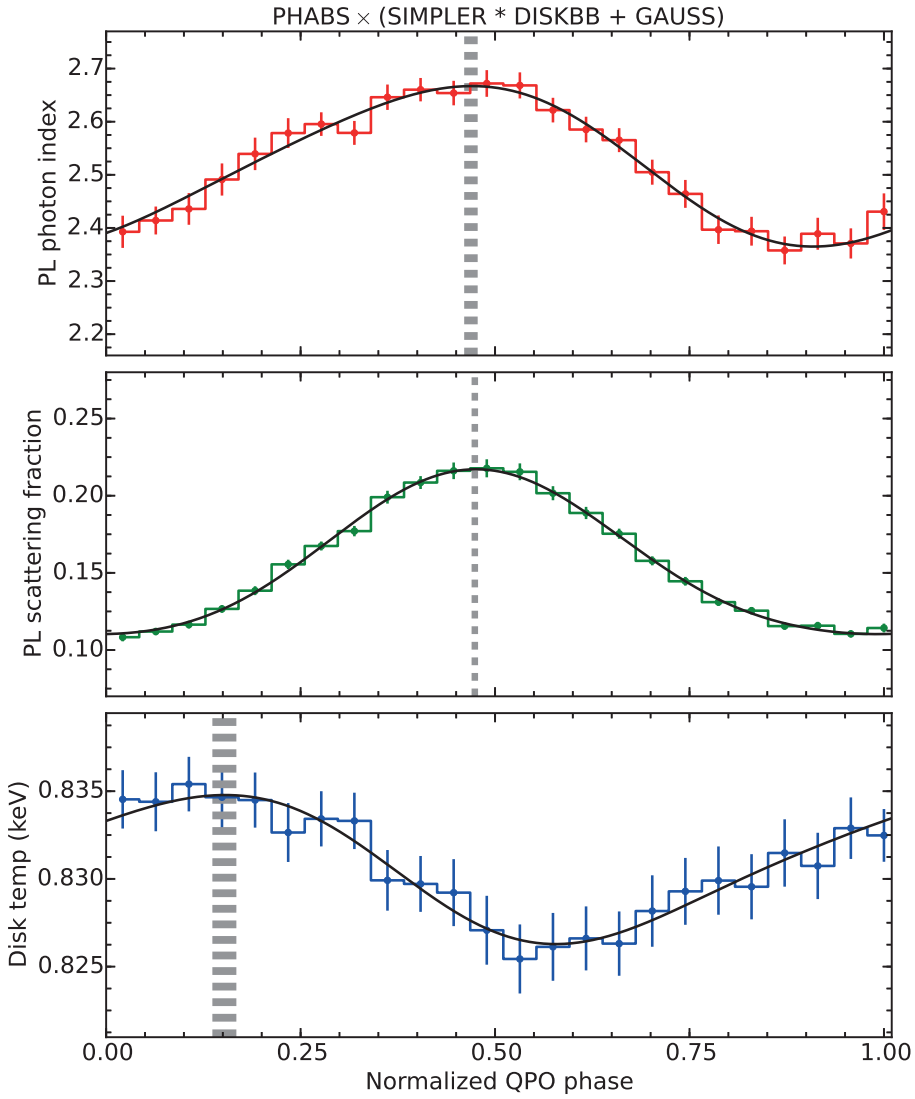

Figure 2.8: Best-fit values for untied spectral parameters at each phase in a QPO cycle for spectral model 1. The (1-σ) error bars were determined using bootstrapping (see Section 2.4.4). The top panel shows Γ values in red, the middle panel shows $f_{\rm scatt}$ values in green, and the bottom panel shows $T_{\rm disc}$ values in blue. The error bars represent 1σ errors (see Section 2.4.4 for details). The fractional rms amplitude variation of Γ is 4.2 per cent, of $f_{\rm scatt}$ is 24 per cent, and of the blackbody flux is 1.4 per cent (from luminosity $\propto T^4$ with constant blackbody normalisation). The best-fitting function (Equation 2.7) of each parameter is shown in solid black (the parameters of the best-fit function for the untied parameters in each spectral model are listed in Table 2.5). The phase of the maximum of each fit function is marked with a dashed gray line, the thickness of which represents the 1-σ error. The normalised phase difference between $f_{\rm scatt}$ and Γ is 0.01 ± 0.02, and between $f_{\rm scatt}$ and $T_{\rm disc}$ is 0.32 ± 0.02; this is discussed further in Section 2.4.4.

Spectral model 2

In the final best-fitting parameterisation of model 1, $T_{\rm disc}$ was required to vary, but the values had a very small variation ($0.8254 - 0.8354\,$keV). For the model component DISKBB, which has a fixed radial temperature profile, this variation implies that the whole disk blackbody is changing in effective temperature. It seems more physically plausible that a patch of the disk is varying, e.g. due to heating by the varying power law, while the intrinsic disk emission remains constant (particularly since the temperature variation is so small). To model the temperature variations required by this model parameterisation, we next tested a spectral model with an added single-temperature blackbody component BBODYRAD that is allowed to vary between phase bins, while both the DISKBB normalisation and temperature are tied across phase. Following that, in the next subsection, we tested a varying radial temperature dependence with the disk blackbody model DISKPBB (spectral model 3).

Our spectral model 2, incorporating an additional single-temperature blackbody, is defined to be PHABS $\times$ (SIMPLER * DISKBB + BBODYRAD + GAUSS). To determine the most effective parameterisation we followed the same procedure as in the previous section. For all fits, $T_{\rm disc}$ was allowed $0.6 - 1.0\,$keV, $T_{\rm bb}$ was allowed $0.1 - 1.0\,$keV, and all other parameters were allowed the same range of values as in model 1 (unless frozen).

Starting with all parameters tied across QPO phase (and $N_{\rm H}$ frozen), we found $\chi^2 = 7079.5$ for 567 d.o.f. As with model 1, the first and second spectral parameters that needed to be untied were $f_{\rm scatt}$ ($\chi^2 = 392.2$ for 544 d.o.f.) and Γ ($\chi^2 = 157.5$ for 521 d.o.f.). By then untying $T_{\rm bb}$ we see further improvement of the fit, to $\chi^2 = 125.7$ for 498 d.o.f. Again, an untied $N_{\rm line}$ gave a slightly better statistical fit (χ^2 / d.o.f. = 123.3/498) than an untied $T_{\rm bb}$, but it also returned an implausibly low value for the line energy. So, we selected the untied $T_{\rm bb}$ as the best fit for this set and froze $\sigma_{\rm line} = 0.82$ (the best-fitting value found using `steppar` on the mean event-mode spectrum) in the final parameterisation. The fits slightly preferred untying $T_{\rm bb}$ to untying $N_{\rm bb}$. In order to break the degeneracy between $T_{\rm bb}$ and $N_{\rm bb}$, we froze $N_{\rm bb}$ at 8857.68 (the best-fitting value). The best parameterisation for model 2 and the values of each parameter are stated in Table 2.3. The value of the untied parameters at each phase bin is shown in Figure 2.9. The iron line equivalent width varied in the range $0.23 - 0.32\,$keV with the opposite phase of the power-law variations, due to the varying power-law continuum.

Spectral model 3

We also considered a different radial temperature dependence or temperature profile of the disk. Typically for a standard disk, and assumed by DISKBB, the radial temperature dependence of the observed multicolour disk blackbody is $T(r) \propto r^{-p}$ where $p = 0.75$. In the DISKBB model component, this value of p is fixed. How-

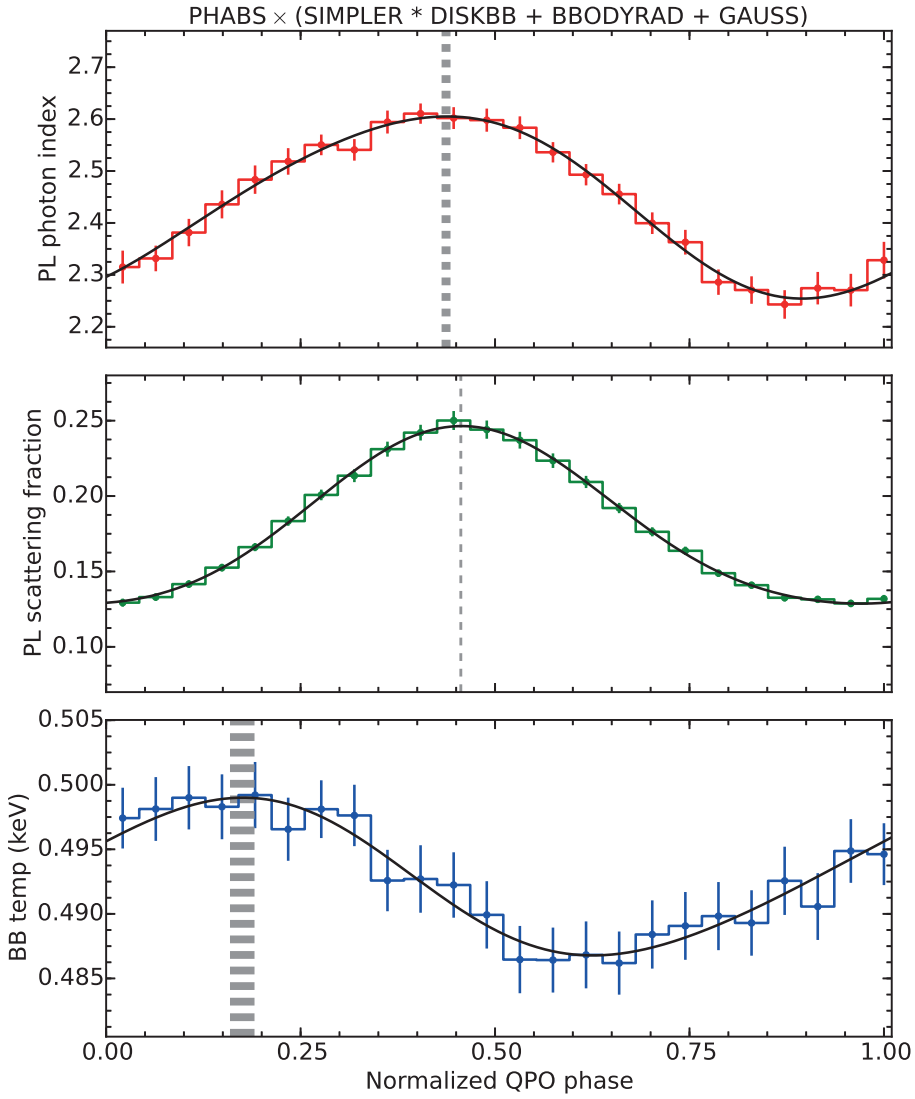

Figure 2.9: Same meaning as Figure 2.8, for model 2. The top panel shows Γ parameter values in red, the middle panel shows $f_{\rm scatt}$ parameter values in green, and the bottom panel shows $T_{\rm bb}$ parameter values in blue. The fractional rms amplitude variation of Γ is 4.9 per cent, of $f_{\rm scatt}$ is 23.5 per cent, and of the blackbody flux is 3.5 per cent (or 1.4 per cent in total blackbody flux, allowing for the constant disk blackbody component). The normalised phase difference between $f_{\rm scatt}$ and Γ is 0.019 ± 0.009, and between $f_{\rm scatt}$ and $T_{\rm bb}$ is 0.28 ± 0.02.

Component	Parameter	Value	Notes
PHABS	$N_{\rm H}$ ($\times 10^{21}$ cm^{-2})	6.0	Frozen
SIMPLER	Γ	2.44 †	Untied
SIMPLER	$f_{\rm scatt}$	0.179 †	Untied
SIMPLER	U	1	Frozen
DISKBB	$T_{\rm disc}$(keV)	0.94 ± 0.01	Tied
DISKBB	$N_{\rm disc}$	1090 ± 90	Tied
BBODYRAD	$T_{\rm bb}$ (keV)	0.49279 †	Untied
BBODYRAD	$N_{\rm bb}$	8857.68	Frozen
GAUSS	$E_{\rm line}$(keV)	6.40 ± 0.05	Tied
GAUSS	$\sigma_{\rm line}$ (keV)	0.82	Frozen
GAUSS	$N_{\rm line}$	$1.42 \pm 0.08 \times 10^{-2}$	Tied

Table 2.3: Values of the best-fitting parameterisation for model 2, with symbols and errors the same as in Table 2.2. XSPEC returned a fit statistic of $\chi^2 = 128.4$ for 500 d.o.f.

ever, as discussed in Kubota & Makishima (2004) and Kubota et al. (2005), general relativistic effects, electron scattering, and/or advection can give a different value of p. Heating by the power-law component may also cause the disk blackbody to deviate from the standard temperature profile. So, for a varying temperature dependence caused by, e.g. a varying illumination pattern, we tested the model PHABS × (SIMPLER * DISKPBB + GAUSS). In the component DISKPBB, the exponent of the radial dependence of the disk temperature p is a free parameter. Since DISKBB is a multicolour blackbody, and p changes how strongly the spectrum weights the component blackbodies, a smaller p increases the effect of the inner disk radii blackbodies on the total disk blackbody spectrum. Using DISKPBB also affects values of $f_{\rm scatt}$: since the spectrum of DISKPBB is flatter than DISKBB for $p < 0.75$, i.e. there are more lower energy photons, the Comptonising component modelled by SIMPLER must scatter a smaller fraction of the seed spectrum to produce the same shape as in models 1 and 2. For all fits, parameters were given the same value ranges as in model 1, and p was allowed the default range.

The starting χ^2/ d.o.f. was 7113.9/568 with all parameters tied. As with the first two models, untying $f_{\rm scatt}$ and then Γ gave the most significant improvements to the fit ($\chi^2 = 387.3$ for 545 d.o.f., and $\chi^2 = 167.0$ for 522 d.o.f., respectively). The next parameter that, when untied, gave the most significant improvement to the fit is p, returning $\chi^2 = 127.9$ for 499 d.o.f. Supervising the fitting, it became clear that reducing p effectively lowered $T_{\rm diskP}$ and $f_{\rm scatt}$. We note that untying $N_{\rm line}$ instead of p did not reduce the χ^2 contrary to what was seen for models 1 and 2. Untying additional parameters did not significantly improve the fit. For the final parameterisation, $\sigma_{\rm line}$ was frozen at 0.76 and $N_{\rm diskP}$ was frozen at 560.907 (the best-

Component	Parameter	Value	Notes
PHABS	$N_{\rm H}$ ($\times 10^{21}$ cm^{-2})	6.0	Frozen
SIMPLER	Γ	2.44 †	Untied
SIMPLER	$f_{\rm scatt}$	0.130 †	Untied
SIMPLER	U	1	Frozen
DISKPBB	$T_{\rm diskP}$(keV)	0.947 ± 0.003	Tied
DISKPBB	p	0.507 †	Untied
DISKPBB	$N_{\rm diskP}$	560.907	Frozen
GAUSS	$E_{\rm line}$(keV)	6.49 ± 0.05	Tied
GAUSS	$\sigma_{\rm line}$ (keV)	0.76	Frozen
GAUSS	$N_{\rm line}$	$1.23^{+0.07}_{-0.06} \times 10^{-2}$	Tied

Table 2.4: Values of the best-fitting parameterisation for model 3, with symbols and errors the same as in Table 2.2. XSPEC returned a fit statistic of $\chi^2 = 126.8$ for 501 d.o.f.

fitting values) following the same motivation as with models 1 and 2. A summary of the best parameterisation for model 3, with the value of each parameter, is shown in Table 2.4.

For model 3, the untied parameters are Γ (the power law photon index), $f_{\rm scatt}$ (the power law scattering fraction), and p (the exponent of the disk's radial temperature dependence). The value of the parameters in each phase bin is shown in Figure 2.10. Comparing the values of the untied parameters in this model with the previous two models, Γ has the same value as in model 2 and is lower than in model 1, $f_{\rm scatt}$ is lower here than in models 1 and 2, $T_{\rm diskP}$ is the same as $T_{\rm disc}$ in model 2 and is higher than the mean $T_{\rm disc}$ in model 1, and $N_{\rm diskP}$ is much lower than $N_{\rm disc}$ in models 1 and 2. These differences in the values are consistent with what we expect from the model component DISKPBB with $p < 0.75$. The iron line equivalent width varied from $0.21 - 0.29$ keV.

2.4.4 Parameter variation fitting

In order to study how the spectral parameters inferred from the CCF vary and to what significance, we must account for any error that is introduced due to the spectral data not being independent in time. Since the spectral fitting software assumes that the simultaneously fit data and errors on the data are fully independent, we cannot trust the errors from our spectral analysis. Therefore, to estimate the accuracy of the untied parameter variations, we implemented bootstrapping from our data set. Bootstrapping requires selecting (with replacement) M times from M segments of data. For our data set, 198 segments requires 5537 bootstrap iterations to converge on the

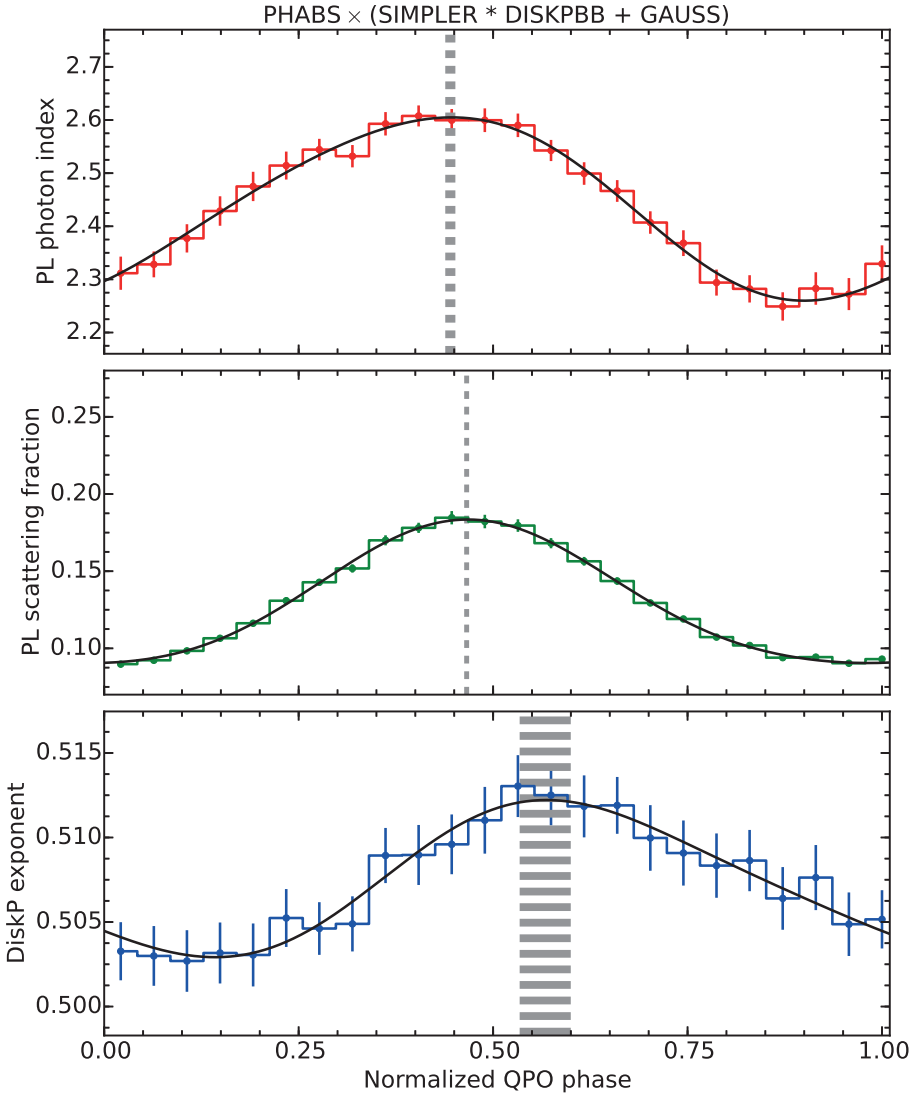

Figure 2.10: Same meaning as Figure 2.8, for model 3. The top panel shows Γ in red, the middle panel shows $f_{\rm scatt}$ in green, and the bottom panel shows p in blue. The fractional rms amplitude variation of Γ is 5.0 per cent and $f_{\rm scatt}$ is 26 per cent. The normalised phase difference between $f_{\rm scatt}$ and Γ is 0.02 ± 0.01, and between $f_{\rm scatt}$ and p is -0.10 ± 0.04.

true error distribution.[2] Each bootstrap iteration consisted of selecting CCF segments to make a new average energy-dependent CCF, through to the phase-resolved spectroscopy. Each bootstrapped average energy-dependent CCF was used to determine a set of phase-binned spectral parameters for each spectral variability model, in the same way as for the data. The standard deviations for each parameter were determined from the distributions obtained from all bootstrap iterations and are used to determine the errors plotted in Figures 2.8, 2.9 and 2.10.

Even though the untied spectral parameters are free to have any value within the physically allowed ranges given previously, and despite the apparently complex form of spectral variability seen in Figure 2.6, the parameters appear to vary in a roughly sinusoidal fashion. The function fitted to the variations is

$$y(t) = A_1 \cdot \sin(2\pi t + \phi_1) + A_2 \cdot \sin(4\pi t + \phi_2) + D\,, \tag{2.7}$$

where A_1 and A_2 are amplitudes, ϕ_1 and ϕ_2 are phase offsets, and D is a y-axis offset. The function is fixed so that the second part is a harmonic, with twice the frequency of the first part. We refer to this function as a whole as the "fit function". The fit parameters A_1, A_2, ϕ_1, ϕ_2, and D were found for each bootstrap iteration by a least-squares fit (using scipy.optimize.leastsq) to the untied parameter value variations.

If there is a strong harmonic present in the QPO data, the CCF erases any phase offset between the fundamental and the harmonic. This is because the CCF picks out only the *relative* offset of the channel-of-interest fundamental and harmonic with the reference band fundamental and harmonic respectively. So, any harmonic content is added into the CCF such that its phase is relative to the reference band phase of the harmonic and thus the phase dependence of spectral parameters need not be aligned with those of the fundamental. The GX 339–4 data used in this analysis has a weak harmonic in the 3–5 keV power spectrum, but more harmonic content at higher energies (Figure 2.1). We do not exclude the possibility that the stronger harmonic at higher energies could give a harmonic in the untied parameter values, and so we included a harmonic with all parameters free in the fit function. Overall we found that A_2 is at least a factor 4 times smaller than A_1 for all variable parameters and is only formally significant for $f_{\rm scatt}$.

Phase relationships of the spectral parameters

The parameters of the fit function for each untied parameter per model are shown in Table 2.5. The phase differences between the untied parameters were measured from the maximum of the fit function $\phi_{\rm max}$ for each parameter, with the errors corresponding to the standard deviation obtained from bootstrapping. Since $f_{\rm scatt}$ was the best-constrained untied parameter, we computed the phase difference of the two

[2] Using $M \cdot (\ln(M))^2$ (Feigelson & Babu 2012).

Model	Parameter	A_1	ϕ_1	A_2	ϕ_2	D	$\phi_{\max}$
1	Γ	0.15 ± 0.01	0.834 ± 0.009	0.02 ± 0.02	0.1 ± 0.3	2.518 ± 0.008	0.449 ± 0.010
	$f_{\rm scatt}$	0.053 ± 0.003	0.792 ± 0.004	0.007 ± 0.001	0.48 ± 0.02	0.157 ± 0.001	0.454 ± 0.005
	$T_{\rm disc}$	0.0041 ± 0.0004	0.16 ± 0.02	0.001 ± 0.001	0.8 ± 0.7	0.831 ± 0.002	0.13 ± 0.02
2	Γ	0.17 ± 0.02	0.859 ± 0.005	0.01 ± 0.01	0.1 ± 0.2	2.44 ± 0.01	0.417 ± 0.007
	$f_{\rm scatt}$	0.059 ± 0.003	0.811 ± 0.002	0.007 ± 0.001	0.892 ± 0.008	0.180 ± 0.003	0.436 ± 0.002
	$T_{\rm bb}$	0.0060 ± 0.0006	0.12 ± 0.01	0.001 ± 0.001	0.2 ± 0.2	0.493 ± 0.002	0.16 ± 0.02
3	Γ	0.17 ± 0.01	0.851 ± 0.006	0.02 ± 0.01	0.1 ± 0.2	2.44 ± 0.01	0.425 ± 0.008
	$f_{\rm scatt}$	0.047 ± 0.002	0.800 ± 0.003	0.006 ± 0.001	0.37 ± 0.01	0.131 ± 0.001	0.446 ± 0.004
	p	0.0045 ± 0.0005	0.67 ± 0.02	0.001 ± 0.001	0.3 ± 0.5	0.507 ± 0.002	0.55 ± 0.04

Table 2.5: Values of each variable of the fit function (Equation 2.7) for the untied parameters in each spectral model. $\phi_{\max}$ is the phase of the maximum value of the fit function (indicated with the dashed gray lines in Figures 2.8-2.10). The errors are computed as the standard error from bootstrapping. The values of ϕ have been normalised to the range 0 to 1.

other untied parameters with respect to it. The difference in phase between the variations in the untied parameter values are also quoted in Table 2.1. The phases of $f_{\rm scatt}$ and Γ are very close: 0.01 ± 0.02, 0.019 ± 0.009, and 0.02 ± 0.01 for models 1, 2, and 3, respectively (normalised to the range 0 to 1; 1-σ errors). The phase difference for model 1 between $f_{\rm scatt}$ and $T_{\rm disc}$ is $\Delta\phi = 0.32 \pm 0.02$. For model 2, $f_{\rm scatt}$ and $T_{\rm bb}$ and have $\Delta\phi = 0.28 \pm 0.02$. For model 3, $f_{\rm scatt}$ and p are out of phase by $\Delta\phi = -0.10 \pm 0.04$. As previously noted, lowering p gives the same effect as raising $T_{\rm diskP}$ and $f_{\rm scatt}$, so these results are consistent with the previous two models finding that the blackbody increases before the power law increases. If we instead compare the phase of the maximum of $f_{\rm scatt}$ with the phase of the *minimum* of p, we see a phase difference of 0.33 ± 0.02, consistent with the phase lead seen for the blackbody temperature variations in models 1 and 2. In all three models, the blackbody parameter values have a small variation, but as evidenced by the systematic spectral analysis, they are required to vary. Before considering the physical interpretation of the results in Section 2.5, we first checked that the QPO spectral parameter variations inferred from our CCF method can also reproduce the shape of the lag-energy spectrum.

2.4.5 Comparison with the lag-energy spectrum

The last step of the analysis was to use the best-fitting energy spectral models to simulate lag-energy spectra and compare them to the data. This was done as a secondary check on the models in the Fourier domain, where the QPO can be selected directly in frequency and the errors are better-behaved.

We selected untied parameter values from the smooth fit function at each of the 24 QPO phases, and simulated a PCU2 event-mode spectrum for each phase bin using the XSPEC command `fakeit`, with the same instrumental response matrix as the data. The spectra are simulated without Poisson errors, to provide an idealised model. We then tiled the spectra to make a light curve with the same length as a segment used to calculate the CCF, and turned the spectra into photon count-rate light curves in each energy channel. From here, the simulated models were treated in the same way as the data for lag-energy spectral analysis.

In Figure 2.11a we show simulated lag models plotted with the data from Figure 2.2 (model 1, in dark blue, gives $\chi^2 = 37.9$; model 2, in purple, gives $\chi^2 = 45.5$; model 3, in orange, gives $\chi^2 = 47.2$; all have 24 d.o.f.). While all three of the QPO spectral models can match the overall 'broken' shape of the lag-energy spectrum quite well, models 2 and 3 slightly deviate from the slope of the observed lag-energy spectrum at low ($< 4\,$keV) and high ($> 16\,$keV) energies.

For comparison, in Figure 2.11b, we also plot the simulated lag-energy spectra from some alternative parameterisations of model 1. The first, in dark red, has only $f_{\rm scatt}$ and Γ untied. It is clear from the plot that this model cannot re-create the slope below 5 keV or the break in the lag at $\sim$6.5 keV. Likewise, the second and third

parameterisations vary $E_{\rm line}$ and $N_{\rm line}$, respectively, in addition to $f_{\rm scatt}$ and Γ and with the iron line width frozen at $\sigma_{\rm line} = 0.97$. Neither can reproduce the slope below 5 keV, and the second cannot make a bump with the right shape at the right energy. All three of these lag models have the same high-energy slope as model 2 from panel a. These poor matches are reflected in the fits to the data: $\chi^2 = 175.6$, $\chi^2 = 149.1$, and $\chi^2 = 82.1$, respectively, each for 24 degrees of freedom.

The bootstrapping technique accounts for statistical error, but it is possible that some systematic error is unaccounted for. However, since simulations from the bootstrapping-verified parameter variations can reproduce the lag-energy spectrum of the data well, our phase-resolved spectroscopy technique does not seem to introduce a systematic bias to the spectral-timing properties of the data.

The lag-energy spectra simulated from the QPO phase-resolved spectral models further support a spectral model for QPO variability consisting of a varying blackbody preceding (by relative phase ~ 0.3) a varying power law, however the lags show a better fit with model 1 (as opposed to the phase-resolved spectral fits, which preferred models 2 and 3). As demonstrated here, it is very important to account for the Fourier-domain cross spectra as well as the time-domain energy spectra when fitting a model to the data, as both spectral and timing information are pertinent. Furthermore, our results show that it is possible to reproduce the lag-energy spectrum of a Type B QPO by sinusoidally varying spectral parameters with different relative phases.

2.5 Physical Interpretation

We fitted three spectral models to phase-resolved spectra of the Type B QPO from the 2010 outburst of GX 339–4 which we determined using our CCF method. All of our models consisted of a blackbody-like disk component and a power law component produced by Compton up-scattering of the disk blackbody photons, but parameterised the blackbody contribution in different ways: either variations of an entire multiple-temperature disk blackbody, variations of a single-temperature blackbody on top of a constant disk blackbody, or variations in the radial temperature profile of a multi-temperature disk blackbody. All three models showed that the blackbody must vary as well as the power law on the QPO timescale: the varying blackbody temperature leads the varying power law scattering fraction $f_{\rm scatt}$ by a relative phase of ~ 0.3, and the varying power law photon index Γ is very close in phase to the variations in $f_{\rm scatt}$ (with a model-dependent relative phase lead of up to ~ 0.02). These results are checked with and supported by bootstrapping analysis and simulated lag-energy spectral models compared with the data.

In terms of variability amplitudes, the power-law component dominates the QPO variability, showing fractional rms variations in scattering fraction of ~ 25 per cent, when calculated directly from the sinusoidal fit function applied to each model. In

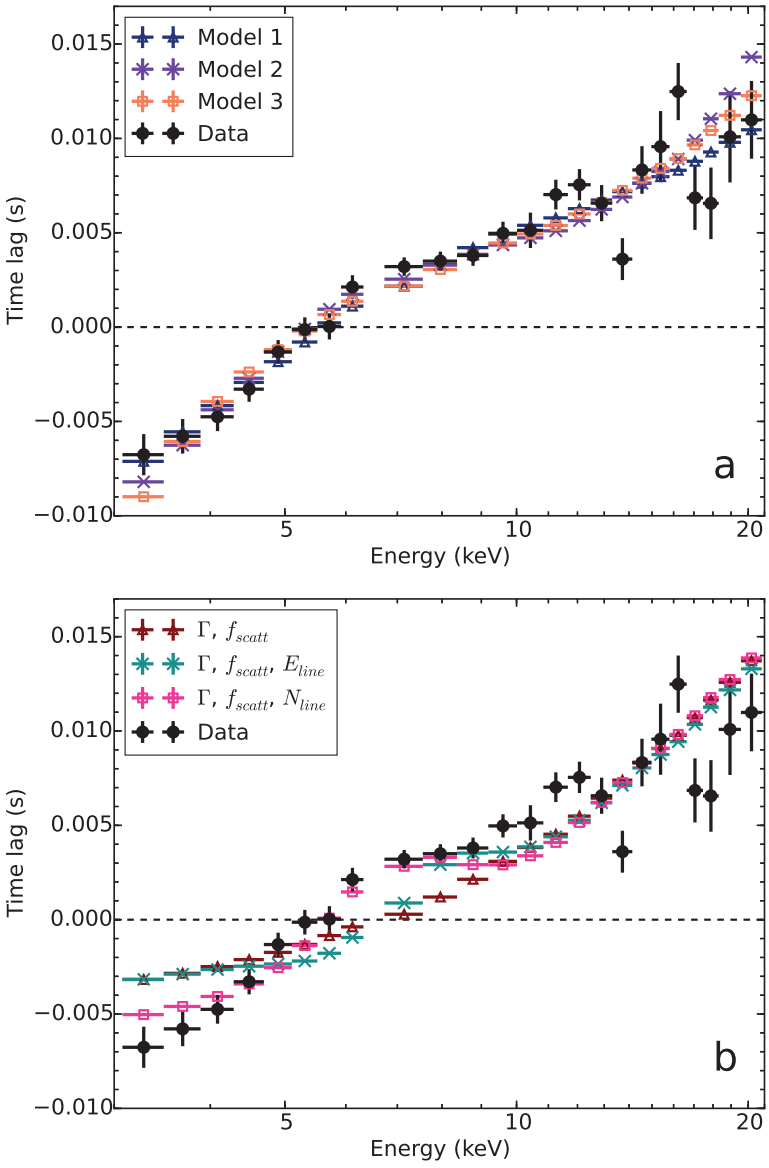

Figure 2.11: Lag-energy spectra obtained from the cross-spectrum averaged over 4–7 Hz. The black dots are the data (same as in Figure 2.2). The black dashed line indicates no time lag with respect to the variability in the reference band. **a**: The dark blue triangles are simulated using model 1 ($\chi^2 = 37.9$). The purple X's are simulated using model 2 ($\chi^2 = 45.5$). The orange squares are simulated using model 3 ($\chi^2 = 47.2$). **b**: These simulations are alternative parameterisations of model 1 that fit the data poorly. The dark red triangles are simulated with Γ and $f_{\rm scatt}$ untied ($\chi^2 = 175.6$). The dark cyan X's are simulated with Γ, $f_{\rm scatt}$, and $E_{\rm line}$ untied and $\sigma_{\rm line} = 0.97$ ($\chi^2 = 149.1$). The pink squares are simulated with Γ, $f_{\rm scatt}$, and $N_{\rm line}$ untied and $\sigma_{\rm line} = 0.97$ ($\chi^2 = 82.1$). All of the model-comparisons in panels **a** and **b** have 24 degrees of freedom.

comparison, the blackbody temperature in model 1 shows a fractional rms of only 0.35 per cent, corresponding to a flux variation of only 1.4 per cent. The total blackbody flux variation is similar for the other models (e.g. after including dilution of the larger single blackbody rms variability by the constant disk blackbody contribution in model 2). These results are consistent with the findings of Gao et al. (2014) that the power law dominates the Type B QPO emission in GX 339–4 (and, more generally, is consistent with the findings of Remillard & McClintock 2006; Sobolewska & Życki 2006 and Axelsson et al. 2014 for LFQPOs). However, despite the dominant role of the power-law emission in the observed QPO variability, the large phase-lag of the power-law relative to the blackbody has a significant effect on the lag-energy spectrum, causing the distinctive 'break' feature, as shown in Figure 2.11.

Although our spectral models allowed for different types of blackbody variability, all point to a very similar phenomenological picture of the QPO spectral variability. The parameter variations show large changes in the power law index and normalisation (measured by Γ and $f_{\rm scatt}$) and small changes in the blackbody emission. For model 1, the $T_{\rm disc}$ variations represent the entire disk blackbody modestly varying in effective temperature, suggesting that it is not only the hottest region of the disk which varies. This picture is supported when we keep $T_{\rm disc}$ constant with phase in model 2 and allow a separate single blackbody component to vary, resulting in modest variations of the single blackbody $T_{\rm bb}$ but with a *lower* temperature than the inner disk temperature. Allowing the disk radial temperature profile to vary in model 3 also supports this view, with the variations in profile mimicking the variations in the disk spectrum down to lower energies that are seen in the other two models.

We now consider the broad physical implications of our results. The simplest explanation for linked blackbody and power-law variability is that the blackbody drives the power-law variation by varying the seed-photon flux which illuminates the Compton scattering region. This interpretation suffers from the problem that the observed phase lag corresponds to a light-travel time of ~ 60 ms or a distance of $\sim 1.8\times10^4$ km: $\sim 1700\,r_g$ for a $7\,M_\odot$ black hole. This distance is much larger than the disk inner radius inferred from the disk blackbody normalisation of $\sim$100–260 km,[3] so any model seeking to explain the lag in terms of light-travel times must assume an unfeasibly large height of the Compton-scattering region above the disk. Furthermore, if the observed photon index fractional rms variation of ~ 5 per cent is produced by changes in the seed to heating luminosity ratio in the Compton scattering region, the required variations in seed luminosity would need to be much larger than the observed disk variability (e.g. as much as ~ 30 per cent, applying the equation from Beloborodov 1999).

Another possibility is that quasi-periodic variations in accretion rate are produced

[3]Here we assume a disk spectral hardening factor $f \sim 1.7$ (see Shimura & Takahara 1995), inclination 40 degrees (Muñoz-Darias et al. 2013) and calculate values for the conservative distance range of 6–15 kpc (Hynes et al. 2004).

in the accretion disk and propagate inwards to produce a lagging power-law component from an inner corona, as seems to be the case for the hard state broadband noise variability in GX 339-4 and other black hole systems (Uttley et al. 2011). This interpretation also seems unlikely, since for the Type B QPO the disk blackbody variations are much smaller than the power-law variations (unlike the situation for the broadband noise). It might be possible that the accretion rate variation originates only in the innermost radii of the disk and so modulates only a small fraction of the blackbody emission, but then the lag would likely be washed out by variable heating effects of the strong power-law variations (see below).

An alternative possibility is that the fluctuations in blackbody emission are driven by X-ray heating from intrinsically varying power-law emission. According to model 1, the average total power-law flux is between 40 and 60 per cent of the total disk blackbody flux (assuming a power-law high-energy cut-off $> 20\,\mathrm{keV}$). If the power-law emitting region is isotropically emitting and the disk albedo is 0.3 (e.g. Gierliński et al. 2008), then given the observed 25 per cent power-law rms variation, the corresponding disk rms variability should be $(0.07\text{–}0.1)\times f_{\mathrm{cov}}$, where f_{cov} is the fraction of the sky covered disk as seen from the power-law. The observed disk variability could then easily be produced by X-ray heating even for modest covering fraction of ~ 0.2. However, we then would expect the blackbody variations to be in phase with those of the power-law, or at least lag by a (small) light-travel time, which is not the case. Therefore we are left with the simplest explanation for our results, which is that the variation in power-law emission is not intrinsic but *geometric*, in which case the observer and the disk need not see the same power-law emission, and large phase lags can be produced by the varying system geometry. This interpretation is consistent with the idea that the inner power-law emitting region is precessing, perhaps via the Lense-Thirring mechanism (e.g. Stella & Vietri 1999). We will therefore consider this scenario in more detail. However, we do not rule out the possibility that other models not considered here can explain the observed energy spectral changes, and we stress again that the phase-resolved spectroscopy technique presented in this paper is applicable regardless of the physical interpretation of the results.

First, we should consider that as seen from the disk, a precessing power-law emitting region should (quasi-)periodically illuminate different azimuthal regions of the disk, producing a local enhancement of the blackbody emission due to X-ray heating. Even if the power-law emission is intrinsically constant, an observer will see periodic blackbody variations as a result, due to the Doppler boosting and deboosting of the region of enhanced blackbody emission as the power-law emitting region illuminates the approaching and receding sides of the disk as seen by the observer.

Type C QPOs, which occur on timescales very similar to Type Bs but are observed across a wider range of outburst states, might arise from Lense-Thirring precession of a hot inner flow inside the inner edge of the accretion disk (Fragile & Anninos 2005; Ingram et al. 2009; Axelsson et al. 2014; Ingram & van der Klis 2015; Fragile

et al. 2016). A hot inner flow would have a small scale height, and the precession would effect more significant changes in illumination and heating of the inner edge of the accretion disk than in the cooler parts of the disk. Since the inner edge of the accretion disk shows the strongest Doppler boosting, the varying illumination pattern should show a large blackbody modulation. For example, at a disk radius of $10\,r_g$ and assuming a disk inclination of 40 degrees, the maximum Doppler factor for gas on the approaching side of the disk is 1.13, so that allowing for boosting and deboosting, the peak to trough apparent temperature shift should be 27 per cent compared with 2.4 per cent observed for model 2 (which shows the most extreme temperature change). Even a large emitting radius in the disk of $100\,r_g$ will produce a peak-to-trough shift of 10 per cent. However, this variation assumes that only a small range of azimuths on the disk is illuminated, to maximise the observed temperature shift.

The problem can therefore be solved by assuming that the illuminated region of the disk is large and only undergoes small shifts due to precession, such that the observed Doppler boost is an average over many azimuths and changes only by a small amount in response to the changing illumination pattern. Such an effect might be obtained if the precessing power-law emitting region which produces the Type B QPO has a relatively large scale height compared to the inner disk radius (up to tens of r_g), i.e. it is more *jet-like* than disk-like (as for the Type C QPOs). This picture is consistent with the findings of Motta et al. (2015), that the Type B QPO integrated rms is larger for more face-on binary systems and the Type C QPO integrated rms is larger for more edge-on systems, implying that the emission geometries are different.

Detailed modelling of the effects of a variable illumination pattern on the disk blackbody emission is beyond the scope of this work, but we illustrate our interpretation by showing in Figure 2.12 the illumination footprint on the disk at four key phases (separated by 0.25 in relative phase). At phase *a*, the apparent blackbody temperature is at a maximum when the power law is preferentially illuminating the Doppler-boosted approaching side of the accretion disk, while at phase *c* the opposite, deboosted side of the disk is preferentially illuminated, leading to a minimum in apparent temperature. Linking the power-law emission to the illumination pattern is less-clear since it depends on the orientation-dependence of the power-law emission. Here we assume that the illumination pattern precesses in the same direction as the orbital motion, so that for a jet-like power-law emitting region, we might expect the jet to point more towards the observer at phase *b* and more away from the observer at phase *d*. Comparing with our observed pattern of disk and power-law parameter variations (e.g. in Figure 2.8 for model 1), the peak in power-law emission and photon index would coincide roughly with the time when the jet-like emitting region points towards us at phase *b*, suggesting that Doppler boosting of the emitting region may play a role in the observed power-law flux variability, which might be expected if the emitting region is jet-like in geometry.

Note that in the above scenario we do not necessarily expect the peak in blackbody

emission and power-law emission to be separated by exactly 0.25 in relative phase, which is close to but significantly different from the phase difference we observe. The difference will also be influenced by the shape of the illuminating region on the disk and hence the relative weighting of emission from different parts of the disk, and we would also anticipate that light-bending effects should play a role. Such effects can be predicted in future using ray-tracing simulations. Our toy model also does not simply explain the tight relation (and very small phase offset) between the scattering fraction (or effectively, power-law normalisation) and photon index, although this result suggests that the spectral shape of the power-law emission is also orientation-dependent. Also we do not observe the variations in the iron line energy (or flux) on the QPO timescale that are predicted by Fragile et al. 2005 and Ingram & Done 2012 for LFQPOs produced by precessing hot flows in XRBs. Since the disk blackbody varies by only a very small fraction, if this effect were due to variable illumination combined with Doppler boosting, we would expect a similarly small fractional amplitude of variation in the iron line emission. Such a variation would not be detectable given the lack of sensitivity to a very small iron line flux variation and the coarse energy resolution at the expected iron line energy (5.5–7.0 keV).

An obvious interpretation of our toy model geometry is that the jet-like power-law emitting region is actually the base of the larger-scale radio-emitting jets seen in the hard and intermediate states of black hole X-ray binary systems. A jet-base origin for the X-ray power-law emission in X-ray binaries was first proposed by Markoff et al. (2005). A number of works support the possibility of a connection between Type B QPOs and jets. Type B QPOs are only seen at the crossover between the soft and hard spectral states, in the soft intermediate state (Belloni 2010; Heil et al. 2015b). During this transition in previous outbursts, there are observations of radio flaring and evidence of the radio-emitting jet switching off (Fender et al. 2004, 2009). Moreover, Kalamkar et al. (2016) have reported the discovery of an infra-red LFQPO in GX 339-4, coincident with an X-ray Type C QPO at the first harmonic of the IR QPO. Their results suggest that low-frequency quasi-periodic phenomena such as precession could extend to or affect the IR–producing part of the jet, which would make it plausible that precession also affects the X-ray emitting jet base in the case of the Type B QPO.

2.6 Conclusions

We have demonstrated a new spectral-timing technique for phase-resolved spectroscopy using the cross-correlation function, to enable deeper understanding of the QPO mechanism in black holes and neutron stars. The technique allows us to quantitatively probe how the spectrum of a source changes as a function of QPO phase. For the Type B QPO in GX 339–4's 2010 outburst observed with *RXTE*, the spectral shape

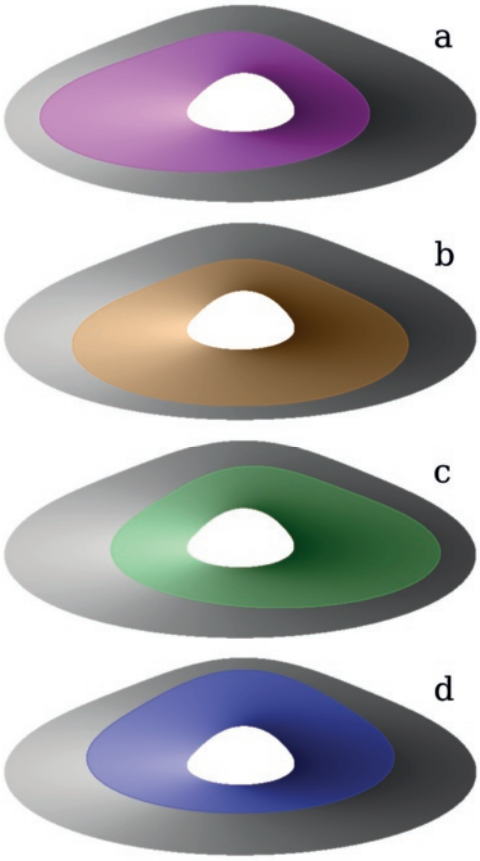

Figure 2.12: A simplified depiction of the illumination footprint on the inner portion of the accretion disk by a large-scale-height precessing Comptonising region as it sweeps around. The grayscale colour-mapping of the disk shows the intensity of the emission; light shading means higher observed intensity (which varies as a function of Doppler blueshift and intrinsic emission). The accretion flow is orbiting in a counter-clockwise direction. The illuminated region of the accretion disk is shown in magenta, orange, green, and blue for phases *a*, *b*, *c* and *d* respectively in a QPO cycle (which we assume are each separated by 0.25 in relative phase).

changes on the QPO timescale. Specifically, we find that the power-law normalisation has a ~ 25 per cent fractional rms amplitude variation, the power-law photon index has ~ 5 per cent rms variation, while the blackbody varies significantly, but with an rms variation in flux of only 1.4 per cent. Crucially, the varying blackbody-like component leads the varying power law component by ~ 0.3 of a QPO cycle. This combination of large power-law variability with small blackbody variability, combined with a large phase lag of the power-law relative to the blackbody variation, implies that the QPO

flux variations are not primarily intrinsic to the emitting regions but geometric. The spectral and flux variability and lag suggest a large-scale-height Comptonising region (such as the base of a jet) which illuminates and heats different but overlapping regions of the inner accretion disk as it precesses during a QPO cycle.

Regardless of the physical interpretation, our method may be used as a guide to understand quasi-periodic variability in terms of the different spectral components. The statistical errors across phase-bins are not statistically independent however, so caution should be applied when interpreting the detailed results, with bootstrapping or other Monte Carlo methods applied to determine uncertainties on spectral parameter variations. Furthermore, the spectral component variability inferred from our method should always be compared with statistically robust spectral-timing products, e.g. derived from the cross-spectrum. Looking further ahead, more sophisticated fitting techniques (e.g. using the error covariance between phase bins) could be developed to carry out more robust fitting of the phase-resolved spectra. Alternatively, direct frequency-domain fitting of the QPO cross-spectrum offers be a simple and robust approach, but may suffer from the problem that variable spectral modelling in the frequency domain is not intuitive, compared to time-domain fitting which may be more easily related to conventional spectral-fitting of time-resolved spectra.

Our results highlight the power of the combination of high count-rates and improved spectral and timing resolution in understanding the innermost regions of accreting compact objects. Unfortunately, telemetry constraints limited *RXTE* PCA observations of brighter sources than GX 339–4 to data modes with even coarser energy binning than the event mode data we use here. However, the recently launched *ASTROSAT* mission (Singh et al. 2014) suffers no such constraints on data from its Large-Area X-ray Proportional Counter (LAXPC) experiment, which offers full event telemetry for much brighter sources (and with larger effective area than the PCA). Similarly, we anticipate important breakthroughs from applying these techniques to data from the forthcoming *Neutron Star Interior Composition Explorer* (*NICER*, Arzoumanian et al. 2014), since *NICER* will obtain the first data with a soft X-ray response and CCD-level spectral-resolution for Crab-level and higher fluxes, without significant deadtime or instrumental pile-up effects. Thus *NICER* will allow us to probe the quasi-periodic variability of Fe K shape due to its superior energy resolution, and distinguish between different models for blackbody variability, thanks to its soft X-ray response.

Since the signal-to-noise of X-ray binary spectral-timing measurements scales linearly with count rate (Uttley et al. 2014), future large-area X-ray observatories with dedicated timing capability and large ($>$ few m^2) collecting areas at iron line energies, will enable our technique to perform detailed phase-resolved tomography of the Fe K line, even for weak QPO signals. In fact, our CCF method was originally conceived to demonstrate tomography of the much weaker HFQPO signatures in data from the proposed *Large Observatory for X-ray Timing* (*LOFT*, Feroci et al. 2012), which

could directly reveal the orbital motion in strong field gravity of emission structures producing the HFQPOs.[4] Thus the phase-resolved spectroscopy of QPOs offers a powerful probe of the behaviour of matter in the most strongly curved space-times and we encourage further development of spectral-timing techniques to reveal the inner workings of QPOs.

Acknowledgements

A.L.S. acknowledges support from NOVA (Nederlandse Onderzoekschool voor Astronomie). We thank Adam Ingram, Lucy Heil, Victoria Grinberg, Michiel van der Klis, and the participants of 'The X-ray Spectral-Timing Revolution' Lorentz Center workshop (February 2016) for useful discussions that contributed to the development of this paper. We also thank the anonymous referee for their helpful comments.

This research has made use of data and software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC); NASA's Astrophysics Data System Bibliographic Services; NumPy v1.9.3 and Scipy v0.16.0 (Jones et al. 2001); Astropy v1.0.4 (Astropy Collaboration et al. 2013); Matplotlib v1.4.3 (Hunter 2007); iPython v3.2.0 (Perez & Granger 2007); and the AstroBetter wiki.

Appendix 2A: Software

The software developed for the phase-resolved spectroscopy technique presented in this paper will be made publicly available 6 months after this paper's publication. The repository links are as follows:

- power spectra:
  `https://github.com/abigailStev/power_spectra`
- lag spectra:
  `https://github.com/abigailStev/lag_spectra`
- cross-correlation:
  `https://github.com/abigailStev/cross_correlation`
- phase-resolved spectroscopy:
  `https://github.com/abigailStev/energy_spectra`
- SIMPLER XSPEC model:
  `https://github.com/abigailStev/simpler`
- simulating spectral-timing data:
  `https://github.com/abigailStev/simulate`

[4] See Figure 2-17 in `http://sci.esa.int/loft/53447-loft-yellow-book/#`

3

Neutron Star Mass-Radius Constraints using Evolutionary Optimization

A. L. Stevens, J. D. Fiege, D. A. Leahy, and S. M. Morsink

The Astrophysical Journal, 2016, 833, 244

Abstract

The equation of state of cold supra-nuclear-density matter, such as in neutron stars, is an open question in astrophysics. A promising method for constraining the neutron star equation of state is modelling pulse profiles of thermonuclear X-ray burst oscillations from hotspots on accreting neutron stars. The pulse profiles, constructed using spherical and oblate neutron star models, are comparable to what would be observed by a next-generation X-ray timing instrument like *ASTROSAT*, *NICER*, or *LOFT*. In this paper, we showcase the use of an evolutionary optimization algorithm to fit pulse profiles to determine the best-fit masses and radii. By fitting synthetic data, we assess how well the optimization algorithm can recover the input parameters. Multiple Poisson realizations of the synthetic pulse profiles, constructed with 1.6 million counts and no background, were fitted with the Ferret algorithm to analyze both statistical and degeneracy-related uncertainty, and to explore how the goodness-of-fit depends on the input parameters. For the regions of parameter space sampled by our tests, the best-determined parameter is the projected velocity of the spot along the observer's line-of-sight, with an accuracy of $\leq 3\%$ compared to the true value and

with $\leq 5\%$ statistical uncertainty. The next best-determined are the mass and radius; for a neutron star with a spin frequency of 600 Hz, the best-fit mass and radius are accurate to $\leq 5\%$, with respective uncertainties of $\leq 7\%$ and $\leq 10\%$. The accuracy and precision depend on the observer inclination and spot co-latitude, with values of $\sim 1\%$ achievable in mass and radius if both the inclination and co-latitude are $\gtrsim 60°$.

3.1 Introduction

Neutron stars are an astrophysical laboratory for studying cold supra-nuclear-density matter. Accreting millisecond X-ray pulsars, a particular subset of neutron stars in low-mass X-ray binaries (LMXBs), are rapidly spinning accretion-powered neutron stars with spin periods of a few milliseconds (e.g., SAX J1808.4–3658, Wijnands & van der Klis 1998b). Their pulsed X-ray emission originates from material striking the surface of the neutron star during regular accretion and warming an area on the surface so that it emits blackbody radiation. Then, as the neutron star rotates, it gives periodic oscillations in brightness as the emitting region faces towards and away from the observer. Since these photons originate from the surface of the neutron star itself, physical properties like its mass and radius are encoded in the detected pulse profile. Fitting these pulse profiles with realistic models can then yield neutron star mass and radius estimates (Watts et al. 2016).

In addition to regular pulsed X-ray emission, some neutron stars in LMXBs exhibit thermonuclear (Type I) X-ray bursts (Watts 2012) . In a fraction of thermonuclear X-ray bursts, we observe brightness oscillations, where the frequency corresponds strongly with the spin frequency of the neutron star; these are referred to as thermonuclear burst oscillations. The pulse profile models that we discuss in this paper refer specifically to models of these burst oscillations.

The emission area on the surface of the neutron star is referred to as the hotspot or spot. Theories suggest two different surface hotspot models: one that ignites nuclear burning at one point and spreads across the whole neutron star, and another that ignites at one point and begins to spread but remains limited to a smaller area (referred to as a "persistent hotspot") (see Watts 2012 and references therein). The persistent hotspot on the surface of a rotating neutron star has been demonstrated to be an effective model for Type I X-ray bursts from the source 4U 1636–536 (Artigue et al. 2013), and so we use a persistent spot model with no size variation over the course of the burst. The fixed spot model is used for convenience here; a changing spot model could be incorporated for observations that show evidence for such behaviour.

The spectral model depends on the physics of the spot production and includes both the energy and angular dependence of the emitted radiation. In the case of rotation-powered X-ray pulsars (Bogdanov 2013), a hydrogen atmosphere model is appropriate (e.g., Heinke et al. 2006). The hydrogen atmosphere model depends on the spot's temperature and the local surface gravity. Since the surface gravity only depends on the mass, radius, and spin of the star (AlGendy & Morsink 2014), the local temperature is the only additional free parameter introduced by the spectral model.

For accretion-powered X-ray pulsars, an empirical model including a blackbody plus Comptonized photons has been used, motivated by spectral observations (Poutanen & Gierliński 2003; Leahy et al. 2009, 2011; Morsink & Leahy 2011). The Comp-

tonization model includes photon power law indices as well as a parametrized fan-beaming model. These models required two free parameters, one for the energy dependence and the second for the angular dependence of the emitted radiation. One of the issues seen with fitting the data from the accretion-powered X-ray pulsars is that the extra degree of freedom in the radiation's angular dependence leads to extra degeneracies amongst the geometric parameters. The result is fairly large regions of parameter space allowed by the fits, which do not strongly constrain the neutron star's equation of state.

Since thermonuclear X-ray burst oscillations can be well-reproduced with a simpler spectral model, we can fit models in a reduced but still physically motivated parameter space. Fewer free parameters gives fewer degeneracies among the parameters, and therefore allows for better constraints on the neutron star's mass and radius.

In the setup of our models, we assume that the inner boundary of the accretion disk is the neutron star's co-rotation radius (Ghosh & Lamb 1979). From this assumption, the accretion disk would block emission from a possible second antipodal hotspot (from an observer's perspective), so we only test models for one spot in the northern hemisphere of the neutron star. If the signature of an antipodal spot is detected, our code can be easily adapted to include a second hotspot.

We created synthetic data for a variety of different neutron star and spot parameters. Fitting multiple Poisson realizations for each synthetic pulse profile allows us to disentangle what uncertainty is due to random statistical fluctuations and what is due to inherent degeneracy between the parameters; understanding both is crucial for placing proper constraints on neutron star masses and radii. The pulse profile fitting was carried out with the Ferret optimization algorithm (Fiege 2010) to determine the acceptable range of masses and radii.

In this paper we show that evolutionary optimization algorithms are a powerful method of fitting neutron star pulse profiles, and we test the effects of changing various input parameters on how well we can recover the true neutron star mass and radius. In Section 2 we explain the details of constructing the pulse profiles, our parameter choices, and the Ferret algorithm. The results of the pulse profile fitting are described and examined in Section 3, and the conclusions are discussed in Section 4.

3.2 Pulse Profile Models

We construct the pulse profiles within the Schwarzschild + Doppler (S+D) approximation (Miller & Lamb 1998; Poutanen & Gierliński 2003) and the Oblate Schwarzschild (OS) approximation (Morsink et al. 2007). In the S+D and OS approximations, the metric exterior to the rotating neutron star is approximated by the Schwarzschild metric as described by Pechenick et al. (1983), adding the appropriate Doppler boost factors arising from the rotation of the star. In the S+D approximation, the surface of

the star is a sphere, while in the OS approximation the surface is an oblate spheroid. Cadeau et al. (2007) compared the S+D and OS results with pulse profiles generated from the exact metric and showed that the OS approximation is a good approximation for stars spinning with frequencies above 300 Hz. However, we continue to use the S+D approximation in many of our models in order to further explore the effect of using the wrong shape on the fits. At spin frequencies higher than 600 Hz, it may be necessary to use higher order approximations that make use of the star's quadrupole moment (Psaltis & Özel 2014), however this level of approximation is not necessary for the stars studied in this paper.

Pulse profiles can be constructed once 8 geometric parameters and a spectral emissivity model are specified. The 8 geometric parameters are the neutron star's mass M; equatorial radius R; spin frequency $\nu_{\rm spin}$; the observer's inclination angle i (as measured from the spin axis); the hotspot's co-latitude θ; the angular radius of the spot ρ; the distance to the star d; and a phase offset ϕ. In practice, the star's spin frequency will always be known, so there are only 7 geometric parameters. It is possible to add parameters describing a more complicated shape for the spot (Poutanen et al. 2009), but in this paper we only consider the simplest spot models, which are circular spots with uniform temperature. Due to the approximately universal nature of a spinning neutron star's shape (Morsink et al. 2007; Bauböck et al. 2013), inclusion of the star's oblate shape does not require any additional free parameters.

Thermonuclear X-ray burst oscillations can be spectrally modelled as a single-temperature blackbody with limb-darkening (Artigue et al. 2013). Once a prescription for the angular dependence has been selected, such as the Chandrasekhar (1960) limb-darkening model (approximated by the Hopf function), the only free parameter is the hotspot's temperature. The reduced parameter space required by this spectral model results in less degeneracy with the geometric parameters. For this reason, the oscillations from Type I X-ray bursts will be a major target for large-area, high-time-resolution X-ray telescopes like *ASTROSAT* (Singh et al. 2014), the soon-to-be launched *NICER* mission (Arzoumanian et al. 2014), and a future *LOFT*-like mission (Feroci et al. 2012).

The geometric parameters M, R, i, and θ have inherent degeneracies, so it is useful to refer to less-degenerate combinations of these parameters. In the spherical S+D approximation, i and θ only appear in the formulae in the combinations $\sin i \sin\theta$ and $\cos i \cos\theta$. As a result, in all fits there is a simple degeneracy that allows i and θ to be switched. Likewise, M and R are somewhat degenerate, so the dimensionless compactness ratio M/R can be better-constrained. These parameter combinations factor into the approximate bolometric pulse amplitude Amp (Beloborodov 2002),

$$\mathrm{Amp} = \frac{(1-2M/R)\sin i \sin\theta}{2M/R + (1-2M/R)\cos i \cos\theta}, \tag{3.1}$$

and the dimensionless projected velocity of the spot β,

$$\beta = \frac{2\pi R \nu_{\rm spin}}{\sqrt{1-2M/R}} \sin i \sin\theta, \tag{3.2}$$

in geometric units ($G = c = 1$). Due to the reduced degeneracy, it is possible to fit for $\sin i \sin\theta$, $\cos i \cos\theta$, M/R, Amp, and β better than the individual parameters. In our models, we consider an infinitesimally small spot for simplicity. However, the $i - \theta$ degeneracy can be partially broken for models with a large spot, in which case the spot would span a range in θ but not in i.

Equation (3.1) is only an approximate relation for the bolometric pulse amplitude. In reality, the pulse amplitude depends both on the emitted energy spectrum and the energy bands at which the observations are made. There is no simple formula for the dependence of the pulse amplitude on energy, but it can be computed numerically. Since the pulse amplitude Amp depends on the observed photon energy, observations of the burst oscillations in two or more energy bands can provide stronger constraints than suggested by equation (3.1); for this reason, two energy bands are used in this work. Additional information could be extracted with more energy bands, but we were unable to accommodate more bands. The projected velocity β controls the asymmetry in the pulse profile through the Doppler boosting effect, as well as the phase lags between the hard and soft energy bands. For higher values of β, the pulse profiles are more asymmetric in rise and fall times, showing the effects of higher harmonics.

Previous work (Morsink & Leahy 2011; Stevens 2013; Lo et al. 2013; Bauböck et al. 2015) has shown that for smaller spots, it is sufficient to compute the pulse profile assuming a point-source spot instead of an extended region. By adopting the point-source approximation, we simplify the calculation and do not make use of the angular radius parameter ρ, reducing the number of free geometric parameters to 6. Furthermore, we normalize the synthetic pulse profiles to have a mean of 1 to remove the dependence on the distance d, giving five free geometric parameters for each model.

Our goal is to constrain the neutron star's mass and radius based on fitting models to the pulse profile, and determine how the shape of the pulse profile affects how well we can constrain its M and R. By fitting synthetic data, the input parameters are known, so we can analyze how well the fitting can recover the true M and R. In this paper we investigate two sources of error in the constraints: degeneracy-related uncertainty and statistical uncertainty. Although the two types of uncertainties are coupled we have introduced two measures of the uncertainties $\sigma_{\rm degen}$ and σwhich are affected differently by the degeneracy and the statistics. The degeneracy-related uncertainty $\sigma_{\rm degen}$ arises from parameter degeneracies in each fit, whereas the standard deviation σ is examined by simulating multiple Poisson realizations of a model and determining the mean and standard deviation over all the fits. We explore how σ and

$\sigma_{\rm degen}$ are affected by the values of Amp and β for each synthetic pulse profile. We also quote the accuracy of the fit in M and R, comparing the mean best-fit value for each model with the true value.

3.2.1 Properties of Test Models

We computed pulse profiles for a set of nine test models with the parameter values listed in Table 3.1. The properties that were kept the same for all test models are the spot's temperature, the phase offset ϕ, the observer's inclination angle i and the distance to the star d. The spot emission model is a 2 keV blackbody (in the frame comoving with the neutron star's surface) with a limb-darkening atmosphere, approximated by the Hopf function (Chandrasekhar 1960), appropriate for a Thompson-scattered atmosphere. Since for real data we would independently have the spot temperature *at infinity* instead of in the star's comoving frame, this parameter should be allowed to vary within a narrow range (explored in Section 3.3.7), where the appropriate range would be determined by the temperature at infinity. The model assumes that any emission from the surface of the neutron star outside the spot is negligible. The observers inclination angle $i = 60^\circ$ and the phase offset $\phi = 0$ for all cases. The computed pulse profiles are normalized to an average flux of 1, so that the distance d to the star does not affect the pulse profiles.

The parameters that were changed for different test models are M, R, θ, the star's shape (spherical or oblate), and $\nu_{\rm spin}$. For oblate neutron star models, R is defined to be the radius of the neutron star at the spot. The formalism for the oblate model is detailed in Morsink et al. (2007), and does not require the addition of any extra parameters.

The parameters for the fiducial model, Model A, were chosen to be representative of the masses and radii of accreting millisecond neutron stars. Our fiducial mass $M = 1.6\,{\rm M}_\odot$ is larger than the mass $M = 1.4\,{\rm M}_\odot$ typically measured for slowly rotating radio pulsars, since we expect the neutron star has been spun up by and gained mass from accretion. The radius of 12 km is consistent with other radius estimations (Leahy 2004; Steiner et al. 2010). Rapid rotation with $\nu_{\rm spin} \geq 550\,{\rm Hz}$ is seen in burst oscillations for at least seven neutron stars that exhibit Type I X-ray bursts (Watts 2012). We have chosen $\nu_{\rm spin} = 600\,{\rm Hz}$ for the fiducial model since it is representative of these rapid rotators. The angles were chosen to provide a pulse amplitude (Amp = 0.373) comparable to the largest pulse amplitudes seen, in order to reduce the parameter degeneracy. In particular, the neutron star 4U 1636-536 has rms pulse amplitudes as high as 0.25 after subtracting the pre-burst emission (Strohmayer et al. 1998; Galloway et al. 2008), where Amp $= \sqrt{2}$ rms. The full set of theoretical parameters are given in the row labelled "A" in Table 3.1. The resulting pulse profiles in two energy bands are shown in Figure 3.1 with black and red solid curves labelled "True". Our fiducial model is similar to the model described as "low

Table 3.1: Summary of Test Models

Name	Description	Shape	$\nu_{\rm spin}$ Hz	M $M_\odot$	R km	i deg.	θ deg.	ϕ	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β
A	Fiducial	Sphere	600	1.6	12	60	20	0	0.197	0.296	0.470	0.373	0.057
AO	Oblate	Oblate	600	1.6	12	60	20	0	0.197	0.296	0.470	0.373	0.057
θ_{37}	$\theta = 37°$	Sphere	600	1.6	12	60	37	0	0.197	0.521	0.399	0.720	0.101
θ_{37}O	$\theta = 37°$	Oblate	600	1.6	12	60	37	0	0.197	0.521	0.399	0.720	0.101
θ_{60}	$\theta = 60°$	Sphere	600	1.6	12	60	60	0	0.197	0.750	0.250	1.305	0.145
$(M/R)_{hi}$	High M/R	Sphere	600	1.68	11	60	20.8	0	0.226	0.308	0.467	0.350	0.057
$(M/R)_{lo}$	Low M/R	Sphere	600	1.35	13	60	19.8	0	0.153	0.293	0.471	0.423	0.057
$\beta_{hi}(M/R)_{hi}$	High β, high M/R	Sphere	600	1.68	11	60	64	0	0.226	0.778	0.219	1.235	0.145
ν_{400}	Low $\nu_{\rm spin}$	Sphere	400	1.60	12	60	20	0	0.197	0.296	0.470	0.373	0.038

inclination" by Lo et al. (2013), except that they consider a slightly smaller R and a slower $\nu_{\rm spin}$. As will be shown in Section 3.3.5, small changes in R do not qualitatively change the results of this paper. Our choice of a faster $\nu_{\rm spin}$ improves the accuracy of the fits, but the faster $\nu_{\rm spin}$ is also more appropriate for the neutron stars that we hope to apply this method to.

Theoretical pulse profiles were calculated for the nine different test models. Each theoretical pulse profile was converted to a set of twenty synthetic observations by adding noise from a Poisson distribution to the pulse profile. The standard case (a low-noise model) assumes 25,000 photon counts per phase bin and no background count rate, as if the background were negligible in comparison with the hot spot emission. By not including a background, we are underestimating the Poisson fluctuations, and thus underestimating the error bars. This is done so we can test the suitability of evolutionary optimization on the best possible quality of synthetic data.

For a variation on the model θ_{60}, a higher noise level was used: $\theta_{60}\mathrm{C}_{6250}$ assumes 6250 photon counts per phase bin with no background. In total, there are ten sets of twenty simulated observed pulse shapes.

Each of the 200 simulated pulse profiles was then fit using the Ferret algorithm (described in the next subsection) which searches for different values of the free parameters in order to minimize the χ^2 fit statistic. Five parameters were allowed to vary within physical ranges: $1.0\,\mathrm{M}_\odot \leq M \leq 2.5\,\mathrm{M}_\odot$, $6.0\,\mathrm{km} \leq R \leq 16.0\,\mathrm{km}$, $0^\circ \leq (i,\, \theta) \leq 90^\circ$, and ϕ defined cyclically from 0 to 2π. The star's shape (spherical or oblate) was fixed to be the same in the fitting procedure as in the theoretical waveform. The temperature is kept fixed at 2 keV since it has been shown by Lo et al. (2013) that observations in multiple energy bands allows for a good determination of the gravitationally-redshifted spot temperature. We do not allow the spot size to vary, as explained in the previous subsection. For applications to real data, it would not be difficult to add variations in spot size (and shape), temperature, and to use unnormalized fluxes to allow for a distance measurement. However, the addition of the extra free parameters for the number of synthetic waveforms considered in this study would be impractical. We carried out preliminary trials (summarized in Section 3.3.7) allowing these parameters to vary but did not find much change in our final results.

Since there are two energy bands with 32 phase bins each, and five free parameters, there are 59 degrees of freedom (dof). The fitting yielded best-fit model parameters and also allowed computation of confidence regions in the M-R plane, from which we can assess the effect of various parameters on the uncertainties in parameter determination.

3.2.2 Parameter Fitting using Evolutionary Optimization

Evolutionary optimization algorithms provide a useful but heuristic approach to search through large parameter spaces, to optimize the fit of a model to data. Such algorithms use principles inspired by biology to evolve a population of candidate parameters sets, over many generations, toward an optimal solution. The heuristic nature of the search aims to sample the parameter space efficiently, but not completely, since exhaustive search is not practical for problems involving many parameters. Therefore, it cannot be guaranteed that the true, globally optimal solution has been found, although this would also be true for any other optimization algorithm, besides exhaustive search. Evolutionary algorithms have been well studied and found to be useful in many fields of science and engineering.

The oldest, and most commonly known, type of evolutionary optimizer is the genetic algorithm (Holland 1975; Goldberg 1989, 2002). Classic genetic algorithms encode their search parameters on a (usually) binary string (genotype), which is decoded into a model (phenotype). A population of candidate parameter sets is normally initialized as a set of random bit strings, which are expressed as a model, and in turn evaluated by a fitness function. This information is used to probabilistically select good parameter sets (individuals) to propagate to the next generation, in analogy to the "survival of the fittest" principle of Darwinian evolution. Parameter sets undergo random bitwise mutations to help perturb solutions into previously unexplored regions of parameter space. Information is shared between individuals by means of a crossover operator that cuts a pair of bit strings at a random position, and recombines them into a new "offspring" configuration, in analogy with sexual reproduction. This process of evaluation, selection, mutation, and selection is performed iteratively, over many generations, until a convergence criterion terminates the search. The genetic algorithm can be viewed as a directed stochastic search, which makes use of random noise to evade local minima, while a mostly deterministic selection operator pushes solutions toward an optimal solution.

In this paper, we used versions 5.3–5.5 of Ferret from the Qubist Global Optimization Toolbox for MATLAB (Fiege 2010; Rogers & Fiege 2011), a commercially available software package, to find the best fits to our synthetic data. This is an alternative to the Markov chain Monte Carlo approach in Lo et al. (2013) to fit pulse profiles. Ferret's development began in 2002, as a variant of the multi-objective genetic algorithm, which made use of a real-valued parameter encoding and real-valued mutation and selection operators, rather than the traditional binary encoding discussed above. Numerous other features have been added to Ferret since then, which go beyond the usual genetic algorithm paradigm. Most notably, the code contains a unique linkage-learning algorithm, which detects a certain type of non-linearity (linkage) between parameters, with the goal of dividing a large parameter space into several smaller and nearly independent subspaces, thus greatly simplifying the prob-

lem. This capability is discussed at length in the software user's manual (Fiege 2010). Ferret also contains an algorithm that allows several of its most important control parameters, including the typical strength of mutation and crossover events, to be automatically adapted and optimized during a run. This auto-adaptation capability is similar in spirit to the self-adaptation used in Evolution Strategies (ES) codes, although ES codes do not typically employ a crossover operator. Ferret does not adhere to the strict definition of a genetic algorithm, due to these enhancements and others, but the code still remains closer to the genetic algorithm paradigm than to others within the family of evolutionary optimizers.

As an example of Ferret's inner workings, consider the parameters i and θ. Ferret selects sets of values from the allowed region of parameter space ($0.01^\circ \leq i, \theta \leq 90^\circ$) for the population of a generation, and fits pulse profile models. It then keeps the best-fitting parameter sets and selects new sets to explore more of the parameter space. Within a few generations, Ferret discovers the degeneracy between i and θ, since their values can be swapped with minimal impact on the χ^2 fit. A more detailed explanation of Ferret in relation to pulse profile modelling can be found in Stevens (2013) Chapter 4. Evolutionary optimization algorithms have also been used in gravitational lensing (Rogers & Fiege 2011, 2012), medical physics (Fiege et al. 2011), star formation (Franzmann 2014), and X-ray spectral fitting (Rogers et al. 2015) applications.

Our problem is quite easy for Ferret, which consistently finds the global minimum χ^2 fit statistic within a few generations. We note that the true global minimum is known, since our results are based on artificial data tests. After finding the χ^2 minimum, the algorithm was allowed to run for approximately 100 generations, as Ferret accumulated solutions within the confidence region. No specific convergence criterion was implemented; rather, we terminated the run manually when it was evident from the software's graphical user interface that no further improvements were being made to the minimum χ^2, and enough solutions had been accumulated within the confidence region to make contours maps of acceptable quality. Ferret was executed in MATLAB R2011a on 12-core AMD Linux servers (dual socket Opteron 2439 SE, with 32Gb RAM, running Red Hat version 4.1.2), using Ferret's built-in parallel computing features. Each fit took approximately 8–20 hours.

3.3 Results

Pulse profiles were simulated as a set of Poisson realizations of a theoretical pulse profile model with known input parameters (see Section 3.2.1). Each Poisson realization was fit with Ferret to produce a set of best-fit parameters for the pulse profile along with confidence regions for the parameters. Since each theoretical pulse profile has multiple Poisson realizations, the average and standard deviation for the best-fit

parameters of each theoretical pulse profile are computed.

Here we discuss the fit results of the pulse profiles, which illustrate the effects of changing different system properties. Table 3.2 shows a summary of the fits to the 20 different Poisson realizations for each of the input models. The first row of each pair of rows shows the input parameters of the theoretical pulse shape model. The second row shows the means and standard deviations of the best-fit values from the fits to the 20 different Poisson realizations of each model. In addition to the fit parameters (M, R, i, θ, and ϕ), other useful measures of the model (M/R, $\sin i \sin\theta$, $\cos i \cos\theta$, Amp, and β) are given.

A brief overview of the entire set of fits is given first. Some parameters are well-determined and others poorly-determined. The standard deviation of any given parameter for each set of Poisson realizations is given as the error σ of that parameter. We take the difference between the mean of a parameter for each set and the input parameter as a measure of the accuracy of that parameter fit. The means and standard deviations are listed in Table 3.2, from which the percent error for precision and accuracy can be determined. The Ferret algorithm also computes contour regions as a measure of the degeneracy-related uncertainty $\sigma_{\rm degen}$. The σ errors are sometimes smaller than the 1-$\sigma_{\rm degen}$ region as shown in figures in the following subsections. These two values of error are measures of the degeneracy ($\sigma_{\rm degen}$) and the statistical fluctuations (σ), but they need not be the same. The 1-$\sigma_{\rm degen}$ contours determined by Ferret are models with different i, θ, and ϕ parameters that all provide an equivalently good fit for that M and R. In this way, the 1-$\sigma_{\rm degen}$ limit can be thought of as a measure of the degeneracy of the parameter space near the best-fit model. The standard deviations σ computed from the ensemble of Poisson realizations provide a measure of the error from the statistical fluctuations. However, the strong parameter degeneracy is certainly still an important factor in the standard deviation computation, and as such the reported σ values cannot be assumed to be purely statistical error.

We find that the projected velocity β is generally the most accurate and precise, with accuracy of 0.5–3% and precision of 1–6%. M and R are the next best, with with accuracy of 1–8%, but R has worse precision (1–21%) than M (1–13%). Compactness (M/R), $\sin i \sin\theta$, and Amp all have similar accuracies ($\sim$1–8%) and precisions ($\sim$1–16%). The accuracy of $\cos i \cos\theta$ is $\sim$4–30% and precision is $\sim$2–20%. The least-well-determined are i and θ, with accuracies of $\sim$4–30% and precisions of $\sim$5–50% due to the degeneracy in angles.

In the following subsections we discuss in detail the results from fitting the fiducial model A, and then discuss trends we find by comparing the fits from model A with the fits for the models with modified parameters.

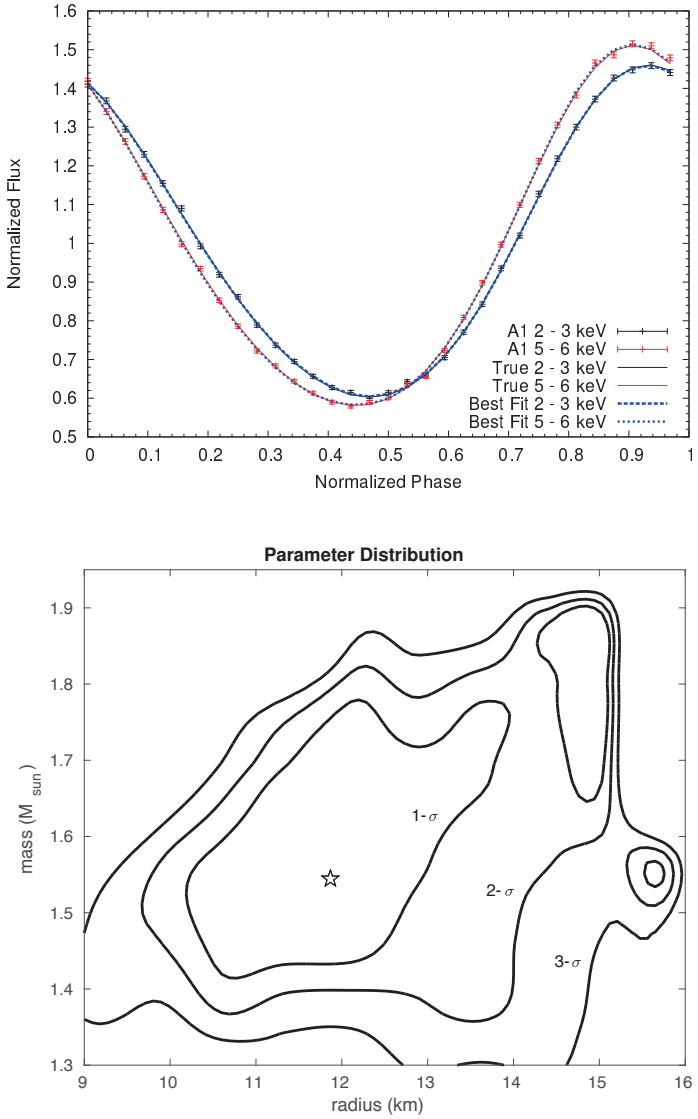

Figure 3.1: Top: Pulse profiles for model A in two energy bands. The solid black and red curves (labeled "True") correspond to the theoretical model A given in Table 3.1. The black and red points (labeled "A1") with error bars correspond to the addition of noise to model A with one Poisson realization assuming an average count-rate of 25,000 per phase bin. The blue dashed and dotted curves correspond to the best fit to model A1 with parameters given by the values in row "A1" in Table 3.3. This best fit has $\chi^2 = 58.2$ for 59 dof. Bottom: χ^2 contours in the $M - R$ plane for fitting synthetic data from model A1. The best-fit model is shown with a star and corresponds to the row labelled "1" in Table 3.3. The "true" values of mass and radius are $1.6\,M_\odot$ and 12 km. The contours show the 1-, 2-, and 3-$\sigma_{\rm degen}$ confidence regions. From this figure, it can be seen that the 1-$\sigma_{\rm degen}$ error region for radius corresponds to approximately 1.9 km, while the 1-$\sigma_{\rm degen}$ limit for mass spans $0.18\,M_\odot$. The contours for 1-, 2-, and 3-$\sigma_{\rm degen}$ correspond to values of $\Delta\chi^2$ (above $\chi^2_{\rm min}$) of 2.3, 6.2, and 11.8 respectively.

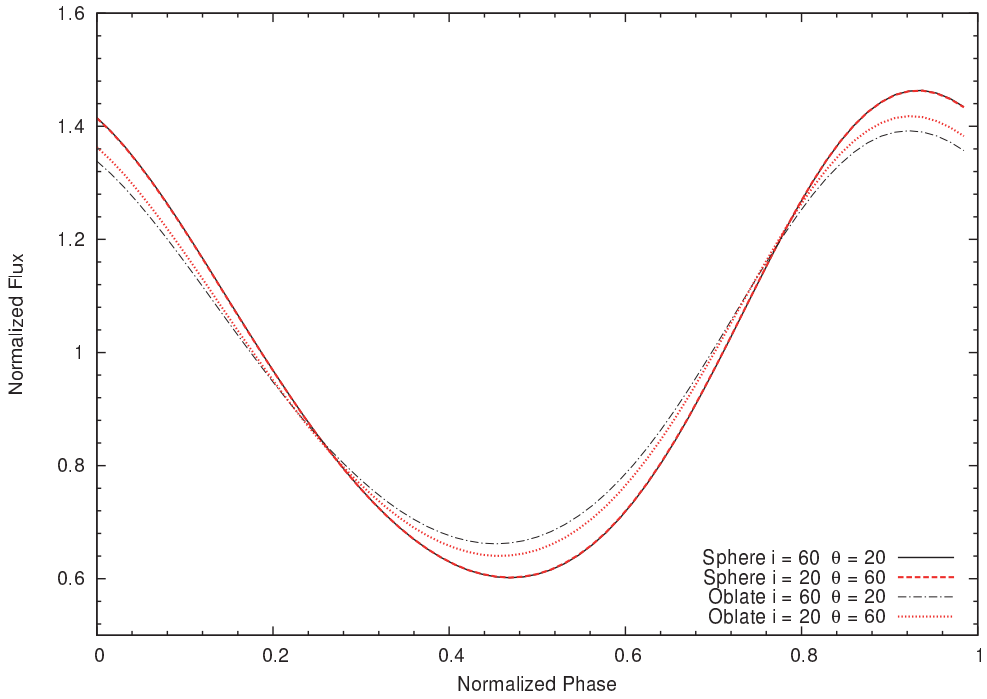

Figure 3.2: Normalized low-energy band pulse profiles for four neutron star models showing the observer inclination–spot co-latitude degeneracy for spherical and oblate stars. The solid black curve represents Model A, while the overlapping dashed red curve represents a star with the same parameters as Model A, but with inclination i and co-latitude θ swapped. The black dot-dashed curve represents Model AO, the oblate version of Model A. The red dotted curve represents a star with the same parameters as Model AO, again with inclination i and co-latitude θ swapped.

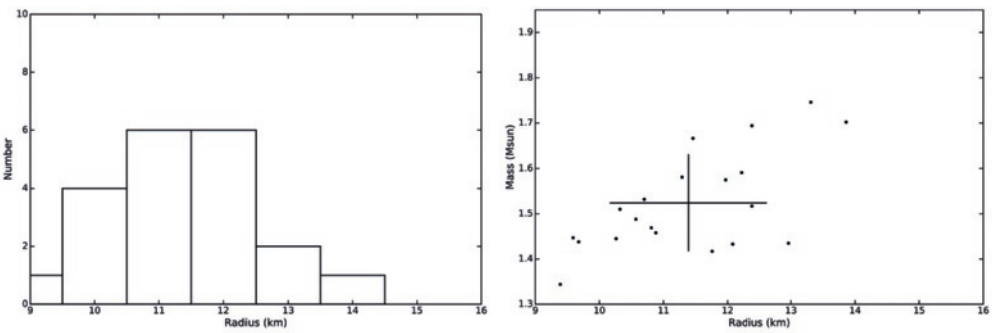

Figure 3.3: Left: Histogram of best-fit values of R for model A. Right: Best fit values of M and R for model A. Each dot is one fit to a Poisson realization, as shown in Table 3.3. The black cross shows the average with 1-σ error bars.

Table 3.2: Summary of Statistical Properties of Fits

Model χ^2(59 dof)	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
A 59.3±7.6	1.60 1.52±0.11	12.0 11.4±1.2	60.0 34.7±15.0	20.0 41.8±15.2	0.0 2.8±4.6	0.1969 0.199±0.016	0.296 0.318±0.032	0.470 0.549±0.120	0.373 0.363±0.025	5.74 5.81±0.20
AO 59.1±7.0	1.60 1.65±0.06	12.0 11.7±0.9	60.0 49.3±18.2	20.0 28.3±11.9	0.0 0.8±5.3	0.1969 0.209±0.016	0.296 0.299±0.019	0.470 0.511±0.142	0.373 0.353±0.060	5.74 5.76±0.18
θ_{37} 61.1±9.9	1.60 1.62±0.08	12.0 12.4±0.9	60.0 49.1±13.8	37.0 47.9±13.0	0.0 -0.7±2.9	0.1969 0.194±0.009	0.521 0.509±0.033	0.399 0.386±0.064	0.720 0.726±0.029	10.10 10.09±0.17
θ_{37}O 58.9±8.1	1.60 1.58±0.04	12.0 11.8±0.4	60.0 49.8±9.7	37.0 46.3±7.8	0.0 1.1±1.5	0.1969 0.198±0.008	0.521 0.528±0.016	0.399 0.424±0.029	0.720 0.704±0.037	10.10 10.08±0.14
θ_{60} 57.7±8.4	1.60 1.61±0.01	12.0 12.1±0.1	60.0 61.5±4.0	60.0 58.5±3.7	0.0 -0.3±0.6	0.1969 0.196±0.003	0.750 0.745±0.008	0.250 0.245±0.010	1.305 1.309±0.011	14.54 14.53±0.09
$\theta_{60}C_{6250}$ 58.2±7.7	1.60 1.62±0.03	12.0 12.5±0.5	60.0 58.6±6.6	60.0 60.2±6.6	0.0 -0.7±1.5	0.1969 0.192±0.008	0.750 0.728±0.027	0.250 0.247±0.023	1.305 1.303±0.028	14.54 14.52±0.18
$(M/R)_{hi}$ 60.3±7.7	1.68 1.61±0.11	11.0 10.4±1.0	60.0 34.9±13.1	20.8 42.7±16.2	0.0 2.1±4.8	0.2255 0.228±0.012	0.308 0.328±0.029	0.467 0.542±0.116	0.350 0.343±0.022	5.74 5.80±0.22
$(M/R)_{lo}$ 62.5±10.1	1.35 1.28±0.10	13.0 13.0±1.2	60.0 34.3±17.2	19.8 42.8±17.1	0.0 1.9±3.6	0.1534 0.146±0.018	0.293 0.299±0.027	0.471 0.522±0.086	0.423 0.411±0.022	5.75 5.78±0.15
$\beta_{hi}(M/R)_{hi}$ 53.8±10.7	1.68 1.70±0.02	11.0 11.2±0.3	60.0 61.9±6.5	64.0 61.6±6.8	0.0 -0.8±0.8	0.2255 0.223±0.004	0.778 0.762±0.016	0.219 0.211±0.016	1.235 1.241±0.016	14.53 14.47±0.09
ν_{400} 58.1±9.9	1.60 1.55±0.21	12.0 13.0±2.5	60.0 35.7±17.6	20.0 40.7±18.8	0.0 1.0±5.2	0.1969 0.181±0.034	0.296 0.292±0.051	0.470 0.525±0.155	0.373 0.361±0.035	3.83 3.88±0.22

3.3.1 Model A: Fiducial Model

Twenty realizations of the Poisson noise were made to simulate observations of the fiducial Model A. The resulting flux values with error bars for the first Poisson realization of the data are shown as black and red points labelled "A1" on Figure 3.1. The data points with Poisson errors were used as input to Ferret, which searched for the best fit to the A1 dataset by minimizing χ^2. In the case of the A1 dataset, the pulse profiles that best fit the data are displayed with blue dashed and dotted curves in Figure 3.1. The parameters corresponding to the best fit of the A1 dataset ($\chi^2 = 58.2$) are shown in the row labelled "1" in Table 3.3. For reference we have included the derived quantities M/R, $\sin i \sin\theta$, $\cos i \cos\theta$, Amp, and β to this table.

The best fit model (star) and contours of constant $\chi^2 = \chi^2_{\rm min} + \Delta\chi^2$ (black curves) for the A1 dataset are shown in the $M - R$ plane in Figure 3.1. The contours correspond to values of $\Delta\chi^2 = 2.3, 6.2$, and 11.8, corresponding to 1-, 2-, and 3-$\sigma_{\rm degen}$ confidence levels respectively in mass-radius space for two free parameters, M and R. This plot uses the profile likelihood (Murphy & Van Der Vaart 2000) to eliminate the nuisance parameters i, θ, and ϕ; a grid is defined in M and R, and for each (M, R) grid point, Ferret finds the lowest χ^2 within each grid cell, allowing for any values of i, θ, and ϕ. The 1-$\sigma_{\rm degen}$ confidence region (calculated by finding the maximum and minimum values on the 1-$\sigma_{\rm degen}$ contour and dividing by 2) gives an approximate 1-$\sigma_{\rm degen}$ error of 1.9 km for the radius and a 0.18 $M_\odot$ error for the mass. Note that since the contours are not ellipses, these are only approximate 1-$\sigma_{\rm degen}$ limits.

Ferret was used to fit each of the 20 Poisson realizations of model A. The independent best fit results for the 20 different realizations are shown in Table 3.3. The average and standard deviation for each parameter are displayed at the bottom of Table 3.3, and also appear in the second line of Table 3.2. For the individual angles i and θ, it can be seen from Table 3.3 that the determinations of these angles are very poor. This is due to the already well-known degeneracy (for spherical stars; Poutanen & Gierliński 2003) which occurs since the equations for light-bending and the Doppler effect only depend on the combinations $\sin i \sin\theta$ and $\cos i \cos\theta$ and not on their individual values. The $i - \theta$ degeneracy can be seen in Figure 3.2, where the normalized low-energy band pulse profiles for model A, and another model with i and θ swapped (both labelled "Sphere") are indistinguishable. For many of the Poisson realizations, the best-fit values for i and θ shown in Table 3.3 are swapped from their true values, and as a result, the average and standard deviations for the individual angles are really not meaningful, except to illustrate that it is the trigonometric combinations of the angles that can be reliably determined. However, since there are independent methods for constraining i and θ through optical (Wang et al. 2013) or gamma-ray (Venter et al. 2012) observations, it is still useful to discuss these two angles separately.

To ensure that a suitable number of independent realizations of the data were

made, a histogram of the best-fit radii is plotted in Figure 3.3. The resulting histogram is approximately Gaussian (70% are within 1 standard deviation, 95% within 2 standard deviations, 100% within 3 standard deviations), indicating that 20 trials are sufficient to illustrate general trends. A scatter plot, also shown in Figure 3.3, shows the values of mass and radius (red crosses) for the 20 different random realizations of the data. The large black cross indicates the average and standard deviation values for the mass and radius fits.

As previously stated, parameter degeneracy plays a major role in all aspects of uncertainty in our results, even the standard deviation. By testing 20 different realizations of the data for each model (a relatively small number), we still find that some models converge on the true parameter values while others do not. The models with stronger parameter degeneracy have larger standard deviations on the parameters. This provides an initial assessment of which types of models and pulse profile shapes have stronger degeneracies. Even with very small error bars on the pulse profile, there would still be relatively large standard deviations over many models due to the inherent degeneracies.

3.3.2 Effect of Changing Spot Co-latitude

In this section, we examine the effect of changing the hotspot's co-latitude θ on the accuracy and precision of the pulse profile fits. By changing θ while keeping M, R, and i constant, we alter the projected velocity β and the approximate pulse amplitude Amp (see Table 1). Models A (with $\theta = 20°$), θ_{37}, and θ_{60} have increasing values of θ, β, and Amp. Due to the i-θ degeneracy, this is also equivalent to keeping θ fixed and varying i.

The $\theta_{37}1$ pulse profile (i.e., for the first Poisson realization of the θ_{37} model) is shown in the top panel of Figure 3.4 and the constant χ^2 contours are shown in the bottom panel. This case is one of the "outlier" pulse shapes (of the 20 Poisson realizations) with a very large best-fit R that only includes the true R at the 3-$\sigma_{\rm degen}$ level. The 1-$\sigma_{\rm degen}$ error regions for the $\theta_{37}1$ model (from the bottom panel of Figure 3.4) are smaller than the A1 model, close to 1.1 km for R and 0.11 $M_\odot$ for M. The 1-$\sigma_{\rm degen}$ regions for this particular Poisson realization are somewhat larger than the standard deviation 1-σ computed for the ensemble of θ_{37} models.

Similarly, the $\theta_{60}1$ pulse profile (first Poisson realization of the θ_{60} model) is shown in the top panel of Figure 3.5 and the constant χ^2 contours are shown in the bottom panel. The 1-$\sigma_{\rm degen}$ error regions for the $\theta_{60}1$ model are 0.7 km for R and 0.06 $M_\odot$ for M, larger than the ensemble standard deviation 1-σ by a factor of about 6.

We find a strong trend as θ increases: the average M and R fit values are closer to the true values and the standard deviation is smaller. For example, σ in M decreases from 0.11 to 0.08 to 0.01 $M_\odot$ for models Models A, θ_{37}, and θ_{60}. This improvement in fitting accuracy is due to models θ_{37} and θ_{60} having a larger Amp and β than model

A. This is consistent with the general trend seen by Lo et al. (2013) in individual model fits.

During an X-ray burst, θ is expected to change. This set of models roughly approximates this process, which has been explored in detail by Mahmoodifar & Strohmayer (2016). If a series of *different* burst oscillation pulse shapes are found for a neutron star during the *same* burst, future work could include a simultaneous multi-epoch fit, allowing θ (and spot size ρ) to be dependent on epoch.

3.3.3 Effect of Oblateness

Rotation alters the shape of a neutron star, making it an oblate spheroid. Although the change in shape is small for stars spinning at rates seen in accreting systems, the alteration in the star's shape changes the positions on the neutron star's surface for which photons can reach the observer (Morsink et al. 2007), leading to large changes in the pulse profile. As an example, a 1.6 $M_\odot$ neutron star with an interior given by the APR equation of state (Akmal et al. 1998) has an equatorial radius of 11.7 km and a ratio of polar to equatorial radii of 0.93. This oblate geometry has a larger effect on the pulse profile than other effects due to rotation, such as frame-dragging, as has been discussed in detail in, e.g., Morsink et al. (2007).

In this section, we investigate the effect of the oblate shape on the accuracy and precision of pulse profile fitting and reproducing the input parameters. To do this we construct oblate versions of two models, A and θ_{37}. The oblate versions, designated with the letter "O" are constructed so that the radius of the star at the location of the spot is the same as for the corresponding spherical model. This means that the values of the star's compactness M/R and the projected velocity β are the same for the spherical and oblate models. As a result of this definition for the radius, the star's equatorial radius is larger than that listed in Table 3.1. The general trend is for the oblate neutron star's pulse profile to have a smaller pulse amplitude than the spherical neutron star with the same parameters (when i and θ are in the same hemisphere), since the visibility condition makes it easier to see the far side of the neutron star when the star is oblate (Cadeau et al. 2007; Morsink et al. 2007). In cases where the spot is visible for all phases (as in both models A and θ_{37}), consider the emission when the spot and the observer have the same azimuthal angle. At this moment, the light from the spherical star is emitted close to the normal to the surface. For the oblate star, the light is emitted in the same direction in space, but it is at an angle further from the surface's normal, due to the tilt of the surface. Since the intensity of light is proportional to the cosine of the angle between the original direction of emission and the normal to the surface, the spherical star's spot appears brighter at this phase. The opposite is true when the spot is on the opposite side of the star from the observer. In this case the light is emitted close to the tangent to the surface for the spherical star, and closer to the normal for the oblate star, making the spot

appear dimmer for the spherical star at this phase. The overall effect is to create a less modulated pulse profile for the oblate star.

In model AO, the neutron star's equatorial radius is 12.7 km, while the radius at the spot's latitude is 12 km. As a comparison with the spherical model, the low-energy band pulse profiles for both AO and A are plotted as black dot-dashed and solid curves respectively in Figure 3.2. The high energy band (not shown for clarity) has a similar decrease in modulation. In model θ_{37}O the equatorial radius is 12.5 km, and has a lower pulse amplitude than model θ_{37}. Since the pulse profiles and $M - R$ confidence regions for models AO and θ_{37}O look quite similar to those for models A and θ_{37}, they are not shown here. The mean and standard deviation of the best fit parameters for the 20 realizations for each oblate model are given in Table 3.2. For most parameters, the σ and accuracies are smaller for the oblate case than for the fiducial spherical case (we also found that the 1-$\sigma_{\rm degen}$ regions were smaller by about a half). A similar improvement in the precision of the results was also noted by Miller & Lamb (2015), however since they only considered one Poisson realization of the data they could not rule out the possibility that this was due to a statistical fluctuation. In our case, since we are comparing a sample of Poisson realizations for each model, the increase in precision and accuracy is most likely due to the properties of the oblate model pulse profiles.

The improvement in the accuracy and precision in the determination of most of the parameters for oblate models is most likely due to the partial lifting of the degeneracy between i and θ. For an oblate star, the direction that the normal to the surface points in depends on the shape of the star which is a function of θ, but is independent of i. This introduces small differences between the normalized pulse profiles for models with i and θ switched, while for spherical stars, the normalized pulse profiles are the same when the angles are switched. This is illustrated in Figure 3.2 where the two oblate models representing the swapped inclination and spot angles, shown with black dot-dashed and red dotted curves, are clearly different from each other. However, since they are still fairly similar to each other, there is still a partial degeneracy when the angles are swapped.

There is a similar but much smaller lifting of degeneracy for large hot spots on spherical stars. A large spot extends over a range of θ values while the observer's inclination i is just one fixed value, so swapping the two angles if the spot is large will not yield the same pulse shape. However, the magnitude of the effect is much smaller than the magnitude of the change that occurs when the angles are swapped on an oblate star.

We also tested the effect of using the wrong shape model by using the oblate data corresponding the 20 AO models as input and fitting them with pulse shapes for spherical stars. In this case, χ^2 increased a small amount but the best-fit values for M and R became very inaccurate. For the set of fits to the Poisson realizations, the average fit returned $\chi^2 = 64.7$, $M = 1.91\,{\rm M}_\odot$, and $R = 12.3$ km (which should

be compared with the values in the 4th row of Table 3.2). The average M and R fits are accurate to 19.4% and 2.5%, respectively. Similar results were found by Miller & Lamb (2015). In practice, this is not a problem, since we know that the stars are actually oblate, so the correct shape can be used. The spherical shape model is computationally somewhat cheaper, hence it is used for the other models tested in this paper. Furthermore, the oblate shape model used is still an approximation. It would be worthwhile, in future research, to compare the slightly different shape models that have been used in this paper with other models used by other groups (such as Bauböck et al. 2013; Miller & Lamb 2015).

3.3.4 Effect of Photon Count-Rate

The average photon count-rate per phase bin of 25000 is the expected best case scenario, assuming typical burst flux, large detector effective area, and long observation time (Feroci et al. 2012), and similar to the count-rate used in other pulse profile analyses (Lo et al. 2013; Psaltis et al. 2014). The pulse profile is noisier when fewer counts are detected. We tested the effect of statistical noise on model θ_{60} by generating different Poisson realizations, assuming a reduced count-rate, of the same theoretical pulse profile. Comparing model θ_{60} with the reduced count-rate model, $\theta_{60}C_{6250}$, allows us to explicitly compare the degeneracy-related uncertainty versus error. The average counts per bin are 25000 and 6250 for models θ_{60} and $\theta_{60}C_{6250}$, respectively.

The pulse profile for model $\theta_{60}C_{6250}$ looks very much like that for model θ_{60} but with larger errors, so it is not shown. The χ^2 contours for model $\theta_{60}C_{6250}1$ are shown in Figure 3.6. The size of the 1-$\sigma_{\rm degen}$ contours in Figures 3.5 and 3.6 are very similar, suggesting that the confidence contours for one particular mock observation are mainly dominated by the parameter degeneracy.

From Table 3.2, the standard deviations of the parameters for the low count-rate model are $\simeq 2$–3 times as large as those of model θ_{60}. The factor of two is expected purely on statistics (from the factor 4 reduction in counts). The actual degradation is somewhat worse, likely because partial degeneracy between the parameters makes the parameter determination degrade more rapidly than the errors. Model $\theta_{60}C_{6250}$ still gives well-determined parameters even at the lower count-rate value. The accuracies in M and R determination are 2% and 4% respectively, for the low count-rate case.

3.3.5 Effect of Compactness

Changes in the compactness ratio (M/R) are expected to impact the quality of the fits, since the compactness affects the pulse amplitude Amp (evident in equation (3.1)). Increasing the compactness decreases Amp, which might be expected to decrease the precision and/or accuracy of the fits. To test this effect, we generated models with

differing values for the compactness ratio, but the same values of the projected spot velocity β.

We created two models similar to model A with a larger and smaller compactness ratio, $(M/R)_{hi}$ and $(M/R)_{lo}$, and generated a set of pulse shapes with Poisson noise. The results of the fits can be seen in Table 3.2. All three models with the same value of β have similar accuracy and precision for most of the parameters. For instance, the accuracy of determining the mass ranges from 4–5%, while the precision ranges from 6–7%, so the change in compactness does not greatly affect the fits. However, while there is little change in the precision of the radius measurement (9–10%), in the case of the low compactness model the accuracy was improved.

As a further test on the effect of compactness, model $\beta_{hi}(M/R)_{hi}$ was created to compare with the high spot co-latitude model θ_{60}. The parameters were the same except that θ was adjusted to give the same β as model θ_{60}. This also necessitated that Amp was somewhat lower than model θ_{60} (see equations (3.1) and (3.2)). Previously for model θ_{60}, the accuracy and precision for both M and R were better than 1%; by fitting model $\beta_{hi}(M/R)_{hi}$, we found that increasing the compactness degrades both the accuracy and precision in M and R, but only so that they range from 1–3%.

From these results, given a value of β, it appears that the accuracy and precision to which the other parameters can be determined does not strongly depend on the compactness ratio of the star. Thus, our fit results would not be significantly different if we had chosen a different M and R for the fiducial model A.

3.3.6 Effect of Spin Frequency

The last parameter we adjusted to be different from the fiducial model was $\nu_{\rm spin}$. We compared model A ($\nu_{\rm spin} = 600\,{\rm Hz}$) with a 400 Hz model, labelled ν_{400}. Model ν_{400} has all parameters the same as model A except $\nu_{\rm spin}$, and hence a different projected spot velocity (β) (see Table 3.1). This model has parameter values very similar to the "low inclination" model presented by Lo et al. (2013).

The fits results are summarized in Table 3.2. The precision (standard deviations) and accuracy (difference of mean and true values) are worse for model ν_{400}, as expected. For M, R, M/R, and Amp, the precisions are about twice as large for the 400 Hz model as for the standard 600 Hz model. However, β has the same absolute precision for the 400 Hz model as for the 600 Hz model.

3.3.7 Effect of Incorrect Model or Additional Free Parameters

Along with the sets of trials summarized in the previous subsections, we tested some of the model assumptions used in the fits. These tests include the assumptions we have made about the spectral emissivity, spot size, presence of a constant background flux, and the star's shape. Errors are not reported for the tests of this subsection since

fits were only carried out on the A1 synthetic data set, the first Poisson realization of the Model A theoretical curve.

We first tested the effect of using the wrong atmosphere model in the fits. To do this, we used the synthetic data corresponding to the A1 model (see Figure 3.1) as input to Ferret. The synthetic data were generated using the Hopf limb-darkening function, but Ferret was run assuming a perfectly isotropic blackbody spectrum. The algorithm was unable to find a good fit after 300 generations (about double the number of generations normally required) and the best fit had $\chi^2 = 316.6$ for 59 dof. The best-fit M and R were $1.21\,\mathrm{M}_\odot$ and 11.3 km, which are quite inaccurate (24.4% for M, 5.8% for R). As was seen by Lo et al. (2013), one does not get a good fit using an incorrect atmosphere model.

We next tested fitting a variable spot size model to a pulse profile that was generated with an infinitesimally small spot. The infinitesimal size of the spot is a necessary assumption due to the large number of trials carried out in this work. In order to model a larger spot, the spot has to be cut into a number of segments, and separate computations of the deflection angles must be computed, increasing the run time linearly with the number of segments. The pulse shape for a large spot tends to have a lower pulse amplitude than a small spot centred at the same latitude, due to the effect of averaging over many latitudes. As a test of this assumption, we used the A1 data set as input to the Ferret algorithm, but added the angular radius ρ of the spot as a free parameter that was allowed to vary between 1°and 60°. The resulting best fit has $\chi^2 = 59$ (for 58 dof) and converged on an angular radius of $\rho = 1°$, the smallest angle allowed. The best-fit M and R for this case are $1.51\,\mathrm{M}_\odot$ and 11.6 km. While these values are less accurate (5.6% and 3.3%, respectively) than the best-fit value shown in the first row of Table 3.3 corresponding to a fit with an infinitesimal spot, they are within the 1-σ limits computed for the set of model A fits. Since using multiple spot segments leads to a significantly longer computation time, it was not feasible to use variable spot sizes in this paper. However, in future work when real data are being fitted, it will be necessary to allow the spot size to vary, and to compute the observed flux from multiple spot segments.

We also tested the effect of allowing the temperature of the spot to be a free parameter in the fitting program. The theoretical model was constructed with a spot temperature of 2 keV (as measured in the comoving frame at the star's surface). The Ferret program was used to fit the data with the addition of a local temperature parameter that was allowed to vary between 1 and 3 keV. The best fit returned $\chi^2 = 57.9$ (for 58 dof), $M = 1.67\,\mathrm{M}_\odot$, $R = 11.8$ km, and a temperature of 2.1 keV. The accuracies in M and R are 4.8% and 1.7%, respectively. Although these M and R values are within the 1-σ limits for this data set, the addition of the extra parameter does degrade the accuracy. This is most likely due to the degeneracy introduced since we measure the redshifted temperature with the light curves. For the case of real data, it would be important to allow the temperature to vary in the fits, since the

temperature of the spot in the comoving frame is unknown. Furthermore, making use of multiple energy bands would improve the accuracy of the temperature measured at infinity.

The effect of an unknown background flux has been investigated in detail by Lo et al. (2013), and for comparison we test it for one simple case. The synthetic data set A1 was constructed without a background count rate. We used this as input and allowed Ferret to add two new parameters corresponding to a constant (in time) background flux in each energy band. These background values were allowed to vary between 0 and 1 (recall that our pulse profiles are normalized to 1). The resulting fit had $\chi^2 = 57.7$ (for 57 dof), and the best-fit M and R were $1.52\,\mathrm{M}_\odot$ and 11.7 km. M and R were accurate to 5.0% and 2.5%, respectively. It is possible to add a more realistic background model by adding emission at a lower temperature from the rest of the star (Lo et al. 2013; Elshamouty et al. 2016) or by adding light scattered from the disk (Morsink & Leahy 2011), but at this point it is not obvious what the most realistic model would be. Real data are also likely to contain non-negligible emission from the surface of the neutron star outside the spot region and from the accretion disk.

The effect of these tests on the best-fit values of M and R are summarized in Table 3.4. These tests highlight the importance of having as realistic an atmosphere model as possible, since fitting with the wrong model drastically affected the quality and accuracy of the fits. Thus, this method provides a sensitive test for atmosphere models. Adding a free parameter for the temperature, unknown background flux, or spot size did not significantly detract from the accuracy in M and R, and our code was able to replicate the additional parameter quite well. Ideally, for real data we would allow these parameters to also vary in the fits, even if it is not feasible to do so in the present study.

3.4 Discussion and Conclusion

We calculated pulse profiles for simulated thermonuclear burst oscillations from a rapidly rotating neutron star as would be detected by a next-generation X-ray timing observatory, such as *ASTROSAT*, *NICER*, or *LOFT*. The input neutron star parameters include mass M, radius R, emitting spot co-latitude θ, and observer inclination i. We created 20 Poisson realizations of each test pulse profile, which allows us to determine the errors in the derived model parameters, focusing on M and R. The resulting pulse profiles were fitted with Ferret to analyze how well the input parameters could be recovered (in both standard deviation and degeneracy-related uncertainty).

We find that the best-determined parameter is the projected velocity β of the spot along the observer's line-of-sight. The next best-determined are M and R. Compactness (M/R), $\sin i \sin\theta$, and the pulse amplitude Amp are also well-determined, but

$\cos i \cos\theta$, i, and θ are poorly determined. It is clear that more rapidly rotating neutron stars produce pulse profiles with strong harmonics, which yield better constraints on their parameters. Also, neutron stars viewed at larger i with spots at a larger θ produce more modulated pulse profiles and therefore better constraints. However, we also found that count-rate is important, so the best targets will be those systems which have the optimal combination of large $\nu_{\rm spin}$, large θ and i, and bright pulsations. For our best cases presented here, M and R were determined to 1% accuracy and precision, but even for many of the less optimum cases M and R were determined to $\sim 5\%$ accuracy and precision, which is a very valuable result.

We carried out a number of parameter comparison tests, to see how different input parameters can affect the constraints on M and R. The more asymmetry and larger pulse amplitude, the better the accuracy and precision in M and R from pulse profile fitting. The asymmetry in the pulse profile is controlled by β; a larger β results from increases in θ, i, and $\nu_{\rm spin}$. A larger θ or i also gives a larger Amp. Compactness (M/R) has a small effect on parameter determination; for more compact stars it is somewhat more difficult to determine parameters. This is caused by the increased visibility of the surface by the observer and the resulting decreased pulse amplitude. Including oblateness in the pulse profile model improves the accuracy and precision of M and R determinations. This is due to a reduced $i - \theta$ degeneracy in the oblate models.

Photon count-rate has a critical effect on parameter determination. On simple grounds, increasing the counts by a factor of four reduces the error by a factor of two. However, in practise we found a factor of three improvement in parameter determination. The extra gain is likely the result of reduced parameter degeneracy for data with smaller errors.

We assumed the spot temperature to be a known definite quantity. In practice, this is determined by a spectral fit to the data using multiple energy bands. We calculated the simulated pulse profiles using only two energy bands, so we have underestimated the uncertainties in M and R that enter through the uncertainty of the temperature. In principle, it is not difficult to use more energy bands (e.g., Lo et al. 2013). However, there is a trade off that the signal-to-noise in each energy band falls as more, and thus smaller, energy bands are used. Determining an optimal number of energy bands that allows a determination of the spectrum while providing enough statistics to constrain the star's properties should be the topic of a future study.

We also took the background count rate to be negligible in comparison with the count rate from the neutron star hot spot. For a real observation, the background comes from three sources: instrument background, sky background, and source contribution to background. Instrument background depends strongly on the instrument design and mode of observation. Source background can consist of emission from the surface of the neutron star outside of the spot region and from the accretion disk. However, subtracting the persistent pre-burst emission from the X-ray burst flux is

not an appropriate source-background subtraction. As shown in Worpel et al. (2013), the persistent emission can be different during the burst, possibly due to an increased accretion rate onto the neutron star. A method such as weighted-photons Bayesian Blocks (Worpel & Schwope 2015) should be incorporated into the analysis pipeline to subtract the appropriate level of source-background emission.

A limitation of our study is that we did not include the effect of frequency drift, which is normally seen during the rise of an X-ray burst (Watts 2012). When dealing with real data, it would be important to either remove sections of the data where frequency drift is observed, or to use a reliable method for modelling the pulse shape with phase offsets to account for the drift. An additional complication for the sources that also have accretion-powered pulsations is that the pulsations may contaminate the X-ray burst oscillations. At this time these are open problems associated with modelling X-ray burst oscillations.

In this paper we have demonstrated that an evolutionary optimization and parameter search methods can be used to effectively constrain the mass and radius of a neutron star with pulsed emission. We focussed on burst oscillations seen in Type I X-ray bursts, however the routines can also be used to model accretion-powered pulsations or rotation-powered pulsations once the appropriate spectrum and beaming functions have been changed. In particular, the rotation-powered pulsars that will be studied by *NICER* (see Özel et al. 2016) can be analyzed using the methods discussed in this paper.

Acknowledgements

This research was supported by grants from the Natural Sciences and Engineering Research Council of Canada to J.D.F., D.A.L., and S.M.M. A.L.S. acknowledges a travel grant from the LKBF (Leids Kerkhoven-Bosscha Fonds). We thank Arash Bahramian, Frederick Lamb, M. Coleman Miller, Adam Rogers, and Christoph Weniger for useful discussions. We also thank the *LOFT* Dense Matter science working group for early code comparisons to validate the modelling.

Appendix 3A

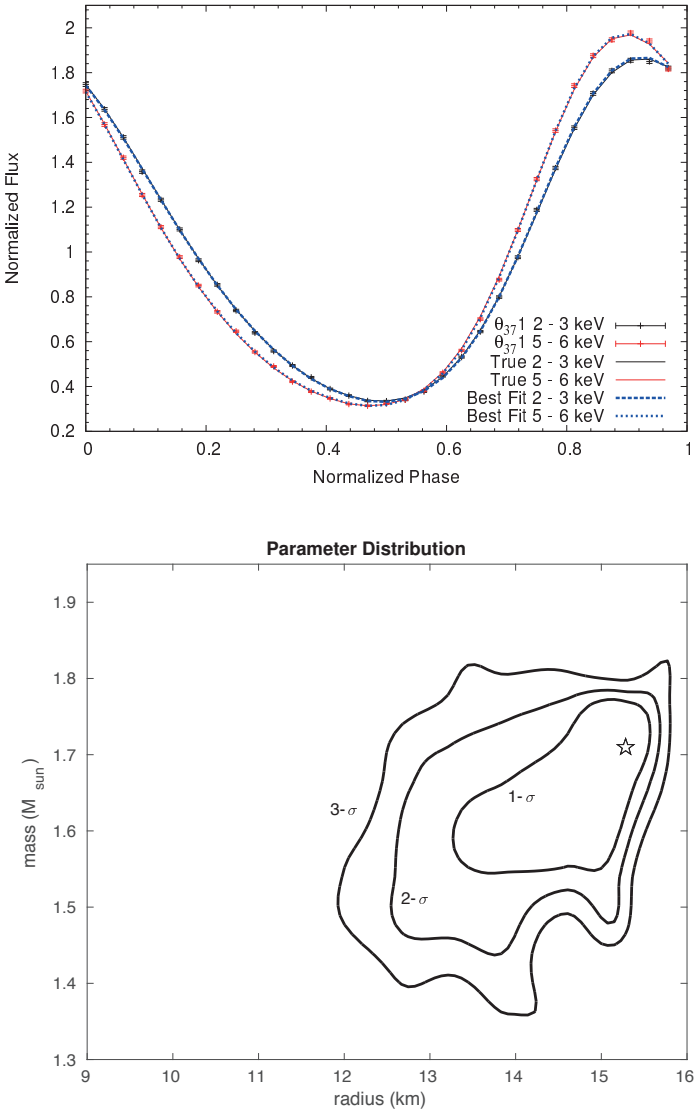

Figure 3.4: Top: Pulse profiles for model θ_{37} in two energy bands. Bottom: χ^2 contours in the $M - R$ plane for fitting synthetic data from model $\theta_{37}1$. Symbols and lines have the same meaning as in Figure 3.1. The best fit values for mass and radius for this Poisson realization are $1.71\,\mathrm{M}_\odot$ and 15.29 km, with $\chi^2 = 58.7$ for 59 dof. The approximate 1-σ_{degen} limits for mass and radius are $0.11\,\mathrm{M}_\odot$ and 1.1 km.

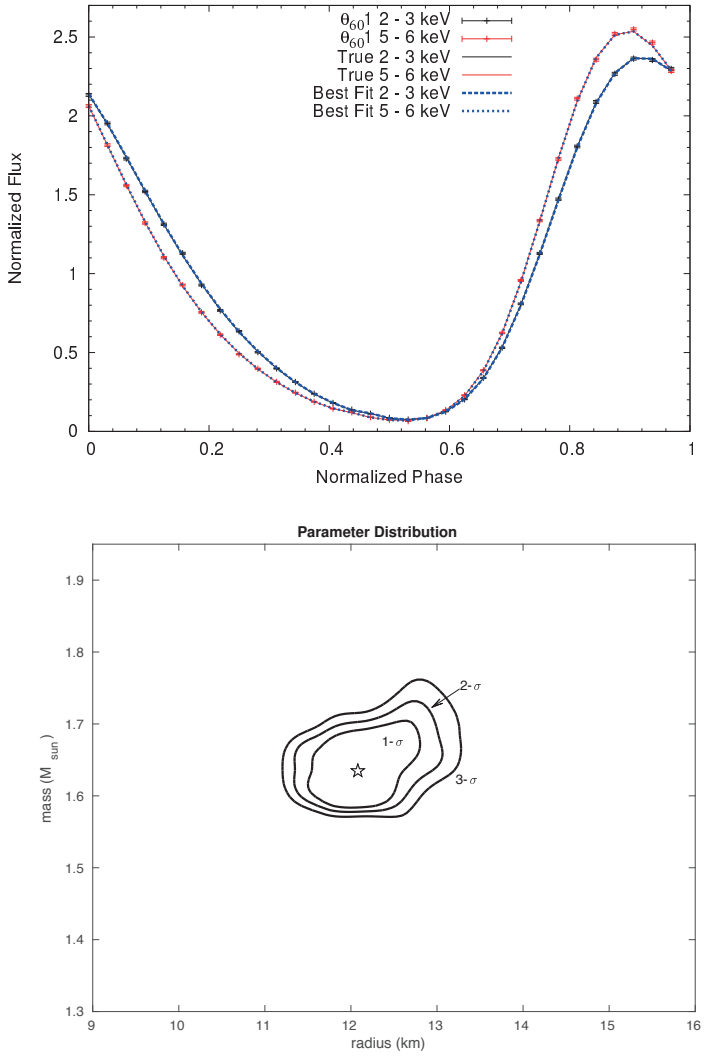

Figure 3.5: Top: Pulse profiles for model θ_{60} in two energy bands. Bottom: χ^2 contours in the $M - R$ plane for fitting synthetic data from model $\theta_{60}1$. The symbols have the same meaning as in Figure 3.1. The best fit values for mass and radius for this Poisson realization are $M = 1.63\,\mathrm{M}_\odot$ and $R = 12.1\,\mathrm{km}$, with $\chi^2 = 49.6$ for 59 dof. The approximate 1-σ_{degen} limits for mass and radius are $0.06\,\mathrm{M}_\odot$ and $0.7\,\mathrm{km}$.

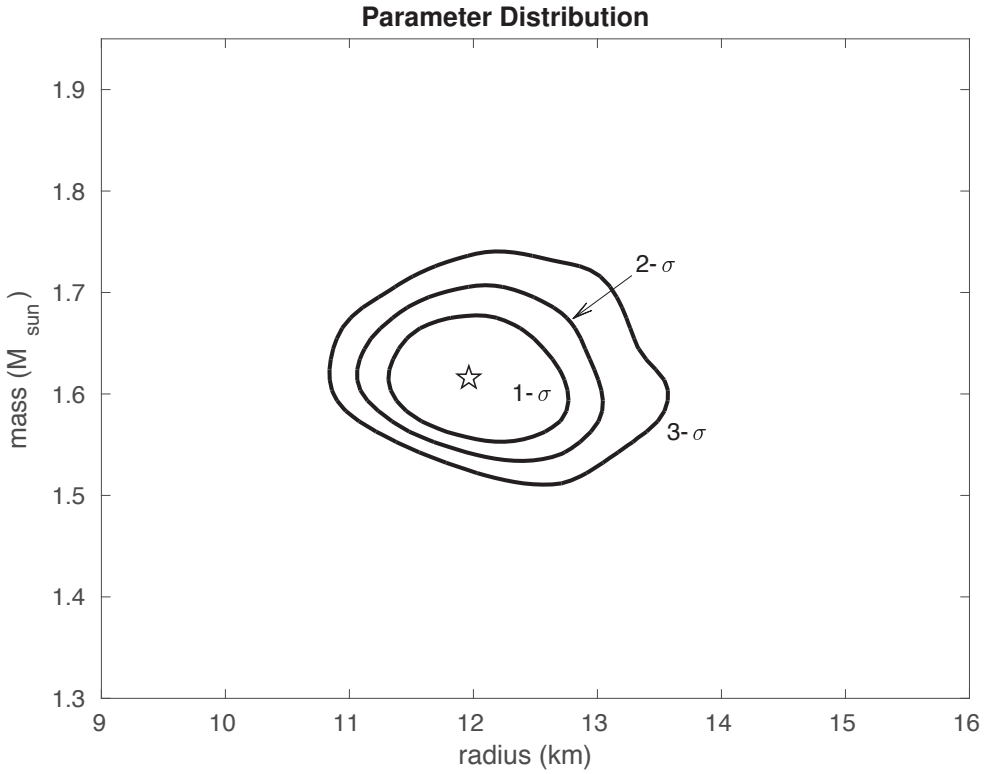

Figure 3.6: χ^2 contours in the $M - R$ plane for fitting synthetic data from model $\theta_{60}C_{6250}1$. The symbols have the same meaning as in Figure 3.1. The best fit values for mass and radius for this Poisson realization are $M = 1.62\,\mathrm{M}_\odot$ and $R = 12.0\,\mathrm{km}$, with $\chi^2 = 53.8$ for 59 dof. The approximate 1-σ_{degen} limits for mass and radius are $0.06\,\mathrm{M}_\odot$ and $0.8\,\mathrm{km}$.

Table 3.3: Summary of Best Fits for Model A

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	20.0	0.0	0.1969	0.296	0.470	0.373	5.741
1	58.2	1.575	11.97	48.7	25.0	7.1	0.1943	0.318	0.598	0.347	6.120
2	60.8	1.517	12.39	19.9	57.3	-1.1	0.1808	0.286	0.508	0.362	5.577
3	58.6	1.344	9.39	39.4	34.4	2.1	0.2114	0.358	0.638	0.357	5.565
4	71.7	1.532	10.70	21.7	59.8	0.0	0.2114	0.319	0.467	0.383	5.658
5	63.9	1.666	11.46	71.8	17.6	-6.2	0.2147	0.287	0.297	0.426	5.471
6	56.7	1.445	10.26	36.7	36.0	6.4	0.2080	0.351	0.649	0.350	5.932
7	64.6	1.469	10.81	37.2	34.4	7.4	0.2007	0.342	0.657	0.345	6.007
8	53.9	1.694	12.39	68.7	17.3	-4.7	0.2019	0.277	0.347	0.404	5.584
9	42.3	1.417	11.76	42.8	27.5	4.4	0.1779	0.313	0.651	0.338	5.773
10	67.0	1.510	10.32	21.9	60.6	-0.7	0.2161	0.324	0.456	0.388	5.583
11	57.1	1.581	11.29	21.7	58.1	2.6	0.2068	0.314	0.491	0.372	5.828
12	68.9	1.433	12.08	29.1	39.9	6.2	0.1752	0.311	0.671	0.331	5.869
13	58.4	1.458	10.88	35.6	35.8	7.5	0.1979	0.341	0.659	0.345	5.994
14	63.7	1.702	13.86	15.9	67.6	-5.7	0.1813	0.254	0.366	0.390	5.541
15	58.7	1.447	9.59	38.1	36.9	6.7	0.2228	0.371	0.629	0.360	6.006
16	45.5	1.746	13.31	18.3	63.1	1.5	0.1937	0.280	0.429	0.375	5.983
17	51.8	1.435	12.96	30.7	36.3	8.6	0.1635	0.302	0.693	0.323	6.009
18	53.1	1.488	10.57	38.4	34.2	7.6	0.2079	0.349	0.648	0.348	6.076
19	72.0	1.591	12.22	19.9	59.0	0.3	0.1923	0.292	0.485	0.367	5.727
20	58.1	1.438	9.68	38.5	36.1	6.0	0.2195	0.367	0.632	0.359	5.962
Ave	59.3	1.524	11.39	34.7	41.8	2.8	0.1989	0.318	0.549	0.363	5.813
Std	7.6	0.107	1.23	15.0	15.2	4.6	0.0159	0.032	0.120	0.025	0.203

Table 3.4: Summary of Tests on A1 Data

Model	χ^2/dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	20.0	0.0	0.1969	0.296	0.470	0.373	5.741
A1 Best Fit	58.2/59	1.575	11.97	48.7	25.0	7.1	0.1943	0.318	0.598	0.347	6.120
Wrong Atmosphere	316.6/59	1.210	11.29	40.2	39.2	10.4	0.1583	0.408	0.592	0.495	7.004
Spot Size Varies	59.0/58	1.508	11.57	33.3	37.0	9.0	0.1925	0.331	0.667	0.338	6.137
Temperature Varies	57.9/58	1.668	11.81	23.1	55.1	5.7	0.2086	0.322	0.526	0.364	6.266
Background	57.7/57	1.522	11.75	45.8	26.7	6.6	0.1913	0.321	0.623	0.344	6.045

Table 3.5: Summary of Best Fits for Model AO

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	20.0	0.0	0.1969	0.296	0.470	0.373	5.741
1	58.8	1.643	12.48	34.2	32.9	11.4	0.1944	0.306	0.694	0.302	6.140
2	58.7	1.590	11.79	63.0	19.2	-3.4	0.1992	0.293	0.429	0.385	5.600
3	58.0	1.721	11.57	20.9	52.1	-3.0	0.2197	0.282	0.573	0.292	5.479
4	70.8	1.584	11.00	42.9	28.0	4.0	0.2127	0.319	0.647	0.314	5.830
5	64.0	1.756	11.29	21.5	52.2	-0.4	0.2297	0.290	0.570	0.292	5.604
6	51.9	1.588	9.52	57.4	23.4	-3.0	0.2463	0.335	0.494	0.342	5.629
7	65.3	1.624	11.36	60.2	20.5	-1.0	0.2111	0.304	0.466	0.365	5.708
8	55.7	1.631	11.29	44.0	27.3	6.8	0.2133	0.319	0.639	0.315	5.980
9	42.5	1.637	13.57	75.4	15.6	-9.7	0.1781	0.260	0.243	0.499	5.521
10	58.5	1.747	12.24	23.5	46.5	5.0	0.2108	0.289	0.632	0.290	5.842
11	62.0	1.751	13.37	21.0	48.9	5.1	0.1934	0.271	0.613	0.291	5.809
12	69.0	1.628	12.13	66.0	18.4	-3.7	0.1982	0.288	0.385	0.403	5.649
13	55.6	1.727	11.95	74.1	17.2	-4.6	0.2134	0.284	0.262	0.448	5.638
14	58.4	1.557	11.70	36.4	31.3	5.9	0.1965	0.308	0.688	0.305	5.822
15	65.3	1.680	11.06	69.2	18.7	-4.7	0.2243	0.300	0.337	0.403	5.612
16	45.1	1.674	12.26	52.7	22.9	8.7	0.2016	0.309	0.559	0.345	6.173
17	56.5	1.686	12.29	72.0	17.2	-4.8	0.2026	0.281	0.295	0.442	5.635
18	56.0	1.667	11.93	62.3	19.8	0.6	0.2064	0.299	0.438	0.379	5.858
19	67.6	1.561	11.65	47.5	24.8	3.4	0.1979	0.309	0.613	0.328	5.820
20	61.6	1.604	9.99	42.3	30.0	2.4	0.2371	0.337	0.640	0.309	5.837
Ave	59.1	1.653	11.72	49.3	28.3	0.8	0.2093	0.299	0.511	0.353	5.759
Std	7.0	0.062	0.93	18.2	11.9	5.3	0.0158	0.019	0.142	0.060	0.183

Table 3.6: Summary of Best Fits for Model θ_{37}

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	37.0	0.0	0.1969	0.521	0.399	0.720	10.101
1	58.7	1.710	15.29	26.4	69.2	-7.6	0.1652	0.416	0.318	0.737	9.765
2	62.1	1.577	11.47	40.6	56.8	1.9	0.2030	0.544	0.416	0.718	10.185
3	52.3	1.594	11.88	43.5	52.1	3.5	0.1981	0.542	0.446	0.700	10.430
4	68.8	1.614	12.43	34.4	62.5	-1.3	0.1918	0.501	0.381	0.724	9.976
5	55.5	1.640	11.78	62.4	36.8	0.7	0.2056	0.530	0.371	0.736	10.237
6	44.1	1.684	12.68	68.1	31.5	-3.5	0.1961	0.485	0.318	0.758	9.925
7	62.3	1.611	11.87	36.6	61.5	0.1	0.2004	0.524	0.383	0.730	10.103
8	51.2	1.651	12.22	35.8	62.1	0.5	0.1995	0.518	0.379	0.728	10.259
9	70.9	1.516	11.91	40.8	53.4	0.0	0.1880	0.524	0.452	0.696	9.935
10	65.4	1.555	11.89	58.0	38.0	-0.8	0.1931	0.521	0.418	0.712	9.953
11	50.3	1.531	12.34	54.2	39.2	0.4	0.1832	0.513	0.454	0.690	9.994
12	70.3	1.567	11.64	56.8	39.8	0.6	0.1988	0.536	0.421	0.714	10.102
13	66.2	1.743	13.42	69.9	29.7	-4.6	0.1918	0.465	0.299	0.763	9.995
14	63.6	1.711	13.02	67.0	31.7	-2.3	0.1941	0.483	0.333	0.743	10.111
15	85.4	1.836	13.52	74.9	28.2	-6.0	0.2005	0.456	0.230	0.808	10.026
16	63.9	1.540	12.21	47.7	45.4	2.1	0.1863	0.527	0.472	0.685	10.209
17	60.4	1.526	12.01	50.5	43.1	1.6	0.1876	0.528	0.464	0.690	10.083
18	48.3	1.586	11.67	46.5	49.4	3.9	0.2007	0.551	0.448	0.703	10.448
19	49.5	1.575	11.69	37.2	60.6	-1.0	0.1990	0.526	0.391	0.729	9.973
20	73.3	1.695	12.81	31.8	67.4	-2.3	0.1954	0.487	0.327	0.752	10.043
Ave	61.1	1.623	12.39	49.1	47.9	-0.7	0.1939	0.509	0.386	0.726	10.088
Std	9.9	0.082	0.88	13.8	13.0	2.9	0.0088	0.033	0.064	0.029	0.165

Table 3.7: Summary of Best Fits for Model θ_{37}O

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	37.0	0.0	0.1969	0.521	0.399	0.720	10.101
1	59.5	1.611	11.86	46.0	48.3	4.1	0.2006	0.537	0.462	0.674	10.351
2	58.7	1.534	11.65	53.7	41.8	1.4	0.1945	0.537	0.441	0.707	10.068
3	59.7	1.614	11.99	37.4	56.7	-0.2	0.1988	0.507	0.436	0.663	9.858
4	74.1	1.546	11.50	56.3	40.6	1.0	0.1985	0.541	0.422	0.721	10.081
5	65.3	1.617	11.65	40.2	54.5	1.5	0.2050	0.526	0.443	0.665	10.032
6	51.4	1.560	11.29	61.6	37.9	-0.8	0.2041	0.540	0.375	0.750	9.968
7	63.5	1.561	11.71	62.1	37.0	-0.2	0.1969	0.532	0.374	0.762	10.064
8	55.8	1.611	11.74	44.2	50.2	3.4	0.2026	0.536	0.459	0.670	10.254
9	42.1	1.520	11.89	58.1	38.6	0.4	0.1888	0.530	0.413	0.740	10.039
10	56.0	1.576	13.48	33.6	58.5	0.3	0.1727	0.472	0.435	0.676	9.893
11	59.5	1.608	11.85	45.8	48.5	3.9	0.2004	0.537	0.462	0.675	10.343
12	57.7	1.562	11.68	59.3	38.7	1.0	0.1975	0.537	0.398	0.741	10.147
13	59.7	1.613	11.98	37.4	56.7	-0.2	0.1988	0.508	0.436	0.663	9.859
14	74.1	1.544	11.50	56.3	40.6	0.9	0.1983	0.541	0.422	0.721	10.072
15	65.3	1.618	11.64	40.3	54.5	1.5	0.2053	0.526	0.443	0.665	10.034
16	51.4	1.564	11.22	62.0	37.9	-0.9	0.2059	0.542	0.371	0.753	9.979
17	63.4	1.566	11.64	61.4	37.6	0.2	0.1987	0.535	0.380	0.754	10.094
18	55.7	1.617	11.74	43.5	51.0	3.4	0.2034	0.535	0.457	0.668	10.246
19	42.1	1.521	11.88	58.0	38.7	0.5	0.1891	0.530	0.414	0.738	10.042
20	63.6	1.652	11.89	38.4	56.6	1.6	0.2052	0.518	0.432	0.665	10.089
Ave	58.9	1.581	11.79	49.8	46.3	1.1	0.1982	0.528	0.424	0.704	10.076
Std	8.1	0.037	0.44	9.7	7.8	1.5	0.0075	0.016	0.029	0.037	0.136

Table 3.8: Summary of Best Fits for Model θ_{60}

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp $\times 10^{-2}$	β
True		1.600	12.00	60.0	60.0	0.000	0.1969	0.750	0.250	1.305	14.536
1	49.6	1.635	12.08	67.7	53.5	-0.5	0.1999	0.744	0.225	1.332	14.584
2	62.0	1.596	12.10	59.8	59.6	-0.1	0.1948	0.745	0.255	1.299	14.516
3	56.2	1.621	12.21	68.0	52.1	-2.0	0.1961	0.732	0.230	1.323	14.411
4	72.0	1.599	12.05	63.3	56.5	-0.4	0.1960	0.745	0.248	1.307	14.479
5	67.5	1.598	11.89	63.8	56.8	-0.3	0.1985	0.751	0.242	1.315	14.453
6	60.7	1.593	12.01	59.6	60.0	-0.2	0.1959	0.747	0.253	1.300	14.472
7	65.9	1.598	12.10	59.7	59.7	-0.1	0.1950	0.746	0.254	1.299	14.527
8	53.9	1.626	12.12	66.3	54.3	-0.3	0.1981	0.743	0.235	1.320	14.579
9	43.0	1.587	12.15	59.6	59.3	-0.4	0.1929	0.741	0.259	1.294	14.449
10	58.8	1.627	12.43	51.1	68.3	-1.4	0.1933	0.723	0.232	1.320	14.424
11	63.0	1.617	12.12	61.0	58.9	0.3	0.1970	0.749	0.250	1.303	14.673
12	45.1	1.589	12.34	63.8	54.4	-1.0	0.1902	0.729	0.257	1.293	14.374
13	50.1	1.613	11.97	56.7	64.0	0.1	0.1990	0.751	0.241	1.314	14.567
14	61.7	1.597	12.04	57.3	62.5	-0.5	0.1959	0.746	0.249	1.306	14.490
15	60.2	1.622	12.08	65.7	55.0	-0.4	0.1983	0.746	0.236	1.322	14.594
16	66.3	1.607	12.06	60.3	59.7	0.2	0.1968	0.750	0.250	1.304	14.598
17	43.7	1.588	11.99	59.6	60.0	-0.3	0.1956	0.747	0.253	1.302	14.441
18	49.9	1.606	12.23	59.3	59.8	0.2	0.1939	0.743	0.257	1.295	14.606
19	55.6	1.612	11.84	63.1	58.3	0.3	0.2011	0.759	0.238	1.322	14.610
20	69.1	1.626	12.00	64.0	57.0	0.4	0.2001	0.754	0.239	1.317	14.689
Ave	57.7	1.608	12.09	61.5	58.5	-0.3	0.1964	0.745	0.245	1.309	14.527
Std	8.4	0.014	0.13	4.0	3.7	0.6	0.0026	0.008	0.010	0.011	0.087

Table 3.9: Summary of Best Fits for Model $\theta_{60}C_{6250}$

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	60.0	0.0	0.1969	0.750	0.250	1.305	14.536
1	69.3	1.634	12.22	60.8	59.3	0.7	0.1975	0.751	0.249	1.305	14.831
2	56.7	1.625	12.83	49.1	68.7	-2.4	0.1870	0.704	0.238	1.310	14.351
3	53.8	1.616	11.97	60.7	60.1	0.8	0.1994	0.756	0.244	1.314	14.678
4	67.0	1.615	12.54	49.9	68.6	-2.4	0.1902	0.712	0.235	1.314	14.269
5	58.0	1.584	12.00	60.5	59.0	-0.5	0.1949	0.746	0.253	1.302	14.415
6	45.2	1.585	12.87	52.3	63.6	-1.8	0.1819	0.709	0.272	1.271	14.378
7	59.1	1.577	13.10	50.5	64.2	-2.0	0.1778	0.695	0.276	1.260	14.270
8	58.3	1.621	12.01	59.1	61.5	0.9	0.1993	0.754	0.245	1.308	14.684
9	64.1	1.615	12.42	56.8	61.6	0.6	0.1920	0.736	0.261	1.286	14.647
10	59.0	1.598	12.49	63.4	54.3	-0.6	0.1889	0.726	0.261	1.286	14.465
11	68.3	1.605	12.86	65.1	51.4	-1.7	0.1843	0.709	0.263	1.278	14.430
12	60.6	1.672	12.64	47.8	72.6	-2.1	0.1953	0.707	0.201	1.356	14.395
13	52.2	1.571	12.50	58.8	57.8	-1.1	0.1856	0.724	0.276	1.268	14.354
14	63.5	1.632	12.42	60.1	58.9	1.2	0.1940	0.742	0.258	1.292	14.819
15	51.0	1.602	12.16	59.9	59.3	01.	0.1946	0.744	0.256	1.296	14.558
16	48.5	1.662	13.05	70.4	47.6	-2.1	0.1881	0.696	0.226	1.319	14.460
17	61.3	1.592	12.30	55.1	63.2	0.0	0.1911	0.733	0.258	1.292	14.415
18	49.8	1.629	11.14	63.8	62.9	2.1	0.2159	0.798	0.201	1.373	14.839
19	45.6	1.697	13.36	73.4	45.0	-4.0	0.1876	0.678	0.202	1.350	14.410
20	72.0	1.620	12.52	54.9	63.3	0.2	0.1911	0.731	0.258	1.288	14.643
Ave	58.2	1.618	12.47	58.6	60.2	-0.7	0.1918	0.728	0.247	1.303	14.516
Std	7.7	0.031	0.48	6.6	6.6	1.5	0.0078	0.027	0.023	0.028	0.179

Table 3.10: Summary of Best Fits for Model $(M/R)_{hi}$

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.680	11.00	60.0	20.8	0.0	0.2255	0.308	0.467	0.350	5.742
1	59.5	1.670	11.16	45.6	27.2	7.5	0.2210	0.327	0.622	0.321	6.141
2	58.8	1.562	9.98	55.1	23.8	0.1	0.2312	0.331	0.524	0.347	5.657
3	58.4	1.680	11.28	17.0	70.2	-9.9	0.2199	0.275	0.323	0.384	5.215
4	72.5	1.569	10.34	29.1	44.0	5.2	0.2241	0.338	0.628	0.327	5.918
5	70.9	1.570	9.71	23.5	57.5	0.7	0.2388	0.336	0.492	0.354	5.681
6	56.6	1.484	8.95	42.3	33.5	3.8	0.2448	0.371	0.617	0.338	5.844
7	64.1	1.610	10.36	24.6	52.7	3.2	0.2295	0.331	0.551	0.339	5.855
8	56.2	1.726	11.07	63.8	19.8	-0.4	0.2303	0.303	0.416	0.360	5.748
9	41.3	1.543	10.91	36.6	33.0	5.0	0.2089	0.325	0.673	0.315	5.848
10	63.7	1.669	10.75	21.8	58.5	1.8	0.2293	0.316	0.485	0.348	5.813
11	62.4	1.949	13.56	14.4	76.1	-10.4	0.2123	0.241	0.232	0.402	5.427
12	70.5	1.520	10.11	36.5	35.0	4.1	0.2220	0.341	0.659	0.322	5.814
13	56.7	1.510	8.99	46.3	30.9	5.1	0.2482	0.372	0.592	0.343	5.922
14	61.9	1.525	10.31	41.0	30.7	4.2	0.2184	0.335	0.649	0.323	5.784
15	64.7	1.578	9.35	22.7	62.4	-0.4	0.2492	0.342	0.428	0.370	5.673
16	45.5	1.733	11.33	48.7	25.8	8.9	0.2259	0.327	0.594	0.325	6.287
17	53.3	1.547	10.54	34.9	35.9	5.8	0.2168	0.335	0.665	0.320	5.895
18	57.0	1.570	9.94	45.1	29.4	5.0	0.2333	0.348	0.615	0.331	5.958
19	69.5	1.629	11.12	23.1	52.7	2.0	0.2163	0.312	0.558	0.332	5.785
20	62.0	1.524	9.19	25.7	54.7	1.0	0.2450	0.354	0.520	0.354	5.724
Ave	60.3	1.608	10.45	34.9	42.7	2.1	0.2282	0.328	0.542	0.343	5.799
Std	7.7	0.105	1.03	13.1	16.2	4.8	0.0117	0.029	0.116	0.022	0.220

Table 3.11: Summary of Best Fits for Model $(M/R)_{lo}$

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.350	13.00	60.0	19.8	0.0	0.1534	0.293	0.471	0.423	5.746
1	72.0	1.299	11.73	53.6	23.8	4.1	0.1635	0.325	0.543	0.413	5.837
2	79.5	1.393	13.44	18.2	63.5	-1.7	0.1531	0.280	0.423	0.435	5.685
3	55.7	1.357	11.28	55.8	24.0	4.3	0.1777	0.336	0.514	0.426	5.945
4	45.9	1.335	14.86	16.6	63.8	-4.0	0.1327	0.256	0.423	0.424	5.578
5	60.9	1.423	14.18	17.5	64.2	-1.9	0.1482	0.271	0.415	0.434	5.765
6	49.1	1.191	12.43	32.8	37.5	8.9	0.1415	0.330	0.667	0.382	6.092
7	73.8	1.275	12.72	56.2	21.2	1.1	0.1480	0.300	0.519	0.412	5.725
8	67.7	1.225	12.85	21.5	54.4	1.1	0.1408	0.297	0.542	0.403	5.672
9	51.4	1.127	12.90	52.5	21.6	0.4	0.1290	0.292	0.566	0.395	5.499
10	71.8	1.236	11.96	52.0	23.8	2.5	0.1526	0.318	0.563	0.406	5.739
11	61.0	1.235	15.25	18.6	56.5	1.6	0.1196	0.266	0.523	0.392	5.854
12	67.3	1.115	13.10	27.3	42.4	6.9	0.1257	0.309	0.656	0.375	5.890
13	81.0	1.307	14.36	19.4	57.3	1.9	0.1344	0.279	0.510	0.402	5.890
14	65.0	1.158	10.84	35.0	37.5	5.7	0.1578	0.350	0.649	0.397	5.760
15	61.6	1.272	13.07	58.3	19.8	-0.1	0.1437	0.288	0.495	0.413	5.606
16	58.2	1.527	11.45	67.0	20.7	2.1	0.1969	0.326	0.366	0.471	6.020
17	47.4	1.311	14.88	17.0	62.7	-4.1	0.1301	0.259	0.439	0.421	5.638
18	69.6	1.200	13.22	24.9	47.0	7.2	0.1340	0.309	0.618	0.385	5.996
19	59.3	1.417	14.13	17.5	64.3	-2.7	0.1481	0.271	0.413	0.435	5.744
20	51.7	1.183	12.26	24.6	49.2	3.7	0.1425	0.315	0.594	0.397	5.741
Ave	62.5	1.279	13.05	34.3	42.8	1.9	0.1460	0.299	0.522	0.411	5.784
Std	10.1	0.105	1.23	17.2	17.1	3.6	0.0178	0.027	0.086	0.022	0.152

Table 3.12: Summary of Best Fits for Model $\beta_{hi}(M/R)_{hi}$

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.680	11.00	60.0	64.0	0.0	0.2255	0.778	0.219	1.235	14.533
1	65.6	1.758	11.85	74.5	48.5	-2.8	0.2191	0.722	0.177	1.274	14.351
2	68.2	1.690	11.13	65.6	58.0	-0.1	0.2242	0.772	0.219	1.234	14.553
3	44.0	1.736	11.41	51.9	72.6	-1.1	0.2247	0.751	0.185	1.267	14.515
4	57.4	1.676	11.32	58.7	62.8	-0.9	0.2186	0.760	0.237	1.215	14.428
5	51.9	1.693	11.17	68.2	55.4	-1.1	0.2238	0.765	0.210	1.242	14.451
6	64.3	1.686	11.21	67.7	55.4	-1.2	0.2221	0.761	0.216	1.237	14.393
7	67.5	1.680	10.96	64.1	60.2	0.0	0.2264	0.781	0.217	1.239	14.548
8	54.0	1.696	11.22	54.3	69.4	-0.9	0.2232	0.760	0.206	1.248	14.407
9	49.4	1.687	11.10	56.1	67.8	-0.4	0.2244	0.768	0.211	1.242	14.443
10	41.0	1.701	11.22	55.3	68.3	-0.8	0.2239	0.764	0.211	1.240	14.505
11	33.6	1.692	11.01	67.4	57.2	0.0	0.2269	0.777	0.208	1.246	14.548
12	44.0	1.693	11.04	64.2	60.1	0.4	0.2265	0.781	0.217	1.237	14.650
13	77.8	1.655	11.27	62.5	58.6	-1.7	0.2169	0.757	0.241	1.212	14.247
14	54.0	1.696	11.22	54.3	69.4	-0.9	0.2232	0.760	0.206	1.248	14.407
15	47.3	1.749	11.90	72.7	49.4	-2.5	0.2170	0.724	0.194	1.254	14.410
16	54.3	1.700	11.06	55.6	69.2	-0.3	0.2270	0.772	0.200	1.253	14.522
17	54.2	1.693	10.88	58.2	67.5	0.5	0.2298	0.785	0.202	1.252	14.612
18	44.0	1.714	11.32	69.3	54.3	-0.6	0.2236	0.759	0.206	1.243	14.540
19	60.3	1.672	11.38	58.9	62.1	-0.9	0.2170	0.757	0.242	1.211	14.397
20	43.6	1.681	11.14	58.0	65.1	-0.6	0.2228	0.769	0.223	1.230	14.470
Ave	53.8	1.697	11.24	61.9	61.6	-0.8	0.2231	0.762	0.211	1.241	14.470
Std	10.7	0.024	0.25	6.5	6.8	0.8	0.0036	0.016	0.016	0.016	0.093

Table 3.13: Summary of Best Fits for Model ν_{400}

Model	χ^2 59 dof	M $M_\odot$	R km	i deg.	θ deg.	ϕ $\times 10^{-3}$	M/R	$\sin i \sin\theta$	$\cos i \cos\theta$	Amp	β $\times 10^{-2}$
True		1.600	12.00	60.0	20.0	0.0	0.1969	0.296	0.470	0.373	3.827
1	45.9	1.082	15.98	30.4	30.5	5.3	0.1000	0.257	0.743	0.296	3.851
2	67.0	1.597	11.03	19.6	65.3	-1.2	0.2138	0.304	0.394	0.396	3.715
3	46.0	1.527	11.63	32.5	38.3	6.8	0.1939	0.332	0.662	0.339	4.142
4	75.8	1.740	14.38	66.5	16.0	-2.8	0.1787	0.253	0.384	0.382	3.805
5	54.8	1.475	9.59	42.5	33.6	5.2	0.2271	0.374	0.614	0.363	4.070
6	45.4	1.898	16.00	76.9	12.6	-10.5	0.1752	0.212	0.222	0.432	3.530
7	59.5	1.905	14.61	15.3	71.8	-2.6	0.1926	0.251	0.301	0.409	3.921
8	46.5	1.581	12.68	37.6	31.8	8.3	0.1841	0.322	0.673	0.334	4.304
9	67.1	1.441	11.71	34.7	34.6	4.8	0.1817	0.323	0.677	0.336	3.977
10	61.3	1.681	16.00	64.2	15.1	-4.3	0.1552	0.234	0.420	0.362	3.778
11	50.8	1.242	11.45	28.7	38.8	0.0	0.1602	0.301	0.683	0.327	3.504
12	50.1	1.754	14.74	14.8	69.7	-6.3	0.1757	0.240	0.336	0.395	3.680
13	57.9	1.343	16.00	35.1	27.8	8.5	0.1240	0.268	0.723	0.302	4.152
14	67.4	1.518	13.79	19.7	53.7	0.9	0.1626	0.272	0.557	0.341	3.827
15	78.2	1.715	16.00	59.0	17.0	1.7	0.1583	0.251	0.493	0.346	4.069
16	64.3	1.312	8.27	38.0	39.4	1.1	0.2342	0.390	0.609	0.372	3.714
17	58.4	1.484	10.98	22.1	56.1	0.4	0.1996	0.312	0.517	0.367	3.701
18	48.5	1.495	8.94	26.8	57.7	4.1	0.2470	0.382	0.477	0.395	4.020
19	50.1	1.754	14.74	14.8	69.7	-6.3	0.1757	0.240	0.336	0.395	3.680
20	67.3	1.497	12.42	34.1	34.5	7.3	0.1780	0.317	0.683	0.331	4.116
Ave	58.1	1.552	13.05	35.7	40.7	1.0	0.1809	0.292	0.525	0.361	3.878
Std	9.9	0.207	2.46	17.6	18.8	5.2	0.0338	0.051	0.155	0.035	0.216

4

Phase-Resolved Spectroscopy of the Lower kHz QPO in 4U 1608–52

A. L. Stevens, P. Uttley, and D. Altamirano

To be submitted to Monthly Notices of the Royal Astronomical Society

Abstract

Kilohertz quasi-periodic oscillations (kHz QPOs) are the most rapid (quasi-)coherent kind of variability that have been detected in the light curves of accreting neutron star X-ray binaries. They often appear in pairs, with the lower kHz QPO the more coherent signal. Whereas the upper kHz QPO may arise from Keplerian motion in the inner accretion disc, the origin of the lower kHz QPO remains unclear. One way to investigate the lower kHz QPO is with spectral-timing, in which the energy and time variability of the photons are simultaneously analyzed. Previous work using the rms spectrum revealed that the variable emission is a Comptonized blackbody, consistent with that expected from the boundary layer between the accretion flow and neutron star surface. Furthermore, the lag-energy spectra and covariance spectra indicate that the lower kHz QPO arises from a more complex energy-dependent variability than just an overall modulation in the flux normalization. To better interpret the spectral variability, we present phase-resolved spectroscopy of a kHz QPO for the first time, using a method based on the energy-dependent cross-correlation function (CCF). The

best-fitting spectral parameterisation requires the power-law index, high-energy cut-off temperature and seed blackbody temperature of the Comptonized emission to vary with QPO phase. Additionally, the variations in these three parameters show small but non-zero phase differences, which together can explain the previously observed lag-energy dependence. We suggest that these spectral variations could be explained by a "breathing" oscillation in the Comptonizing boundary layer; this possibility can be explored in greater detail with current and future X-ray missions such as *AstroSat*, *NICER*, *eXTP*, and *STROBE-X*.

4.1 Introduction

Low-mass X-ray binaries (LMXBs) show a range of variability types and time-scales in their X-ray emission, with quasi-periodic oscillations (QPOs; van der Klis et al. 1985) being particularly enigmatic. LMXBs containing neutron stars are known to exhibit QPOs in the frequency range $300 - 1200$ Hz, known as kilohertz (kHz) QPOs (van der Klis et al. 1996b; Strohmayer et al. 1996b). The frequencies of kHz QPOs are consistent with the dynamical timescale of Keplerian orbital motion at the inner edge of the accretion disc, where the accretion flow interacts with the surface of the neutron star at the so-called boundary layer (e.g., Gilfanov et al. 2003). Dissecting the kHz QPO signal could tell us about the nature of accretion in the strong gravitational field closest to the neutron star, and potentially constrain the physical properties of the neutron star itself.

Often 'twin' kHz QPOs are observed simultaneously (with some weakly varying separation in peak frequency, see, e.g., van der Klis 2006b for a review), leading to a distinction between so-called upper and lower kHz QPOs. These types can also be seen individually and identified according to characteristic properties, such as their coherence. The upper kHz QPOs are relatively incoherent, with quality factors (QPO centroid frequency divided by the Full Width at Half Maximum, corresponding roughly to the number of cycles over which the oscillation remains coherent) $Q \leq 20$ while lower kHz QPOs are more strongly coherent (see, e.g., Berger et al. 1996; Méndez et al. 1998b; Méndez & van der Klis 1999; Méndez et al. 2001; Di Salvo et al. 2001, 2003; Méndez 2006; Barret et al. 2005, 2006), showing Q as high as 250, once random walk variability of the QPO frequency on minutes time-scales is accounted for (see Barret & Vaughan 2012, and references therein).

The first theoretical models for the origin of kHz QPOs assumed orbital motion of bright blobs of accreting material at the inner edge of the accretion disc (e.g., Stella & Vietri 1999; Lamb & Miller 2001). However, it is difficult to explain the very high coherence of the lower kHz QPO with orbital motions of blobs, which would be difficult to maintain in a turbulent accretion flow over the hundreds of orbits required (e.g., Barret et al. 2005, and references therein). Models that do not invoke orbital motion involve quasi-periodic perturbations in the soft photon flux that are Compton up-scattered (e.g., Lee et al. 2001; Kumar & Misra 2014, 2016 and for a more detailed review of kHz QPO models, see van der Klis 2006a,b; Kumar & Misra 2014 and references therein). However, while potentially explaining some of the characteristics of the QPO emission, such models cannot naturally explain the origin of the oscillation signal itself.

Comparison of the spectral and timing evolution of neutron star LMXBs shows that kHz QPOs arise in intermediate states where both disc and Comptonized emission components are strong (e.g., van Straaten et al. 2003; van der Klis 2006b; Altamirano et al. 2008; Zhang et al. 2017, and references therein). Further insight into the

origin of the kHz QPOs can be gained from studying their spectral-timing properties which reveal the timing properties as a detailed function of energy (see Uttley et al. 2014 for a review of spectral-timing methods). For example, Gilfanov et al. (2003) measured the Fourier-frequency resolved rms spectrum of both upper and lower kHz QPOs to show that it is the Comptonized emission, associated with the boundary layer, which is varying at the QPO frequency, and not the disc. More recently, Peille et al. (2015) and Troyer & Cackett (2017) used the more sensitive 'covariance spectrum' approach to demonstrate that the shape of the variable spectrum appears to be harder than that of the time-averaged spectrum, perhaps suggesting an oscillation arising in a specific region of the boundary layer.

Besides studying the spectral shape of the QPO emission, causal information can be obtained using Fourier time-lags, which reveal that for the lower kHz QPO, soft photons lag hard photons (Vaughan et al. 1997, 1998; Kaaret et al. 1999; Barret 2013; de Avellar et al. 2013; Peille et al. 2015; de Avellar et al. 2016; Troyer & Cackett 2017) while there are either the opposite hard lags or no clear lag-energy dependence for the upper kHz QPO (de Avellar et al. 2013; Peille et al. 2015; de Avellar et al. 2016). The soft lags have been variously explained as due to X-ray reverberation of Comptonized emission from the blackbody emitting disc (Barret 2013) or as due to Compton scattering including a substantial back-heated component of seed photon emission (Lee et al. 2001; Kumar & Misra 2014, 2016). Detailed modelling of the expected reflection spectrum appears to rule out a reverberation origin (Cackett 2016), but that explanation may still hold for the hard lags seen for the upper kHz QPO (Peille et al. 2015).

The Fourier-frequency domain spectral-timing methods described above have provided some useful insights into the kHz QPO phenomenon, but their results can be difficult to interpret in the absence of any specific model predictions. Furthermore, these methods treat the variability amplitude and lag information separately. For quasi-periodic signals, *phase-resolved* reconstruction of the spectrum provides a more intuitive approach to studying the spectral variability, since it combines both amplitude and causal information in a manner that is consistent with the (quasi-)coherent nature of the signal. Recently, techniques have been developed to reconstruct the spectral evolution of the QPO signal as a function of its phase, either via reconstruction of the waveform through measuring the QPO phase within a coherence time (Ingram & van der Klis 2015; Ingram et al. 2016), or by using the cross-correlation function (CCF) to 'phase-lock' the QPO variations in narrow energy channels to a broader reference band (Stevens & Uttley 2016).

The latter CCF approach was first used at lower-frequencies to study the Type B QPO in the black hole LMXB GX 339-4, but it is particularly suited to the study of high-frequency QPOs where coherence times are too short to accurately recover the QPO phase via the method of Ingram & van der Klis (2015). The relative coherence and strength of the lower kHz QPOs in neutron star LMXBs makes them particu-

larly suitable targets for such a phase-resolved study, which could shed light on this enigmatic variability signal. Therefore, as a test case for the study of kHz QPOs, in this paper we adapt the CCF method of Stevens & Uttley (2016) to determine, for the very first time, the spectral evolution of a neutron star LMXB on the lower kHz QPO time-scale, as a function of the QPO phase.

In Section 4.2, we explain the data selection and the timing and spectral properties of the data. We then outline our analysis technique for phase-resolved spectroscopy of the lower kHz QPO in Section 4.3 and present our results in Section 4.4. The interpretation and implications of these results for understanding the kHz QPO are discussed in Section 4.5, and our conclusions are found in Section 4.6.

4.2 Timing and Spectral Properties of the Data

The target of our study is 4U 1608–52, a neutron star LMXB transient that regularly goes into outburst (e.g., Christian & Swank 1997), is X-ray bright, and often shows kHz QPOs. It is classified as an atoll source due to the shape of the tracks it follows on the colour-colour diagram (Hasinger & van der Klis 1989). The distance to 4U 1608–52 is estimated to be 3.2 ± 0.3 to 4.1 ± 0.4 kpc (via X-ray burst photospheric radius expansion measurements; Murakami et al. 1980; Nakamura et al. 1989; Galloway et al. 2008) to $5.8^{+2.0}_{-1.9}$ kpc (via interstellar extinction modelling; Güver et al. 2010).

We use data from the March 3 1996 *Rossi X-ray Timing Explorer* (*RXTE*; Bradt et al. 1993) observation of 4U 1608–52 (obsID 10072-05-01-00), obtained with the Proportional Counter Array (PCA; Jahoda et al. 1996). This observation is particularly useful for the study of the lower kHz QPO, since the source was relatively bright during this time and the QPO is relatively stable in frequency, hence it has already been used as a test case for the study of spectral-timing and we can compare our results with the previous works (e.g., Berger et al. 1996; Cackett 2016). The observation consisted of three *RXTE* orbits, during which all 5 Proportional Counter Units (PCUs) were switched on and the PCA collected data in 64-channel event-mode with 122 μs time resolution (`E_125us_64M_0_1s`). No thermonuclear bursts are seen in the data, and only the lower kHz QPO, is detected during this time. We obtained Good Time Intervals (GTIs) corresponding to times when all of the following criteria are satisfied: five PCUs are on, elevation angle $> 10°$, target offset $< 0.02°$, and time since the South Atlantic Anomaly passage > 30 minutes. The background was minimal compared to the source and was not fully subtracted.

To compute the dynamical power spectrum and cross spectrum, the data were split into 32 s segments, which gave 195 segments in total (6.24 ks). We generated a 3–20 keV (event-mode energy channels 6–35 inclusive) dynamical power spectrum of the 5 PCUs combined (Figure 4.1) using our own code.[1] The sudden drop in

[1] See Appendix 4A for URL.

QPO frequency at around segment 30 is due to the gap between the first two orbits that is not shown. The dynamical power spectrum was re-binned to 1 Hz frequency resolution to have adequate signal-to-noise, before being fitted to obtain the QPO centroid frequency per segment. We then used the fitted QPO centroid frequency to determine how much to shift each segment Fourier transform by to line up the QPOs at 835 Hz (the unshifted average centroid frequency; see Méndez et al. 1998a for the shift-and-add technique). The shifting was done by zero-padding the array for each segment in the frequency domain. The effects of our frequency-shifting approach on the QPO power spectrum is shown in Figure 4.2, which shows the time-averaged power spectrum before (in red) and after (in black) the frequency shift is applied. The shifted QPO signal is stronger, narrower and more symmetric than the unshifted signal, making it more suitable for our spectral-timing analysis. The quality factor of the QPO in the shifted power spectrum of all 5 PCUs is $Q = 192$.

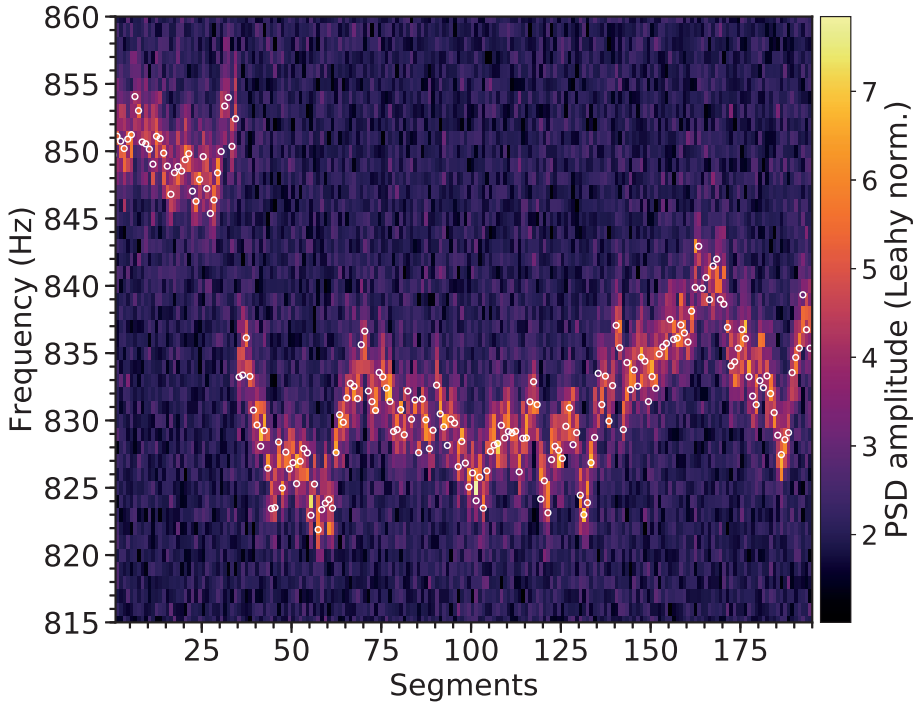

Figure 4.1: Dynamical power spectrum of the lower kHz QPO with all 5 PCUs combined in the energy range 3–20 keV. Each segment of data is 32 seconds long. The colour mapping shows the amplitude of the power spectrum with Leahy normalization. The QPO centroid frequencies in each segment were fitted for in XSPEC, and are indicated with the white circles. We see that the QPO is fairly narrow in each segment of data, and rapidly changes frequency throughout the observation.

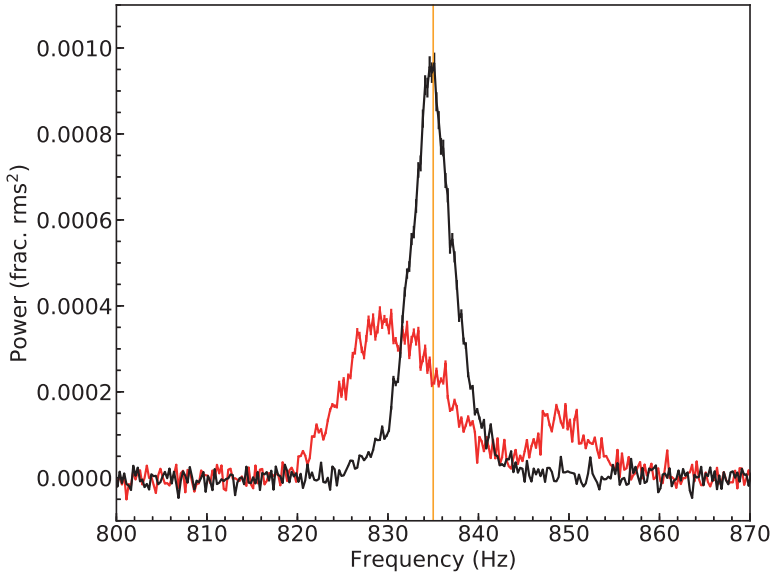

Figure 4.2: Unshifted (red) and shifted (black) time-averaged power spectra, for all 5 PCUs in the energy range 3–20 keV. The vertical orange line corresponds to the assumed QPO centroid at 835 Hz. The quality factor of the shifted QPO shown here is $Q = 192$.

We applied the same shifting to the segments of the cross spectrum in each energy channel. By shifting each segment of the data to line up the QPOs before averaging the segments together, we are assuming that any interesting spectral-timing behaviour is dependent on the relative phase of the QPO cycle, not on the frequency value. Following the spectral-timing approach of Uttley et al. (2014) we improve cross-spectral signal-to-noise by crossing individual 'channels of interest' with a broad reference energy band, for which (following Stevens & Uttley 2016) we use a different set of detectors, to ensure complete statistical independence. Thus, the channels of interest use data from PCUs 2–4 and the reference band data is from PCUs 0–1. Figure 4.3 shows the shifted and averaged cross spectrum amplitude in each channel of interest, with Leahy normalization (Leahy et al. 1983). Each channel's cross spectrum was linearly re-binned to a frequency resolution of 0.25 Hz (from 0.03125 Hz) to improve the signal-to-noise. We see in Figure 4.3 that the lower kHz QPO amplitude is strongest relative to the noise level at around 8–9 keV. Furthermore, the QPO amplitude is weak compared to the noise outside 3–20 keV, the energy range in which we fit the energy spectra in Sections 4.2.2 and 4.4.

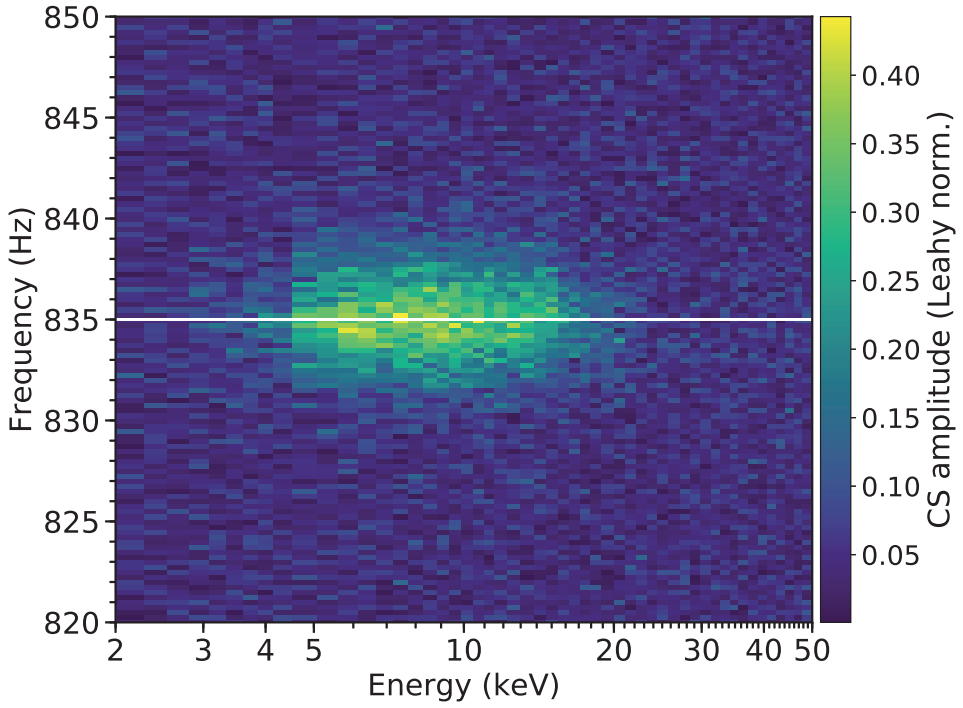

Figure 4.3: Cross spectrum amplitude of the lower kHz QPO per detector energy channel, using PCUs 0–1 for the reference band and PCUs 2–4 for the channels of interest. The averaged cross spectra per energy channel have been shifted per segment using the offset computed from the dynamical power spectrum. The horizontal white line indicates 835 Hz, the value to which the QPO centroids were lined up. It is evident that the shift-and-add computed from the dynamical power spectrum appropriately lines up the QPO peaks in the cross spectrum, too.

4.2.1 Lag-energy spectrum

We computed the lag-energy spectrum for the QPO over 825–845 Hz using our own code,[1] following the outline in Uttley et al. (2014). The lag-energy spectrum indicates the phase information of the cross-spectrum, so it compliments the basic amplitude information shown in Figure 4.3. For completeness, we note again that each segment of the cross spectrum per energy channel was shifted to a centroid value of 835 Hz. The lag-energy spectrum, plotted in Figure 4.4, crosses the zero lag mark (dashed line) at the average energy of the QPO signal in the reference band. The observed lags are in the sense that soft photon variations lag those at harder energies and are consistent with those shown in Vaughan et al. (1997, 1998), Barret (2013) and Cackett (2016) for this same observation. We will come back to the lag-energy spectrum in Section 4.4.3, where we check our phase-resolved spectroscopy results by simulating the corresponding lag-energy spectra and comparing them with the data shown here.

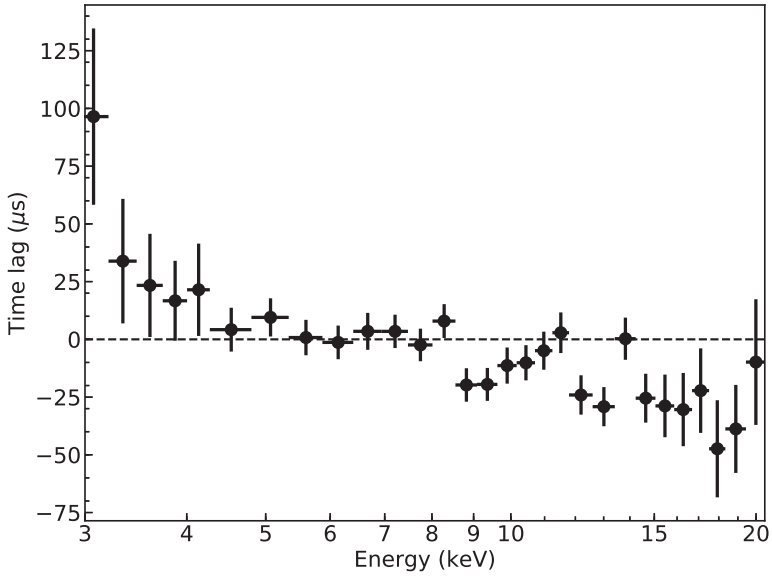

Figure 4.4: Lag-energy spectrum over the frequency range 825–845 Hz, computed from the shifted, averaged cross spectrum. A positive sign indicates that the variation in that energy channel arrived after the variation in the broad reference band.

4.2.2 Time-averaged energy spectrum

We extracted the time-averaged energy spectrum from the three orbits in obsID 10072-05-01-00 in 64-channel event-mode energy binning. We elected to use the event-mode data for fitting the time-averaged energy spectrum since it has the same energy binning as the phase-resolved spectra fitted in Section 4.4.

For fitting the time-averaged and phase-resolved energy spectra we consider two variants based on the 3-component spectral model from Lin et al. (2007) (see Armas Padilla et al. 2017 for the most recent use of this type of model applied to 4U 1608–52 data). Model 1 has a straightforward multi-colour disc blackbody (DISKBB; Mitsuda et al. 1984), a power-law with a low- and high-energy cut-off for the Comptonized emission (NTHCOMP; Zdziarski et al. 1996; Życki et al. 1999), and a broad iron line (GAUSS). Model 2 contains all the components of model 1 and adds to it a single-temperature blackbody (BBODYRAD) with a temperature that is tied to the low-energy cut-off temperature of NTHCOMP.

Physically, the Comptonized component is thought to correspond primarily to the emission from the boundary layer rather than a more diffuse 'corona', since it has a distinctly lower-energy cut-off than the Comptonized components seen in black hole systems and in lower-accretion rate states of accreting neutron stars (Done & Gierliński 2003; Done et al. 2007). The additional single-temperature blackbody com-

#	Time-averaged spectral model	χ^2 / d.o.f.
1	TBNEW_GAS × (DISKBB + NTHCOMP + GAUSS)	13.58/20
2	TBNEW_GAS × (DISKBB + BBODYRAD + NTHCOMP + GAUSS)	13.30/19

Table 4.1: A summary of the model fits to the time-averaged event-mode spectrum. The spectra are plotted with the different models and components in Figure 4.5.

ponent in model 2 is included to allow for the possibility of some visible component of accretion-driven thermal surface emission from the neutron star (e.g., providing the source of seed photons for the boundary layer Comptonizing region). For simplicity, we only include a single Gaussian component to model iron line emission, likely from the disc, but we note that it should be accompanied by a hard reflection continuum. XSPEC models corresponding to reflection of a low-temperature Comptonization continuum do not yet exist, and so rather than including reflection for an incorrect continuum, we choose to neglect this component but note that its variability in the phase-resolved spectra (if any) can be inferred from variations of the iron line modelled by the Gaussian component. Where appropriate, we will comment on the effects on our phase-resolved results of neglecting the reflection continuum.

Using XSPEC version 12.9.1l (Arnaud 1996) we fitted the energy spectra in the energy range 3–20 keV. The solar abundance table was set to `vern` (Verner et al. 1996) and the photoionization cross-section table was set to `wilm` (Wilms et al. 2000). The absorption component used is TBNEW_GAS version 2.3.2.[2] Here and in Section 4.4, the neutral hydrogen column density N_H was frozen at 1.07×10^{22} (as determined by Armas Padilla et al. 2017 for a recent soft-state observation of the source). The systematic error was set to 0.6 per cent as in Cackett (2016). To ensure proper convolution with the broad PCA instrument response, we used the `energies` command to evaluate the models in the range of 0.5–50 keV in 1000 logarithmically-spaced bins.

A summary of the models and their fits to the time-averaged spectra are in Table 4.1. The spectrum of model 1, total fitted model with components, and ratio of the data to the total model (to show residuals) is plotted in Figure 4.5. We find that models 1 and 2 show residuals without any clear structure. Model 2 requires a negligible contribution from the single-temperature blackbody component, implying that the seed photon source for the Comptonized region is not directly visible. However, to allow for the possibility that the single-temperature blackbody component is required once phase-variations are considered, we will continue to apply both models in the phase-resolved spectral fits.

The RXTE PCA spectral correction curve PCACORR (García et al. 2014), which reduces spectral systematic errors to ∼0.1 per cent, is not available for data taken during Gain Epoch 1 of RXTE's operation, so we could not apply this correction

[2] Developed by J. Wilms et al.: `http://pulsar.sternwarte.uni-erlangen.de/wilms/research/tbabs/`

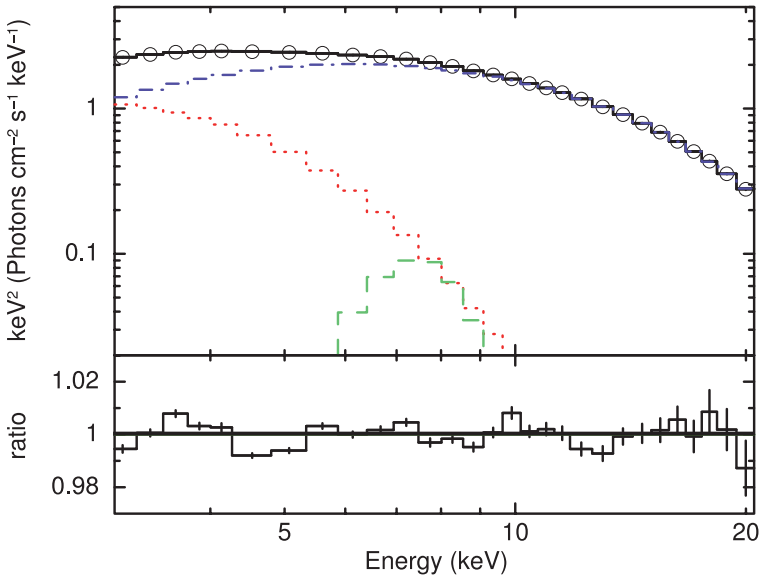

Figure 4.5: The time-averaged event-mode spectrum fit with model 1, showing the ratio of the total model with the data. The data are shown as black open circles, and the combined model is in solid black. The errors on the data represent the statistical 1σ confidence interval. The disc multi-colour blackbody is in dotted red, the boundary layer power-law is in dash-dotted purple and the Gaussian component associated with reflection from the disc is in dashed green. Since for model 2 the surface single-temperature blackbody is negligibly small, the fitted spectrum and residuals are not visibly different than model 1.

directly to our spectra. Due to the instrument requirements for gain shifting, specific event mode channels may not be assigned any counts, thus it was necessary for Stevens & Uttley (2016) to ignore detector channel 11 (at $\sim$4.3 keV in Gain Epoch 1). In the present analysis we can include all detector channels, since we combine data from multiple detectors. Note however that channel 11 still records systematically lower counts than adjacent channels, as can be seen in Figures 4.3 and 4.6.

4.3 Method

We followed a similar method to obtain phase-resolved spectra as in Stevens & Uttley (2016), with some differences needed to best apply this technique to the lower kHz QPO from 4U 1608–52, which we highlight here. For further details and explanation of the method, we refer the reader to Stevens & Uttley (2016).

As explained in Section 4.2, we used the dynamical power spectrum of all 5 PCUs to determine the QPO centroid frequency in each 32-second segment of data. Using

these centroids, we then computed the offset needed to shift the QPO to 835 Hz in each segment and applied it to the cross spectrum in each energy channel. We then averaged the segment cross-spectra together per energy channel to produce a mean cross spectrum per energy channel that has the QPOs lined up. The software to reproduce this can be found at the URL in Appendix 4A.

4.3.1 Filtering

The shifted and combined lower kHz QPO is a very narrow feature in frequency space that stands out from the uncorrelated Poisson noise component in a narrow frequency range (see Figure 4.2). However, to make use of the full time resolution of the data, the CCF method requires that cross spectra are obtained up to high frequencies (4096 Hz). Since the CCF is the inverse transform of the Fourier cross-spectrum, it includes all frequency information, meaning that the kHz QPO signal in the CCF is dominated by the strong Poisson noise component associated with the higher frequencies. Therefore, in order to reduce the noise, we filtered the averaged cross spectra in each energy channel to keep only the lower kHz QPO. Note that this step is additional to the approach used by Stevens & Uttley (2016) to study the Type B QPO in GX 339-4. This is because in the CCF the lower frequency Type B QPO (at 5.2 Hz versus 835 Hz here), dominated over the noise at the 7.8125 ms time-resolution studied, so that no filtering was required.

The cross spectra were re-binned to a frequency resolution of 0.25 Hz for fitting purposes, and a Leahy normalization was applied to the cross spectrum amplitude (Leahy et al. 1983). We fitted the cross spectrum amplitude in XSPEC over the frequency range $700 - 1000$ Hz (Arnaud 1996), with a power-law plus single Lorentzian component. The power-law index was fixed at 0 to match the flat noise level due to Poisson noise (which is close to zero in cross-spectral amplitude, subject only to the small bias noted by Vaughan & Nowak 1997), the Lorentzian centroid energy was fixed at 835 Hz, the Lorentzian width was tied across each cross spectrum, and the power-law normalization and Lorentzian normalization were free.

With the resulting fits, we constructed a ratio at each Fourier frequency of the Lorentzian amplitude to the entire model amplitude, for the shifted and averaged cross spectrum in each energy channel. Then taking this ratio as our filter, we multiplied it by the complex cross spectrum at full (0.03125 Hz) resolution to get a filtered cross spectrum. At this stage, we apply a normalisation to the raw cross-spectrum of $2/N$, where N is the number of time bins per light curve segment used to measure the cross-spectrum (262144). This ensures that our cross-spectrum is in units of $(\mathrm{count\,s^{-1}})^2/\mathrm{Hz}$. We checked that in the filtered cross spectra, the integrated rms of the QPO is preserved in each energy channel. This shifted, averaged, filtered cross spectrum per energy channel was then used to compute the cross-correlation function.

4.3.2 Cross-correlation function

To get the cross-correlation function (CCF), we took the inverse discrete Fourier transform of the filtered cross spectrum, as outlined in Stevens & Uttley (2016). The CCF was normalized by the integrated absolute rms of a Lorentzian model of the QPO in the shifted, averaged reference band power spectrum. This normalization removes the contribution of the reference band amplitude to the CCF, so that the CCF contains only contributions from the energy channels of interest relative to the reference band. To obtain the reference band QPO rms, we fitted the reference band power spectrum in XSPEC with the model POW + LOR. As with the cross-spectral fitting in Section 4.3.1, the power-law index was fixed at 0, the Lorentzian centroid energy was fixed at 835 Hz, and the Lorentzian width, power-law normalization, and Lorentzian normalization were left free. It was not necessary to tie any parameters across spectra from different channels of interest since we only fitted the single reference band power spectrum here.

As discussed in Stevens & Uttley (2016), the CCF yields a count rate per energy channel in each time-delay bin, as deviations from the mean count rate per energy channel. The energy-dependent CCF is plotted in Figure 4.6 as the CCF amplitude divided by the mean count rate per energy channel, i.e.,equivalent to the fractional deviation from the mean spectrum. It shows the now-periodic (since it is locked in phase to the reference energy band) signal of the QPO clearly as a series of flux rises and falls. The amplitude of the variation is much smaller at lower energies than at higher energies, consistent with what is known from the fractional rms spectrum (Berger et al. 1996).

The final step is to add the energy-dependent CCF values in each time-delay bin to the mean energy spectrum in the same count-rate units, to produce phase-resolved spectra that can be fitted using standard spectral models in XSPEC. The fits to the phase-resolved spectra are described in Section 4.4.

4.3.3 Parameter variation fitting and bootstrapping

In fitting the phase-resolved spectra in Section 4.4, we found that a combination of untied spectral parameters give a statistically and physically meaningful result. These untied spectral model parameters are not the same as the flux variations inferred from the CCF, but their variations can still be approximated with a sinusoidal curve, which can be used to estimate the amplitude and relative phase of the parameter variations. Following the protocol of Stevens & Uttley (2016), we fitted those parameter variations with a curve including a fundamental and its harmonic:

$$y(t) = A_1 \sin(2\pi t + \phi_1) + A_2 \sin(4\pi t + \phi_2) + D\,, \tag{4.1}$$

where t is time, A_1 and A_2 are amplitudes, ϕ_1 and ϕ_2 are phase offsets, and D is a y-axis offset. We obtain values of A_1, A_2, ϕ_1, ϕ_2, and D, and further use these

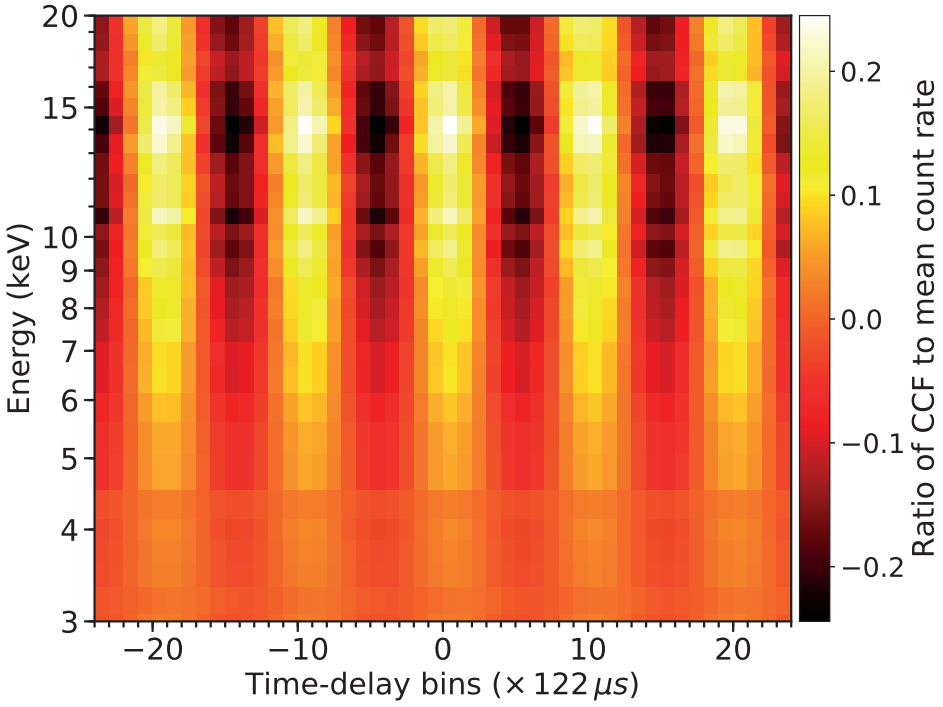

Figure 4.6: The energy-dependent cross-correlation function (CCF). The colour map shows the fractional deviation of the CCF amplitude from the mean count rate in each energy channel. The QPO signal is clear, with a period of 9.8 time-delay bins corresponding to the centroid QPO frequency of 835 Hz (~ 1.2 ms period).

parameters to determine an rms amplitude of the model curve for each untied spectral parameter. Note that, since there are only 9.8 phase bins per QPO cycle and the curves can vary steeply within a bin, we fit the average value of the model curve in each phase bin, rather than the model evaluated at the bin centres. Later, we use the model curves to determine a phase offset between varying parameters, and we use the rms of the curves to determine the amplitudes of the parameter variations.

As emphasized in Stevens & Uttley (2016), the statistical fluctuations between CCF time-delay bins are correlated with one another in a given energy channel. As such, we cannot interpret the χ^2 fit statistic returned from the fits in the usual manner (though we report it for completeness), meaning that we cannot estimate parameter uncertainties from our model fits and uncertainties on the resulting sinusoidal fits to parameter variations, in the usual way. To remedy this in a statistically rigorous way, we follow our approach from Stevens & Uttley (2016) and use a bootstrap analysis to determine the errors in model parameters and sinusoidal fits. Our bootstrap analysis procedure entailed randomly selecting (with replacement) 195 CCF segments of the

195 available to average together, constructing and fitting the phase-resolved spectra with the designated model, and fitting the untied parameter variations as mentioned above. For our 195 segments of data, we carried out 5422 iterations of this process,[3] yielding errors and confidence contours on fit parameters and rigorous estimates of the uncertainties on the relative phases of the model-parameter variations.

4.4 Results

We selected energy spectra at each time-delay bin of the CCF (hereafter, "phase bin") to fully cover four QPO cycles. Since the lower kHz QPO has a centroid frequency of 835 Hz and the data has a time resolution of 2^{-13} s ($\sim 122\,\mu$s), one QPO cycle covers 9.8 phase bins. We used phase bins -20 to +20 (inclusive and including 0), giving 41 phase-resolved spectra (hereafter, "phase spectra"). The zero mark was defined to be at phase bin -20, when the QPO is at a maximum in the CCF (see Figure 4.6). Note that the choice of this zero mark is arbitrary, since only the relative phases of the model component variations are physically relevant. For all phase spectra, the *RXTE* PCA instrumental response matrix was determined for the combination of PCUs 2–4 (all layers and anodes) to match the data selection used for the channels-of-interest.

To fit the phase spectra in a systematic way, we followed the procedure explained in Stevens & Uttley (2016), which we briefly re-cap here. We loaded the phase spectra to fit simultaneously in XSPEC and began with all spectral parameters tied across phase spectra. Then we systematically untied each parameter individually and assessed the fit. The parameterisation with the lowest χ^2 remained untied, and we investigated whether a second parameter could be untied. This continued until the improvement in χ^2 was less than the change in the degrees of freedom (d.o.f.) of the fit. Note that although χ^2 cannot be used in a statistically rigorous way due to the correlation of errors between phase bins, here we use it simply as a guide to whether a fit is improved relative to the errors or not and hence set a criterion as to when to stop fitting. See Table 4.2 for a summary of the best-fitting parameterisations to the phase spectra. The details of the best parameterisation of each model are discussed in the following subsections.

4.4.1 Model 1

Model 1, which in XSPEC terms is TBNEW_GAS $\times$ (DISKBB + NTHCOMP + GAUSS), is the simplest of our two models. To effectively constrain the fits for so many spectra, we chose constrained parameter ranges informed by our fits to the time-averaged spectra. The disc blackbody temperature $T_{\rm disc}$ was tied across all phase spectra and was left free to fit in the range 0.5–1.5 keV. In initial fits to the time-averaged event-

[3] Using $M \cdot \big(\ln(M)\big)^2$, where M is the number of segments (Feigelson & Babu 2012).

#	Phase-resolved spectral model	χ^2 / d.o.f.	Untied	Lag χ^2 / d.o.f.
1	TBNEW_GAS $\times$ (DISKBB + NTHCOMP + GAUSS)	1981.2/1060	T_{bb}, T_e, Γ	32.4/29
2	TBNEW_GAS $\times$ (DISKBB + BBODYRAD + NTHCOMP + GAUSS)	1983.8/1059	T_{bb}, Γ, T_e	29.9/29

Table 4.2: A summary of the best parameterisation for the two spectral models tested. The first column gives the number of each spectral model, for reference throughout the paper, and the second column states the full model. The third column gives the χ^2 fit statistic over degrees of freedom (d.o.f.) (this may only be interpreted in a meaningful statistical sense by assuming that the data are statistically independent, which they are not; see Section 4.3.3). The "Untied" column lists which model parameters are untied between phase spectra and therefore allowed to vary with phase, to give the best parameterisation. The column "Lag χ^2 / d.o.f." gives the fits for lag-energy spectra simulated from the phase-resolved spectral model compared with the data (see Section 4.4.3 and Figure 4.14).

mode spectrum, we found that the parameter space had many local minima with different $N_{\rm disc}$ values, so it was initially frozen at 265 (the estimated normalization for an inner disc radius of $7\,GM/c^2$, binary inclination of 40°, and a distance of 7.8 kpc), all the other parameters fitted, then $N_{\rm disc}$ freed to fit in the range 150–750 and then all parameters fitted again. The power-law index Γ was allowed the range 1.5 to 3.0, the high-energy cut-off temperature $T_{\rm e}$ was allowed the range 2–20 keV, the low-energy cut-off temperature $T_{\rm bb}$ was allowed the range 0.7–3.5 keV, the input seed spectral type was frozen at 0 for a single-temperature blackbody seed spectrum and the power-law normalization $N_{\rm PL}$ was left to be completely free. The iron line centroid energy $E_{\rm line}$ was free in the range 5.5–7.5 keV (given such a broad range around the 6.4 keV rest energy because the GAUSS model doesn't account for Doppler shifting), the iron line width $\sigma_{\rm line}$ was free in the range 0.1–1.0, and the iron line normalization $N_{\rm line}$ was free in the range 0–1.

The full table of the systematic fitting for model 1 is shown in Appendix 4B in Table 4.5. In this procedure, we tied all the free parameters together across the phase spectra (Table 4.5 model ID 1.0.1), then stepped through untying one parameter and fitting the spectra (model IDs 1.1.1–1.1.9). Then, we selected the best fit for one free parameter (model IDs 1.1.5), kept that parameter untied, and stepped through the other free parameters untying them one at a time (Table 4.5 model IDs 1.2.1–1.2.8). Since the parameter space has many local minima, we repeated this procedure for the second best fit with one free parameter (Table 4.5 model ID 1.1.6), shown in model IDs 1.2.9–1.2.15. From there, we selected the best fit with two free parameters (Table 4.5 model ID 1.2.4), and stepped through untying a third parameter. We continued adding untied parameters to the fit until the reduction in degrees of freedom was greater than the improvement in the χ^2 fit statistic. The best-fitting parameterisation for model 1, model ID 1.3.3 from Table 4.5, has $T_{\rm bb}$, $T_{\rm e}$, and Γ untied. The values of each parameter for this best parameterisation are listed in Table 4.3.

To illustrate how the untied parameters vary between phase spectra, we plotted their values in Figure 4.7. Figure 4.8 (left panel) shows two example spectra with model-decomposition and data/model ratios, for the phase bins closest to phases 0 and 0.5 (respectively maxima and minima in the flux cycle). The spectra show that the flux maximum corresponds to a harder spectrum, with Γ at a minimum in the parameter variation curve, and vice versa for the flux maximum. The other parameters, $T_{\rm bb}$ and $T_{\rm e}$ are also correlated with the variations in Γ.

As outlined in Section 4.3.3 and shown in solid black in Figure 4.7, we fitted the parameter variations with a curve to determine the phase of the parameter variation. The first maximum of the curve, within the normalized phase 0–1, is marked with a dashed black line, and the darker gray lines to either side of the black line mark the error region from the fitted curve phase of the bootstrap samples. Our bootstrap samples of the fitted parameter variations show that phases of different parameters are correlated, thus it is more important to consider the phase differences (obtained from

Component	Parameter	Value	Notes
TBNEW_GAS	$N_{\rm H}$ ($\times 10^{22}$ cm^{-2})	1.07	Frozen
DISKBB	$T_{\rm disc}$ (keV)	$0.89^{+0.01}_{-0.03}$	Tied
DISKBB	$N_{\rm disc}$	270^{+28}_{-13}	Tied
NTHCOMP	Γ	$2.024^{\dagger}$	Untied
NTHCOMP	$T_{\rm e}$ (keV)	$2.709^{\dagger}$	Untied
NTHCOMP	$T_{\rm bb}$ (keV)	$0.869^{\dagger}$	Untied
NTHCOMP	$N_{\rm PL}$	0.20 ± 0.01	Tied
GAUSS	$E_{\rm line}$ (keV)	$7.07^{+0.08}_{-0.03}$	Tied
GAUSS	$\sigma_{\rm line}$ (keV)	$1.0^{+0}_{-0.2}$	Tied
GAUSS	$N_{\rm line}$	$4.7^{+0.1}_{-0.8} \times 10^{-3}$	Tied

Table 4.3: Values of the best-fitting parameterisation for model 1. Mean values are listed for the untied parameters, indicated by †. Errors on the tied parameters were computed for the 90% confidence interval with XSPEC's MCMC error routine. XSPEC gives a fit statistic of $\chi^2 = 1981.2$ for 1060 d.o.f. (fitting 41 phase spectra simultaneously).

the best fit to the data) and their distribution (which we obtain via our bootstrap sample of parameter curves). For completeness sake, however, we show the histograms of the curve phase and the maximum phase in Appendix 4B in Figures 4.16 and 4.17.

We found that the variation in $T_{\rm bb}$ leads the variation in Γ by 0.026 (1σ confidence interval from bootstrapping, centred on the median value: [0.015–0.047]), the variation in $T_{\rm bb}$ leads the variation in $T_{\rm e}$ by 0.039 [0.031–0.06], and the variation in Γ leads the variation in $T_{\rm e}$ by 0.013 [0.011–0.018]. Histograms of the phase offsets between Γ and $T_{\rm bb}$ and $T_{\rm e}$ and Γ are plotted in Figure 4.9. These distributions show that the phase offsets between the physical parameters are significant. We will show in Section 4.5 that these phase offsets are crucial for understanding how the energy-dependent lags are related to the measured spectral variations.

We further computed the rms amplitude of each fitted curve, and found the following (measured from our best-fit curve, with 1σ confidence interval from the distribution of bootstrap sample rms): Γ has an rms amplitude of 9.3% [8.2%–10.8%], $T_{\rm e}$ has an rms amplitude of 8.1% [7.4%–9.5%], and $T_{\rm bb}$ has an rms amplitude of 2.1% [1.9%–2.4%].

The error bars on $T_{\rm bb}$ are quite large, but our systematic fitting procedure shows that $T_{\rm bb}$ is indeed required to vary with QPO phase (see Table 4.5), and the large error bars are due to degeneracy of the mean value of $T_{\rm bb}$ with the value of $T_{\rm disc}$ which is tied across phase spectra. The parameter degeneracies for this and other parameter combinations are demonstrated for a single phase bin (bin 5) in Figure 4.10. Similar correlations are found between parameters in other phase bins. Due to these degeneracies, absolute uncertainties on parameter values can be relatively large, but

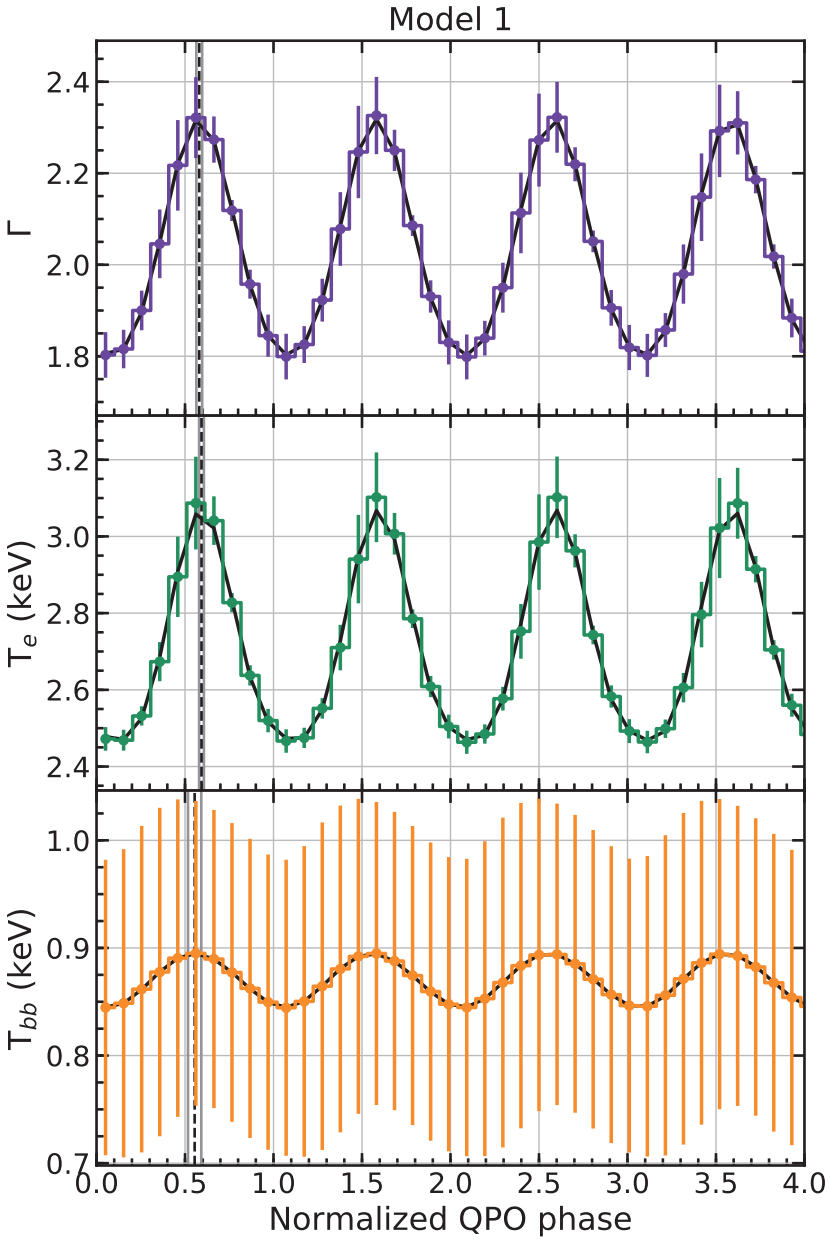

Figure 4.7: The varying parameters of model 1 that best fit the phase-resolved spectra across four QPO cycles. Error bars are obtained from standard deviations of the bootstrapped distribution of parameter values in each phase bin. In all three panels, the parameter variations are fitted with a curve, shown in solid black, and the phase of the maximum of that curve is marked with a dashed black line. The dark gray lines to either side of the black dashed line mark the 1σ error region on the phase.

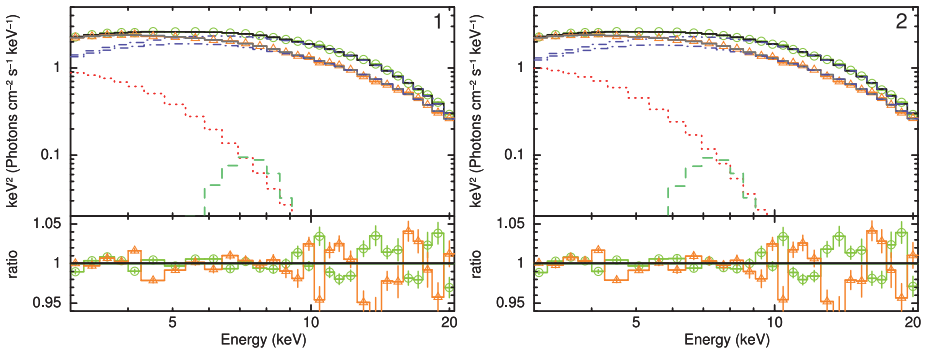

Figure 4.8: Spectra with models at normalized phase 0 (in green circles) and 0.5 (in orange triangles) overplotted in each window for models 1 (LEFT) and 2 (RIGHT). Lines of the model components have the same meanings as in Figure 4.5. In model 2, the BBODYRAD component has a maximum of $\sim 2^{-4}$ that could not be shown on a scale that appropriately displayed the other components for comparison. These spectra visually demonstrate the parameter variations explained in Sections 4.4.1–4.4.2.

the *relative* variations between parameters remain well-constrained, as shown by the small but significant phase offsets that we are able to measure between the parameter variation curves.

Despite the large degeneracy between $T_{\rm disc}$ and $T_{\rm bb}$ there is only a minimal correlation between the fractional rms of the $T_{\rm bb}$ variation and the value of $T_{\rm bb}$ (Figure 4.11), indicating that the fractional amplitude of $T_{\rm bb}$ variations does not depend strongly on the specific realisation of model parameters. The same applies to the phase differences, i.e., in response to changes in best-fitting $T_{\rm disc}$ the $T_{\rm bb}$ curve is shifted in normalisation but not in phase, therefore the phase offset between blackbody seed temperature and other parameters remains well-constrained.

Thus we conclude that, according to the model 1 parameterisation of the data, the electron temperature, photon index and seed photon temperature of the Comptonizing region all vary with QPO phase and are positively correlated with one another (i.e., close to being in phase), but with seed photon temperature varying first followed by the photon index and then a change in electron temperature.

4.4.2 Model 2

Model 2, which in XSPEC terms is TBNEW_GAS × (DISKBB + BBODYRAD + NTHCOMP + GAUSS), builds on model 1 by also including a separate seed blackbody for the neutron star surface (BBODYRAD). In this sense, the model assumes that some part of the blackbody component providing seed photons in model 1 is also visible to the observer. The seed blackbody normalization $N_{\rm bb}$ was free to fit in the range 0–100, and all other parameters were given the same ranges as in model 1. As with model

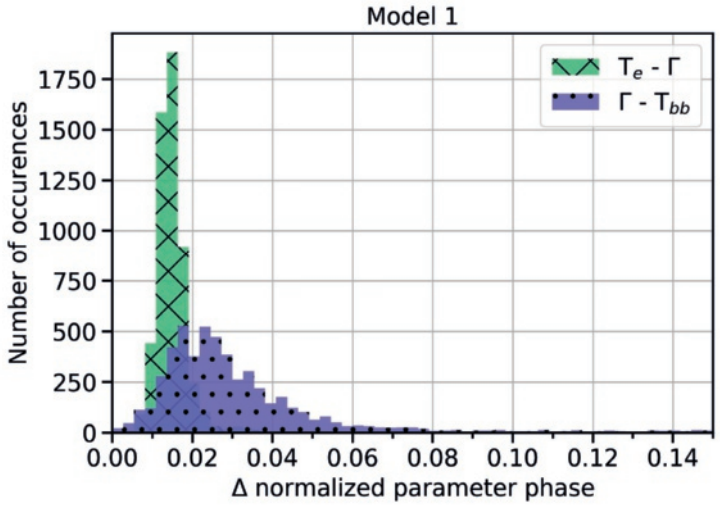

Figure 4.9: Distribution of bootstrap sample phase offsets between parameter variations of untied parameters for the best parameterization of model 1 (out of 5421 good bootstrap samples). The phase offset between $T_{\rm e}$ and $T_{\rm bb}$ is the sum of the offset between $T_{\rm e}$ and Γ and between Γ and $T_{\rm bb}$. As measured from the best-fit to the data, Γ leads $T_{\rm e}$ by 0.013 (1σ confidence interval from bootstrapping, centred on the median value: [0.011–0.018]), $T_{\rm bb}$ leads Γ by 0.026 [0.015–0.047], and $T_{\rm bb}$ leads $T_{\rm e}$ by 0.039 [0.031–0.06].

1, initial fits with model 2 to the time-averaged event-mode spectrum showed that the parameter space preferred a small $N_{\rm disc}$, so we initially froze $N_{\rm disc}$ at 265, fitted the rest of the parameters, then freed it and fitted again. The seed blackbody temperature of the power-law was tied to the surface blackbody temperature within each phase spectrum. The full table of systematic spectral fitting for model 2 is given in Table 4.6.

The spectral parameter values for the best-fitting parameterisation of model 2 are listed in Table 4.4 (model ID 2.3.4 from Table 4.6). The values of the untied parameters ($\rm T_{bb}$, Γ, and $\rm T_e$) at each phase bin are shown in Figure 4.12, with example spectra and data/model ratios from phase bins closest to 0 and 0.5 shown in Figure 4.8. We note again that the error bars on spectral parameter variations can be large, particularly on the values of $T_{\rm bb}$. This is due to similar degeneracies with the disc temperature $T_{\rm disc}$ as for model 1, but also as for model 1, do not lead to significant uncertainties in the amplitude or phase offset of the parameter variations. As with model 1, we show the histograms of the curve phase and the maximum phase for model 2 in Appendix 4B in Figures 4.20 and 4.21 for completeness.

We found that the variation in $T_{\rm bb}$ leads the variation in Γ by 0.036 (1σ confidence interval from the bootstrap sample, centred on the median value: [0.041–0.098]), the variation in $T_{\rm bb}$ leads the variation in $T_{\rm e}$ by 0.05 [0.055–0.111], and the variation in Γ

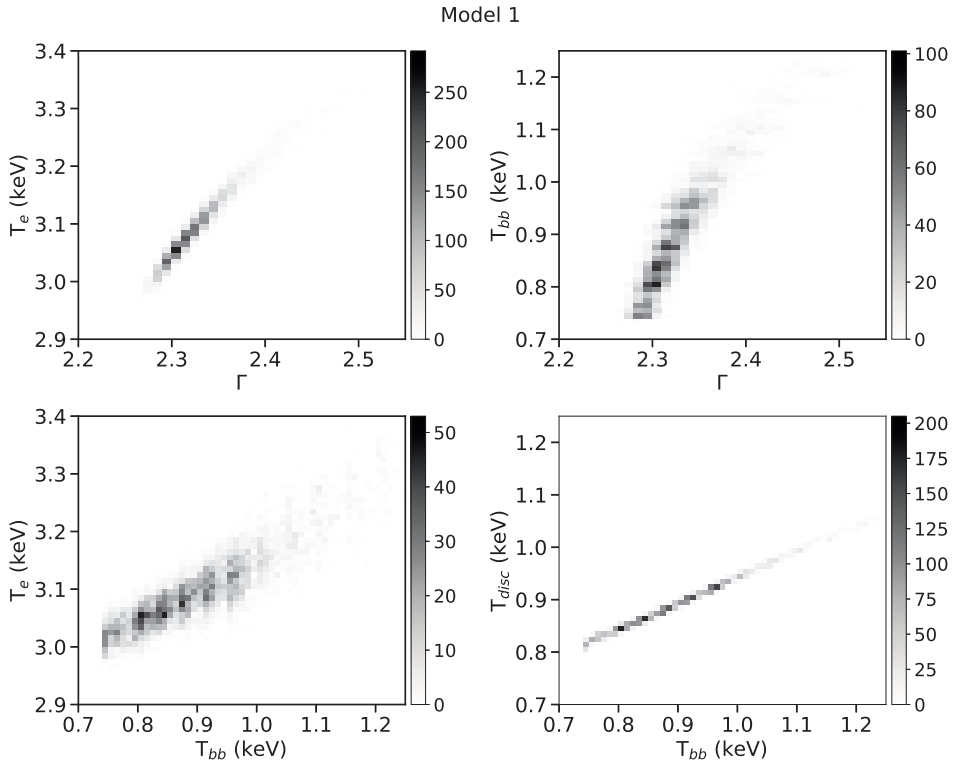

Figure 4.10: Model 1: 2-D histograms of spectral parameter values in phase bin 5 for each bootstrap sample, showing the correlations between spectral parameters. The colour bars indicate the number of occurrences, out of the 5421 good bootstrap samples for model 1. While $T_{\rm disc}$ was tied across phase in model 1, we still found a very strong correlation between $T_{\rm disc}$ and $T_{\rm bb}$.

leads the variation in $T_{\rm e}$ by 0.014 [0.011–0.016]. A histogram of the phase differences between the untied parameters in each bootstrap sample is shown in Figure 4.13. In terms of variability amplitudes, Γ has an rms amplitude of 9.9% [9.9%–12.0%], $T_{\rm e}$ has an rms amplitude of 8.6% [8.1%–10.0%], and $T_{\rm bb}$ has an rms amplitude of 2.0% [1.2%–1.7%]. Note that the overall trends in phase offset are very similar to those seen for model 1, likely because the additional single-temperature blackbody component contributes only a small amount to the spectrum (as also seen in the time-averaged spectra). From this we conclude that both the time-averaged and phase-resolved spectra are consistent with there being no visible blackbody component associated with, e.g., the neutron star surface.

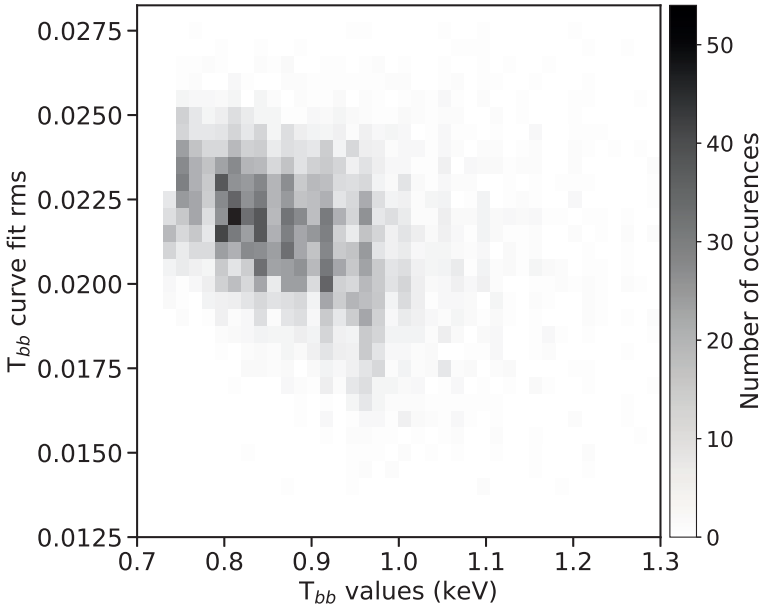

Figure 4.11: Model 1: 2-D histogram of the $T_{\rm bb}$ value in phase bin 5 with the rms amplitude of the $T_{\rm bb}$ variation for each bootstrap sample. As with Figure 4.10, the colour bars indicate the number of occurrences. We see that there is not a strong correlation between the value of $T_{\rm bb}$ and the rms amplitude of the variation.

4.4.3 Comparison with the lag-energy spectrum

For a final check on the best-fitting parameterisation of the spectral models, we simulated lag-energy spectra and compared them to the data via the χ^2 (listed in the final column in Table 4.2). We followed the same procedure as Stevens & Uttley (2016) to simulate QPOs with energy spectral parameters sampled from the best-fitting curves fitted to the parameter variations, and computed time lags from the cross spectrum averaged over the same Fourier frequency range as was used for the data, 825–845 Hz. Both model parameterisations match the lag-energy data well, with χ^2 values of 32.4 and 29.9 for models 1 and 2, respectively, for 29 degrees of freedom. Neither model is clearly preferred, consistent with their similar spectral fit parameters. Figure 4.14 shows the lag-energy spectra simulated from the best parameterisations of both models plotted with the data from Figure 4.4. This plot and the accompanying fits illustrate how even though the bootstrapped error bars on the untied parameter variations allow for no variation of some parameters, the untied parameters are required to vary to fit the data

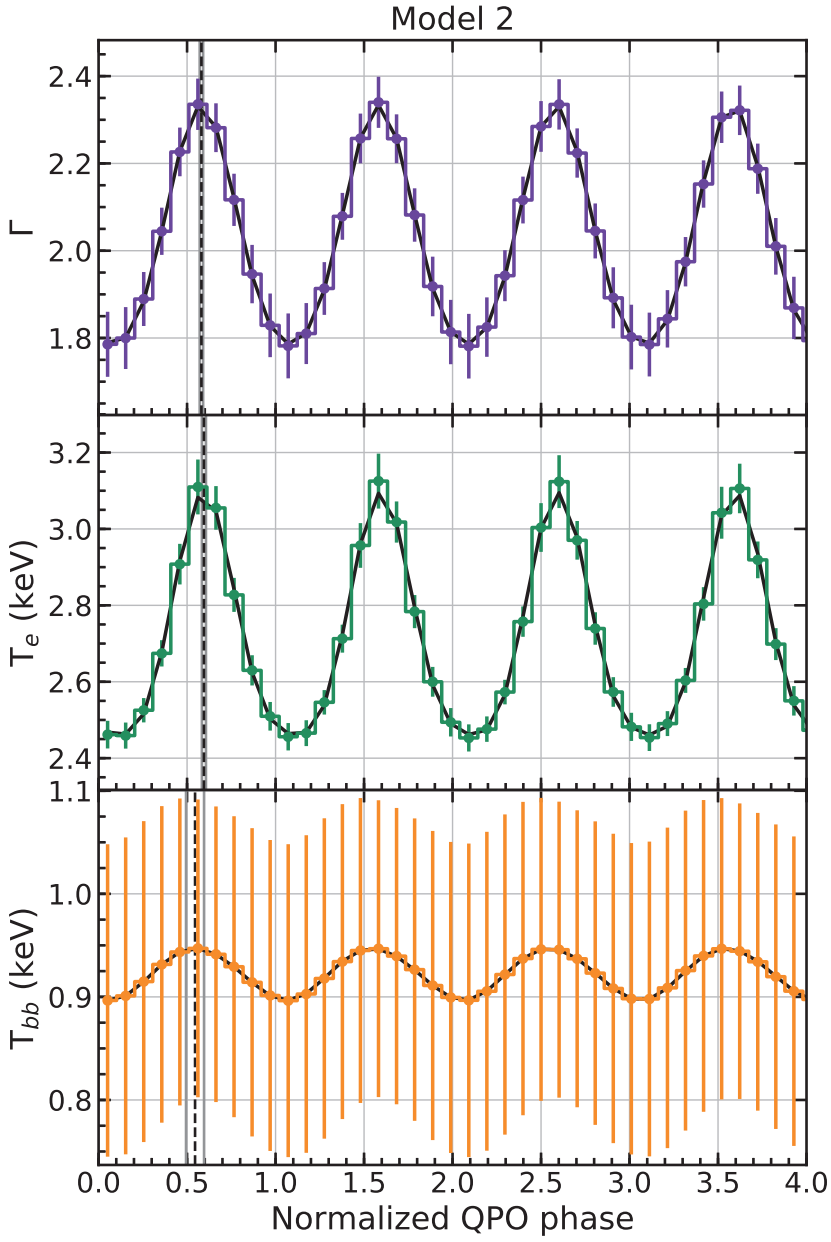

Figure 4.12: The varying spectral parameters of the best parameterisation of model 2. Lines and error regions have the same meanings as in Figure 4.7.

Component	Parameter	Value	Notes
TBNEW_GAS	$N_{\rm H}$ ($\times 10^{22}$ cm^{-2})	1.07	Frozen
DISKBB	$T_{\rm disc}$ (keV)	$0.916^{+0.005}_{-0.029}$	Tied
DISKBB	$N_{\rm disc}$	265^{+27}_{-3}	Tied
BBODYRAD	$T_{\rm bb}$ (keV)	$0.921^{\dagger}$	Untied
BBODYRAD	$N_{\rm bb}$	0.05^{+17}_{-0}	Tied
NTHCOMP	Γ	$2.019^{\dagger}$	Untied
NTHCOMP	$T_{\rm e}$ (keV)	$2.709^{\dagger}$	Untied
NTHCOMP	$T_{\rm bb}$ (keV)	$0.921^{\dagger}$	Untied
NTHCOMP	$N_{\rm PL}$	$0.171^{+0.008}_{-0.010}$	Tied
GAUSS	$E_{\rm line}$ (keV)	$7.08^{+0.07}_{-0.01}$	Tied
GAUSS	$\sigma_{\rm line}$ (keV)	$1.0^{+0}_{-0.2}$	Tied
GAUSS	$N_{\rm line}$	$4.67^{+0.05}_{-0.64} \times 10^{-3}$	Tied

Table 4.4: Values of the best-fitting parameterisation for model 2. Mean values are listed for the untied parameters, indicated by †. Errors on the tied parameters represent the 90% confidence interval as computed by the XSPEC MCMC error routine. The NTHCOMP blackbody temperature was tied to the BBODYRAD blackbody temperature. XSPEC gives a fit statistic of $\chi^2 = 1983.8$ for 1059 d.o.f.

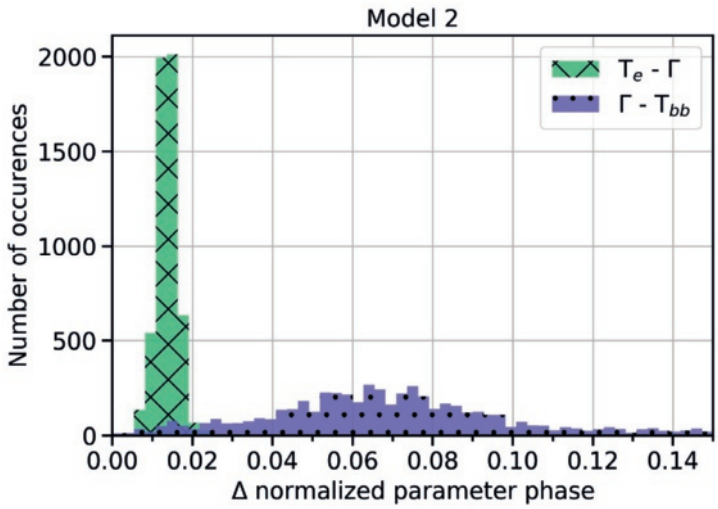

Figure 4.13: Phase differences of the parameter variations of untied parameters for the best parameterization of model 2, computed for each bootstrap sample. The phase offset between $T_{\rm e}$ and $T_{\rm bb}$ is the sum of the offset between $T_{\rm e}$ and Γ and between Γ and $T_{\rm bb}$.

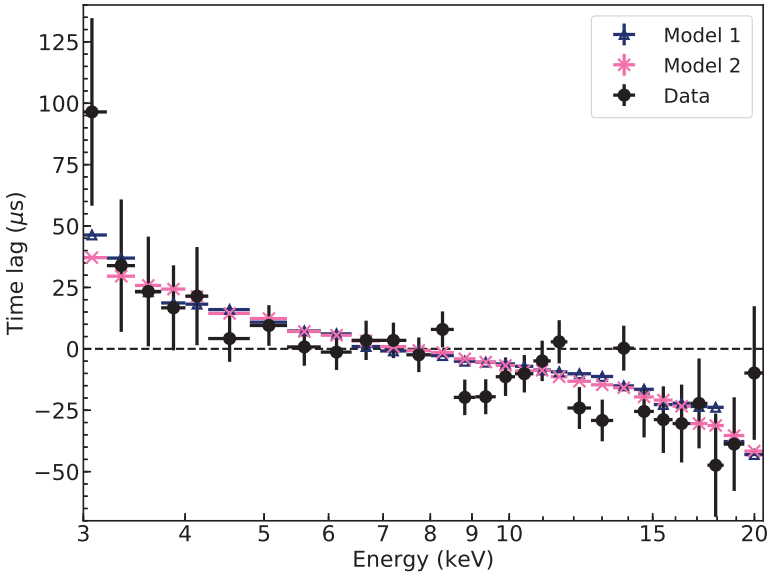

Figure 4.14: Lag-energy spectra from the cross-spectra averaged over 825–845 Hz. The black dots are the data (same as in Figure 4.4). The black dashed line indicates no time lag with respect to the variability in the reference band. The navy blue triangles are simulated from the model 1 ($\chi^2 = 32.4$), the purple X's are simulated from model 2 ($\chi^2 = 29.9$). All model-data comparisons have 29 degrees of freedom. Both models adequately reproduce the amplitude and slope of the time lags.

4.5 Discussion

We have been able to use the CCF method to reconstruct the phase-dependent spectral variability associated with the lower kHz QPO, in observations of 4U 1608–52 obtained in March 1996. The chosen models, dominated by a constant disc blackbody and variable Comptonized component, were able to fit the time-averaged spectrum well. We have shown that a combination of three sinusoidally varying Comptonization parameters (the seed blackbody temperature $T_{\rm bb}$, the power-law index Γ, and the electron temperature $T_{\rm e}$) can model the phase-resolved spectral variations observed via the CCF method and hence also explain the observed dependence of lag on energy. In this section we first consider the observed spectral variability in empirical terms for comparison with the previously known spectral-timing properties of the QPO, before considering any physical interpretation. Then we consider the physical interpretation based on our fits, but we note that the assumed NTHCOMP Comptonization model is simply one parameterization of the variable continuum spectrum, so we should be careful not to assign too much physical significance to the parameter variations we observe.

Our results show that the lower kHz QPO variations correspond to significant changes not only in the flux but also the spectral shape of the Comptonized emission. This finding can shed light on a number of aspects of the QPO variability. Firstly, as can be seen in Figure 4.8, the Comptonized component is only weakly variable at low energies, appearing to pivot so that higher energies show intrinsically larger variations. This finding implies that the well-known drop in lower kHz QPO fractional rms towards low energies (e.g., originally discovered by Berger et al. 1996 for the same 4U 1608–52 data) is not simply due to dilution by the constant disc blackbody emission, but also due to intrinsic spectral evolution of the Comptonized component, which has an intrinsically harder variable spectrum than the time-averaged shape of the Comptonized component. This result could potentially explain the harder shape (higher inferred seed photon temperature) of the covariance spectrum of the QPO compared to the time-averaged Comptonized component shape, which has been noted by Peille et al. (2015) and Troyer & Cackett (2017). Those works suggested that the difference might originate if the QPO emission is localised to a hotter part of the boundary layer, but our new findings suggest that instead, the harder covariance spectrum results from spectral evolution during the QPO cycle.

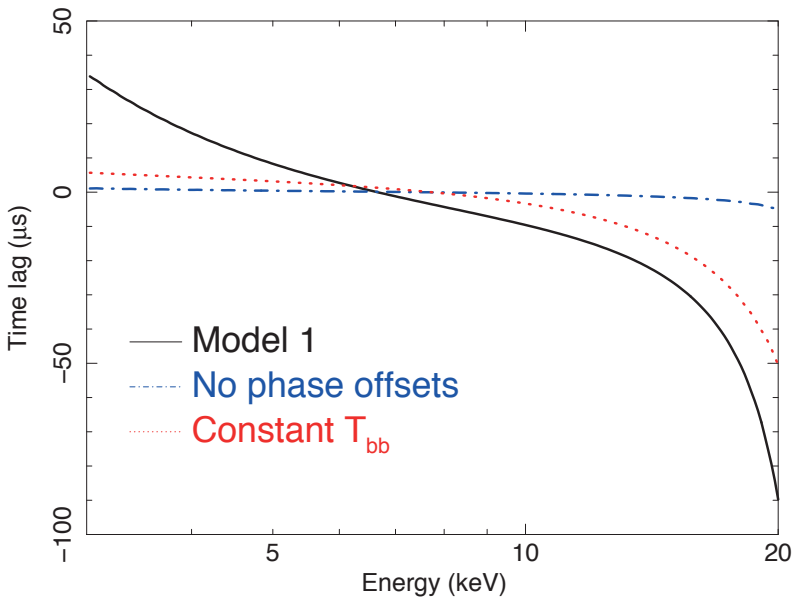

Figure 4.15: Lag-energy model curves for spectral variability models based on model 1, assuming: best-fitting model 1 parameter variations shown in Figure 4.7 (black solid line); best-fitting model 1 variations but with zero phase offset between parameter variations (blue dot-dashed line); best-fitting model 1 variations but with $T_{\rm bb}$ set to be constant (red dotted line). See text for details.

It is important to bear in mind that the observed lags do not automatically arise

from significant spectral evolution during the QPO cycle. To examine the origin of the lags in more depth, we used the best-fitting Γ, $T_{\rm bb}$ and $T_{\rm e}$ parameter curves for model 1 (Figure 4.7) to generate model flux spectra in XSPEC and hence produce flux light curves, which we could use to predict the hypothetical lag-energy dependence in the absence of any instrumental effects. The lag-energy dependence for the best-fitting model 1 parameters is shown as a solid black line in Figure 4.15. It broadly matches the dependence predicted and observed in the PCA data shown in Figure 4.14, albeit with a sharper drop close to 20 keV, since the lags there are not contaminated by redistributed photons from lower energies as is the case for the observed and simulated PCA measurements. The best-fitting parameter variations show significant phase offsets between all the variable parameters. These offsets can be removed in our simulations, by aligning the parameter variation curves to all have the same phase at curve maximum, the result of which is shown in Figure 4.15 as the dot-dashed blue line. The lags are significantly suppressed as a result, showing that the parameter evolution on its own does not lead to the observed lags if those parameter variations are in phase: the parameters which describe the Compton spectral shape must intrinsically lag one another, in the order $T_{\rm bb}$ varying first, then Γ and finally $T_{\rm e}$. The effect of forcing $T_{\rm bb}$ to be constant is shown as the dotted red line in Figure 4.15. Clearly the variation in $T_{\rm bb}$ plays an important role in setting the shape of the lag-energy spectrum below ~ 7 keV.

In our parameterization of the QPO spectral evolution, the small phase offsets between the spectral parameter variations are responsible for producing the observed energy-dependent lags. To understand this effect better, first consider that the peak in the QPO flux corresponds to phase zero in the parameter evolution curves (Figures 4.7 and 4.12). From this point, Γ rises and the spectrum softens significantly. At the same time, $T_{\rm bb}$ also increases. Since the normalisation at 1 keV is fixed, an increase in $T_{\rm bb}$ corresponds to an increase in the energy of the low-energy roll-over in the spectrum, which produces a relative hardening of the spectrum. Since the peak in $T_{\rm bb}$ is seen *before* the peak in softening due to Γ, we see an overall soft lag in the lag-energy spectrum. At higher energies where the high-energy cut-off dominates, the spectral shape is dominated by the roughly exponential cut-off curve and not the changes in photon index. However, the correlation of $T_{\rm e}$ with Γ means that the soft lags due to varying Γ 'leak' into the higher energies as $T_{\rm e}$ increases (and an additional delay is introduced since $T_{\rm e}$ itself slightly lags Γ).

We now turn to the physical interpretation of our results. The main spectral-variability feature of the QPO, independent of the lags, is the large intrinsic variation in spectral shape, which corresponds to a low-energy pivoting throughout the QPO cycle, such that the peak of the flux cycle corresponds to the hardest photon index and the spectrum is softest at the minimum. The weak intrinsic variability at low energies suggests that seed variations are rather weak and we also do not see any clear evidence for a direct component of surface blackbody emission. These spectral results seem to

rule out variations in seed blackbody luminosity, or the upscattered fraction of the surface blackbody luminosity, as the driving signal for the QPO. Instead, modulation of the heating rate in the Comptonizing medium seems to be the most likely cause of the intrinsically hard flux variability of the QPO. It is interesting to note that Kumar & Misra (2014) reach similar conclusions, on the basis of Comptonization modelling of the observed energy dependence of QPO fractional rms and time lags.

An increase in heating rate relative to seed photon luminosity in a Comptonizing medium can explain the hardening of the spectrum coupled with increasing flux (e.g., Pietrini & Krolik 1995), but at the same time, the inferred electron temperature is seen to drop, which would be inconsistent with an increase in relative heating rate. This result can be understood if the inferred high energy cut-off is due to Compton down-scattering and hence more closely related to the optical depth, τ, of the Comptonizing region than to its intrinsic temperature (following $E_{\rm cut} \simeq 511\ {\rm keV}/\tau^2$, Sunyaev & Titarchuk 1980), which may be the case for high-optical depth plasma in the boundary layer. That is, an increase in heating rate is coupled to an increase in optical depth and hence a lower cut-off energy. Such a mechanism might be linked to variations in accretion rate from the disc, but no evidence for disc variability is seen in our data and, as noted previously, it is difficult to understand how the accretion rate could be modulated so coherently.

An alternative possibility is that the QPO represents some kind of "breathing" oscillation excited in the boundary layer. For example, following the possible connection of heating rate to optical depth described above, let us assume that the heating rate is coupled to the density in the boundary layer. A perturbation in density increases the heating rate but also causes the expansion of the boundary layer, which is expected to be radiatively supported (Inogamov & Sunyaev 1999). Density and heating rate then fall, the boundary layer collapses again to its denser configuration, and the cycle repeats. This behaviour could explain the correlated variation of Γ and $T_{\rm e}$ in terms of the heating and density (and hence optical depth) correlation. The connection to seed photon temperature ($T_{\rm bb}$) is less clear, but might be linked to stronger irradiation of the neutron star surface occurring when the boundary layer is less dense (thus allowing more Comptonized photons to reach that surface). The observed phase offsets between the Comptonization parameters are small compared to the oscillation time and so could be inherent to the oscillation mechanism. One could also consider Comptonization delays as a cause of the lags, as explored by Kumar & Misra (2014, 2016). Besides Comptonization delays, causal limits could also play an important role: a phase offset of 0.02 corresponds to only 7 km in light-travel distance. For a moderately optically thick region of $\sim$km scale, fluctuations in density and optical depth in the core of the emitting region could take a time comparable to the observed phase offsets to be communicated to the observer.

Finally, we note that the lack of evidence for iron line variability during the lower kHz QPO cycle is consistent with the lack of clear evidence for a reverberation signal

for the QPO (Peille et al. 2015; Cackett 2016). The implication is that the boundary layer oscillations seen by the observer are not seen (at least with the same amplitude) by the disc, which could happen if the oscillating boundary layer emission occurs at low latitudes relative to the neutron star surface. Some hints of disc reverberation are seen for the upper kHz QPO however. A natural extension of this work would be to carry out a phase-resolved study of the upper kHz QPO, to identify the similarities and differences with the lower kHz QPO.

4.6 Conclusion

We have presented the first application of phase-resolved spectroscopy to a kHz QPO. For this investigation, we applied the Stevens & Uttley (2016) QPO-phase-resolved technique to the lower kHz QPO observed by *RXTE* in 4U 1608–52 during its 1996 outburst. As this observation was taken in the early months of *RXTE* operation, all 5 PCUs were operational and there were minimal calibration differences between the different PCUs. The lower kHz QPO is the only type of periodic or quasi-periodic variability seen in the observation we use (that is, there were no upper kHz QPOs or coherent pulsations), which made this clean signal an ideal proof-of-principle case.

The phase spectra were fitted with models that contained a multi-colour blackbody accretion disc, a thermal Comptonization component for the boundary layer, a single-temperature seed blackbody that could represent the surface of the neutron star, and a broad iron line to model the most prominent reflection feature. We found that the power-law index Γ, high-energy cut-off temperature $T_{\rm e}$, and seed blackbody temperature $T_{\rm bb}$ varied with QPO phase: the fractional rms variation in Γ was $\sim$9–10%, in $T_{\rm e}$ was $\sim$8–9%, and in $T_{\rm bb}$ was $\sim$1.5–2%. There were non-zero phase differences between all three parameter variations, notably a $\sim$0.014 relative phase lag of $T_{\rm e}$ compared with Γ, and a $\sim$0.02–0.05 relative phase lag of Γ compared with $T_{\rm bb}$. With simulations of light curves fed by the sinusoidal best-fitting spectral parameter variations, we were able to reproduce the lag-energy spectrum of the data in amplitude and in slope/shape.

Since the shape of the Comptonization spectral component is changing and the parameters do not vary in-phase with one-another, we rule out upscattering of variations in the surface blackbody luminosity or accretion disc luminosity as the origin of the QPO. Instead, our results point towards a "breathing" oscillation in the boundary layer propelled by cyclic changes in the density and heating rate of the Comptonizing region, with optical depth changes then changing the amount of irradiation of the neutron star surface which produces seed photons.

Further exploration of this possible "breathing" oscillation mechanism in the boundary layer can be done in finer detail (both temporally and spectrally) with two current X-ray instruments, *AstroSat* and *NICER*. *AstroSat*'s Large-Area X-ray Proportional

Counter (LAXPC) does not have the same telemetry-constrained data modes as with *RXTE*, so it has an energy resolution of 10% across 3–80 keV with 10 μs time resolution (Singh et al. 2014; Yadav et al. 2016; Antia et al. 2017). Verdhan Chauhan et al. (2017) reported the detection of kHz QPOs in 4U 1728–34 with the LAXPC. *AstroSat*'s order-of-magnitude improvement in both energy and time resolution compared to *RXTE* event-mode data will allow more precise modelling of phase-dependent spectral changes. In the softer X-rays, *NICER* avoids telemetry constraints while also minimizing deadtime. It has 40 ns time resolution and CCD energy resolution across its 0.5–12 keV bandpass (Gendreau et al. 2012; Arzoumanian et al. 2014). These specifications make *NICER* well-suited for determining whether blackbody temperature variations or low-energy normalization variations are responsible for the soft X-ray changes in the kHz QPO (as discussed in the previous section). Looking into the future of X-ray observatories, *eXTP* (Zhang et al. 2016) and *STROBE-X* (Wilson-Hodge et al. 2017) will both provide exceptional opportunities to study spectral-timing behaviour of kHz QPOs.

Acknowledgements

A.L.S. acknowledges support from NOVA (Nederlandse Onderzoekschool voor Astronomie), and thanks Ed Cackett for useful discussions. D.A. acknowledges support from the Royal Society. The idea for this paper grew out of discussions at 'The X-ray Spectral-Timing Revolution' Lorentz Center workshop (February 2016).

This research made use of data and software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC); NASA's Astrophysics Data System Bibliographic Services; NumPy v1.13.1 and Scipy v0.19.1 (Jones et al. 2001); Astropy v2.0 (Astropy Collaboration et al. 2013); Matplotlib v2.0.2 (Hunter 2007); and Jupyter Notebook v5.0.0 with iPython v6.1.0 (Perez & Granger 2007).

Appendix 4A: Software

The software and data products will be made publicly available when the paper is published. The data were reduced and processed with the scripts in the GitHub repository abigailStev/rxte_reduction. The reduced and processed data products, analysis scripts, and explanatory Jupyter notebooks for this paper will be found in the GitHub repository abigailStev/kHz_QPO_phase_res.

Appendix 4B: Spectral Fitting Procedure

As described in the text, we follow the systematic procedure from Stevens & Uttley (2016) to determine which spectral parameters must vary to reproduce the phase-

dependent changes in the energy spectrum. The model IDs for each parameterisation, $x.y.z$, are to be interpreted as model number x, with y free parameters, enumerated by z.

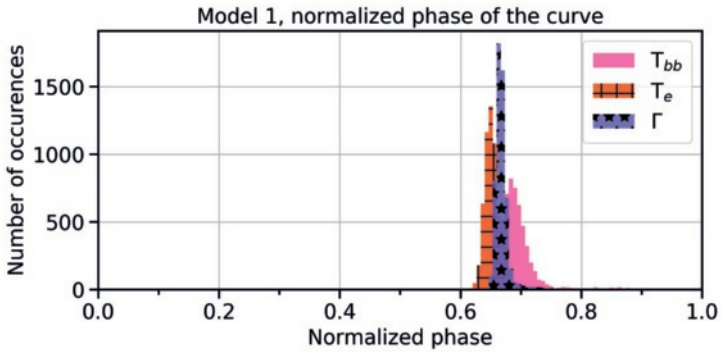

Figure 4.16: Model 1: Histograms of the phases of the fitted curves for each bootstrap sample. Measured from the fit to the data, $T_{\rm bb}$ has a phase of 0.689 (1σ confidence interval from the bootstrap sample, centred on the median value: [0.672–0.711]), $T_{\rm e}$ has a phase of 0.652 [0.640–0.661], and Γ has a phase of 0.667 [0.660–0.674].

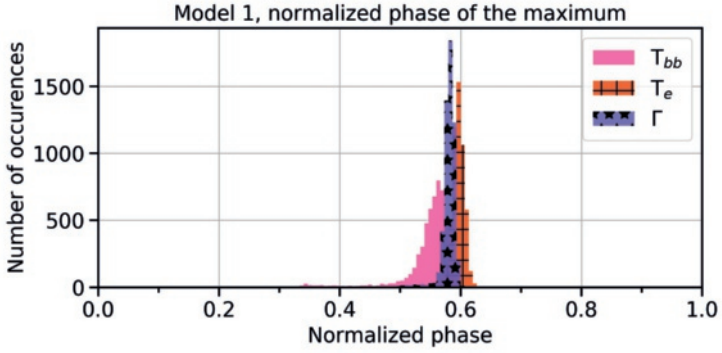

Figure 4.17: Model 1: Histograms of the normalized phases of the maximum of the fitted curves for each bootstrap sample. Measured from the fit to the data, the maximum in $T_{\rm bb}$ has a phase of 0.554 (1σ confidence interval from the bootstrap sample, centred on the median value: [0.529–0.573]), the maximum in $T_{\rm e}$ has a phase of 0.593 [0.586–0.606], and the maximum in Γ has a phase of 0.581 [0.574–0.589].

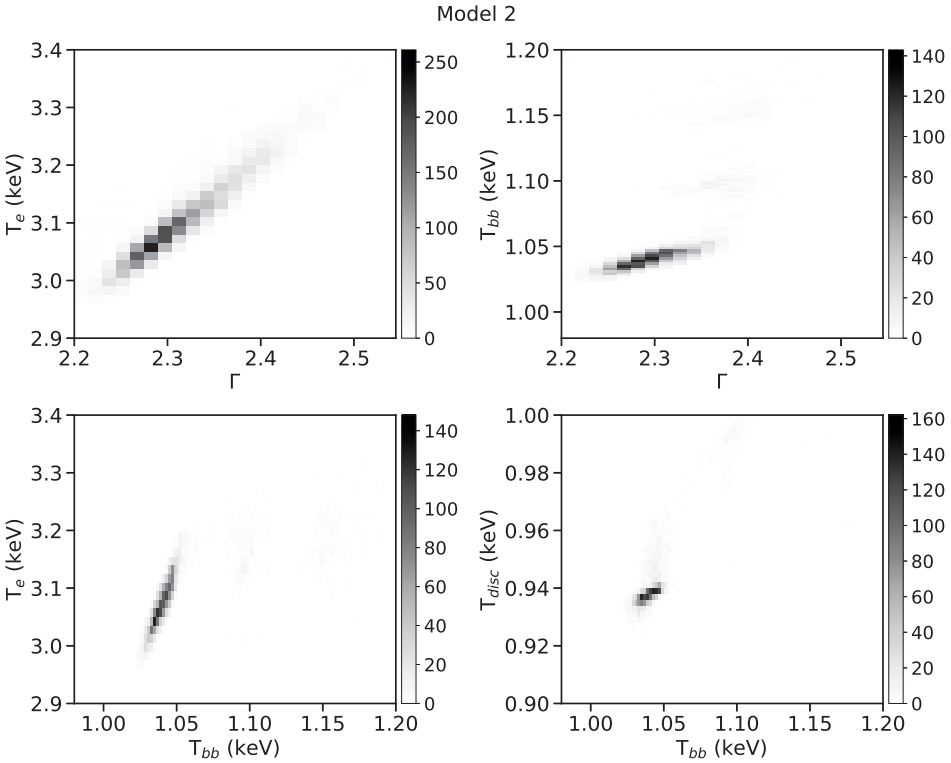

Figure 4.18: Model 2: 2-D histograms of spectral parameter values in phase bin 5 for each bootstrap sample, showing the correlations between spectral parameters. The colour bars indicate the number of occurrences, out of the 5419 good bootstrap samples for model 2. As in model 1, $T_{\rm disc}$ was tied across phase, but here we do not see a strong correlation with it and $T_{\rm bb}$.

Table 4.5: Systematic fitting for model 1 (TBNEW_GAS × (DISKBB+NTHCOMP+GAUSS)). The best parameterization, 1.3.3, has T_{bb}, T_e, and Γ untied, and gives a χ^2 / d.o.f. = 1981.2/1060. Any improvement on that, in the final block of fits 1.4.* with four untied parameters, does not give $\Delta\chi^2 > \Delta$ d.o.f. compared to the best-fit in 1.3.* with three untied parameters.

ID	χ^2 / d.o.f.	Untied
1.0.1	105150.7/1180	-
1.1.1	75497.1/1140	T_{disc}
1.1.2	87690.9/1140	N_{disc}
1.1.3	7126.5/1140	Γ
1.1.4	43674.8/1140	T_e
1.1.5	3854.0/1140	T_{bb}
1.1.6	5021.3/1140	N_{PL}
1.1.7	90572.7/1140	E_{line}
1.1.8	105136.3/1140	σ_{line}
1.1.9	68453.8/1140	N_{line}
1.2.1	3191.1/1100	T_{bb}, T_{disc}
1.2.2	3244.7/1100	T_{bb}, N_{disc}
1.2.3	2233.0/1100	T_{bb}, Γ
1.2.4	2062.9/1100	T_{bb}, T_e
1.2.5	3179.0/1100	T_{bb}, N_{PL}
1.2.6	3609.0/1100	T_{bb}, E_{line}
1.2.7	3832.7/1100	T_{bb}, σ_{line}
1.2.8	3534.3/1100	T_{bb}, N_{line}
1.2.9	3219.5/1100	N_{PL}, T_{disc}
1.2.10	3264.3/1100	N_{PL}, N_{disc}
1.2.11	4997.0/1100	N_{PL}, Γ
1.2.12	4572.9/1100	N_{PL}, T_e
1.2.13	4519.6/1100	N_{PL}, E_{line}
1.2.14	5019.2/1100	N_{PL}, σ_{line}
1.2.15	4794.0/1100	N_{PL}, N_{line}
1.2.16	2063.6/1100	T_e, Γ
1.3.1	2009.3/1060	T_{bb}, T_e, T_{disc}
1.3.2	2008.6/1060	T_{bb}, T_e, N_{disc}
1.3.3	1981.2/1060	T_{bb}, T_e, Γ
1.3.4	2009.3/1060	T_{bb}, T_e, N_{PL}
1.3.5	2043.7/1060	T_{bb}, T_e, E_{line}
1.3.6	2060.1/1060	T_{bb}, T_e, σ_{line}
1.3.7	2059.0/1060	T_{bb}, T_e, N_{line}
1.3.8	2101.9/1060	T_{bb}, Γ, T_{disc}

ID	χ^2 / d.o.f.	Untied
1.3.9	2114.9/1060	T_{bb}, Γ, N_{disc}
1.3.10	2084.5/1060	T_{bb}, Γ, N_{PL}
1.3.11	2200.0/1060	T_{bb}, Γ, E_{line}
1.3.12	2235.6/1060	T_{bb}, Γ, σ_{line}
1.3.13	2208.1/1060	T_{bb}, Γ, N_{line}
1.3.14	1984.7/1060	T_e, Γ, T_{disc}
1.3.15	1982.6/1060	T_e, Γ, N_{disc}
1.3.16	1983.1/1060	T_e, Γ, N_{PL}
1.3.17	2045.3/1060	T_e, Γ, E_{line}
1.3.18	2059.5/1060	T_e, Γ, σ_{line}
1.3.19	2047.9/1060	T_e, Γ, N_{line}
1.4.1	1978.5/1020	T_{bb}, T_e, Γ, T_{disc}
1.4.2	1975.5/1020	T_{bb}, T_e, Γ, N_{disc}
1.4.3	1979.5/1020	T_{bb}, T_e, Γ, N_{PL}
1.4.4	1978.6/1020	T_{bb}, T_e, Γ, E_{line}
1.4.5	1973.4/1020	T_{bb}, T_e, Γ, σ_{line}
1.4.6	1960.1/1020	T_{bb}, T_e, Γ, N_{line}
1.4.7	1999.6/1020	T_{bb}, T_e, N_{disc}, T_{disc}
1.4.8	1996.9/1020	T_{bb}, T_e,, N_{disc} N_{PL}
1.4.9	2001.3/1020	T_{bb}, T_e, N_{disc}, E_{line}
1.4.10	2003.8/1020	T_{bb}, T_e, N_{disc}, σ_{line}
1.4.11	1998.5/1020	T_{bb}, T_e, N_{disc}, N_{line}
1.4.12	1999.1/1020	T_{bb}, T_e, T_{disc}, N_{PL}
1.4.13	2003.3/1020	T_{bb}, T_e, T_{disc}, E_{line}
1.4.14	2005.4/1020	T_{bb}, T_e, T_{disc}, σ_{line}
1.4.15	2001.3/1020	T_{bb}, T_e, T_{disc}, N_{line}
1.4.16	2005.1/1020	T_{bb}, T_e, N_{PL}, E_{line}
1.4.17	2006.9/1020	T_{bb}, T_e, N_{PL}, σ_{line}
1.4.18	2005.8/1020	T_{bb}, T_e, N_{PL}, N_{line}
1.4.19	1971.7/1020	T_e, Γ, N_{disc}, T_{disc}
1.4.20	1981.5/1020	T_e, Γ, N_{disc}, N_{PL}
1.4.21	1976.3/1020	T_e, Γ, N_{disc}, E_{line}
1.4.22	1975.1/1020	T_e, Γ, N_{disc}, σ_{line}
1.4.23	1972.5/1020	T_e, Γ, N_{disc}, N_{line}
1.4.24	1982.4/1020	T_e, Γ, N_{PL}, T_{disc}
1.4.25	1979.3/1020	T_e, Γ, N_{PL}, E_{line}
1.4.26	1976.2/1020	T_e, Γ, N_{PL}, σ_{line}
1.4.27	1971.4/1020	T_e, Γ, N_{PL}, N_{line}

Table 4.6: Systematic fitting for model 2 (TBNEW_GAS × (DISKBB+BBODYRAD+NTHCOMP+GAUSS)). The best parameterisation, 2.3.4, has T_{bb}, Γ, and T_e varying. This fit gives a χ^2 / d.o.f. = 1983.8/1059. Any improvement on that fit, in the final block of fits with four untied parameters, does not give $\Delta\chi^2 > \Delta$ d.o.f. compared to the best-fit with three untied parameters.

ID	χ^2 / d.o.f.	Untied
2.0.1	105157.6/1179	-
2.1.1	75299.4/1139	T_{disc}
2.1.2	87479.2/1139	N_{disc}
2.1.3	3154.0/1139	T_{bb}
2.1.4	4427.1/1139	N_{bb}
2.1.5	7025.8/1139	Γ
2.1.6	42796.6/1139	T_e
2.1.7	4897.7/1139	N_{PL}
2.1.8	90375.5/1139	E_{line}
2.1.9	105138.2/1139	σ_{line}
2.1.10	65550.2/1139	N_{line}
2.2.1	2637.8/1099	T_{bb}, T_{disc}
2.2.2	2646.5/1099	T_{bb}, N_{disc}
2.2.3	2574.8/1099	T_{bb}, N_{bb}
2.2.4	2296.1/1099	T_{bb}, Γ
2.2.5	2405.1/1099	T_{bb}, T_e
2.2.6	2274.8/1099	T_{bb}, N_{PL}
2.2.7	3064.1/1099	T_{bb}, E_{line}
2.2.8	3115.5/1099	T_{bb}, σ_{line}
2.2.9	2958.8/1099	T_{bb}, N_{line}
2.2.10	2851.0/1099	N_{bb}, T_{disc}
2.2.11	3410.6/1099	N_{bb}, N_{disc}
2.2.12	6059.6/1099	N_{bb}, Γ
2.2.13	5002.3/1099	N_{bb}, T_e
2.2.14	3265.5/1099	N_{bb}, N_{PL}
2.2.15	4161.7/1099	N_{bb}, E_{line}
2.2.16	5365.8/1099	N_{bb}, σ_{line}
2.2.17	4288.5/1099	N_{bb}, N_{line}
2.3.1	2056.8/1059	T_{bb}, Γ, T_{disc}
2.3.2	2211.7/1059	T_{bb}, Γ, N_{disc}
2.3.3	2062.5/1059	T_{bb}, Γ, N_{bb}
2.3.4	1983.8/1059	T_{bb}, Γ, T_e
2.3.5	2062.0/1059	T_{bb}, Γ, N_{PL}
2.3.6	2122.5/1059	T_{bb}, Γ, E_{line}

ID	χ^2 / d.o.f.	Untied
2.3.7	2146.0/1059	T_{bb}, Γ, σ_{line}
2.3.8	2991.8/1059	T_{bb}, Γ, N_{line}
2.3.9	2207.7/1059	T_{bb}, T_e, T_{disc}
2.3.10	2377.1/1059	T_{bb}, T_e, N_{disc}
2.3.11	2109.7/1059	T_{bb}, T_e, N_{bb}
2.3.12	2199.0/1059	T_{bb}, T_e, N_{PL}
2.3.13	2042.1/1059	T_{bb}, T_e, E_{line}
2.3.14	2057.8/1059	T_{bb}, T_e, σ_{line}
2.3.15	2055.1/1059	T_{bb}, T_e, N_{line}
2.3.16	2092.8/1059	T_{bb}, N_{PL}, T_{disc}
2.3.17	2108.7/1059	T_{bb}, N_{PL}, N_{disc}
2.3.18	2066.1/1059	T_{bb}, N_{PL}, N_{bb}
2.3.19	2222.4/1059	T_{bb}, N_{PL}, E_{line}
2.3.20	2258.5/1059	T_{bb}, N_{PL}, σ_{line}
2.3.21	2245.4/1059	T_{bb}, N_{PL}, N_{line}
2.4.1	1976.9/1019	T_{bb}, Γ, T_e, T_{disc}
2.4.2	1972.2/1019	T_{bb}, Γ, T_e, N_{disc}
2.4.3	1983.4/1019	T_{bb}, Γ, T_e, N_{bb}
2.4.4	1980.5/1019	T_{bb}, Γ, T_e, N_{PL}
2.4.5	1990.6/1019	T_{bb}, Γ, T_e, E_{line}
2.4.6	1991.1/1019	T_{bb}, Γ, T_e, σ_{line}
2.4.7	1992.2/1019	T_{bb}, Γ, T_e, N_{line}
2.4.8	2362.6/1019	T_{bb}, T_e, E_{line}, T_{disc}
2.4.9	2001.3/1019	T_{bb}, T_e, E_{line}, N_{disc}
2.4.10	2094.0/1019	T_{bb}, T_e, E_{line}, N_{bb}
2.4.11	2005.1/1019	T_{bb}, T_e, E_{line}, N_{PL}
2.4.12	2339.1/1019	T_{bb}, T_e, E_{line}, σ_{line}
2.4.13	2033.0/1019	T_{bb}, T_e, E_{line}, N_{line}

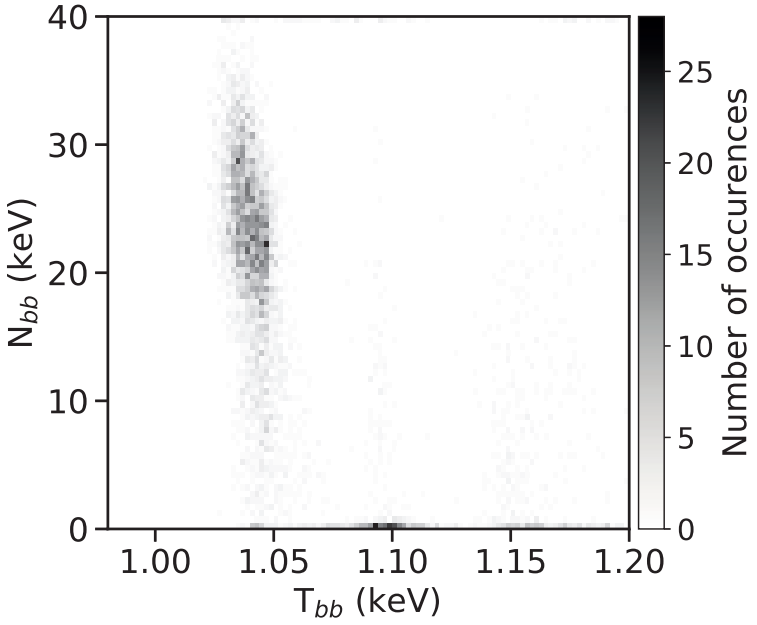

Figure 4.19: Model 2: 2-D histogram of the $T_{\rm bb}$ value in phase bin 5 with the tied $N_{\rm bb}$ value for each bootstrap sample. The colour bar indicates the number of occurrences, out of the 5419 good bootstrap samples for model 2. As in model 1, $T_{\rm disc}$ was tied across phase, but here we do not see a strong correlation with it and $T_{\rm bb}$.

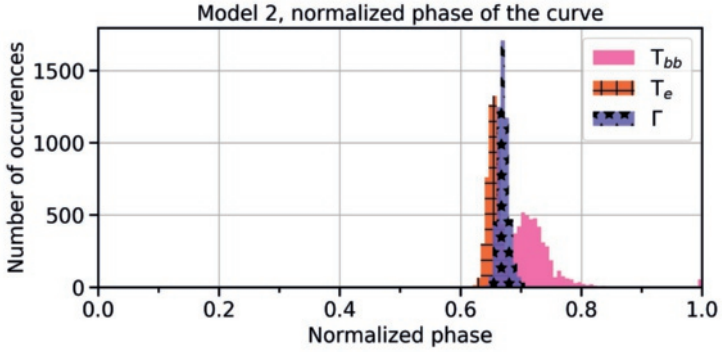

Figure 4.20: Model 2: Histograms of the phases of the fitted curves for each bootstrap sample. As measured from the fit to the data, the maximum in $T_{\rm bb}$ has a phase of 0.695 (1σ confidence interval from the bootstrap sample, centred on the median value: [0.687–0.745]), the maximum in $T_{\rm e}$ has a phase of 0.652 [0.645–0.665], and the maximum in Γ has a phase of 0.668 [0.662–0.679].

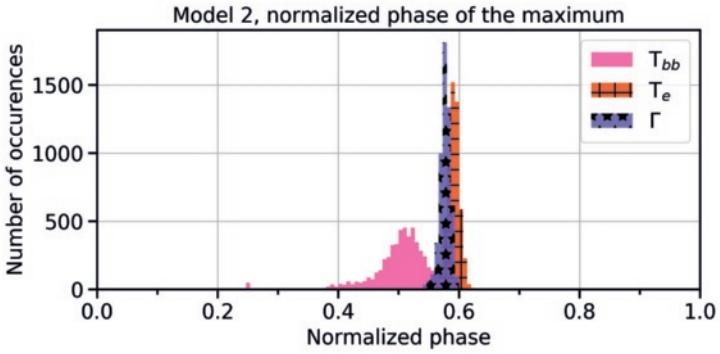

Figure 4.21: Model 2: Histograms of the normalized phases of the maximum of the fitted curves for each bootstrap sample. As measured from the fit to the data, the maximum in $T_{\rm bb}$ has a phase of 0.545 (1σ confidence interval from the bootstrap sample, centred on the median value: [0.471–0.544]), the maximum in $T_{\rm e}$ has a phase of 0.593 [0.581–0.600], and the maximum in Γ has a phase of 0.595 [0.569–0.586].

5

NICER Spectral-Timing Studies of a Low-Frequency QPO in the Soft-Intermediate State of MAXI J1535–571

A. L. Stevens and P. Uttley

To be submitted to The Astrophysical Journal

Abstract

We present the discovery and preliminary spectral-timing analysis of a low-frequency ~ 5.7 Hz QPO feature in observations of MAXI J1535–571 in the soft-intermediate state, obtained in September–October 2017 by the *Neutron Star Interior Composition Explorer* (*NICER*). The feature is relatively broad (quality factor $Q \sim 2$) and weak (1.8% fractional rms from 3–10 keV), with a weak harmonic and additional low-amplitude broadband noise. These characteristics likely make it a weak Type B QPO, like one previously identified in the soft-intermediate state of the transient black hole X-ray binary XTE J1550–564. The lag-energy spectrum of the QPO shows soft lags of about 50 ms with an inflected spectral shape, flattening above ~ 4 keV. Preliminary QPO-phase-resolved spectral analysis suggests that the inflected lag-energy spectrum is due to the disk blackbody component lagging the Comptonized power-law emission by $\sim 27\%$ of a QPO cycle. This large phase shift has similar amplitude but opposite

sign to that seen in *Rossi X-ray Timing Explorer* data for the Type B QPO from the transient black hole X-ray binary GX 339–4. We suggest that the Type B QPOs from these systems may have the same origin, in the form of a precessing jet-like corona illuminating the disk, with the different lag signs depending on the observer inclination angle.

5.1 Introduction

MAXI J1535–571 is a newly-discovered transient X-ray binary that went into outburst starting on 2 September 2017. It was co-discovered in the X-rays by *MAXI* (Negoro et al. 2017b) and the *Neil Gehrels Swift Observatory* (Markwardt et al. 2017; Kennea et al. 2017). Its multi-wavelength properties strongly suggest that the source is an accreting stellar-mass black hole (Scaringi & ASTR211 Students 2017a,b; Negoro et al. 2017a; Russell et al. 2017b; Dincer 2017; Britt et al. 2017). One week following its discovery, X-ray, radio, and sub-mm detections showed the source brightening and entering the hard-intermediate spectral state (Nakahira et al. 2017; Kennea 2017; Palmer et al. 2017; Tetarenko et al. 2017; Shidatsu et al. 2017b). In the hard-intermediate state, low-frequency quasi-periodic oscillations (LF QPOs) were detected in the X-ray light curve by *Swift* (Mereminskiy & Grebenev 2017; Russell et al. 2017a) and *NICER* (Gendreau et al. 2017). The source continued to transition into the soft state in October and November (Shidatsu et al. 2017a). Although heavily absorbed (neutral column density $N_{\rm H} \sim$ few$\times 10^{22}$ cm^{-2}), MAXI J1535–571 reached extremely bright flux levels, up to 5 Crab in 2–20 keV flux (Shidatsu et al. 2017b). As of mid-January 2018, the source was beginning to transition back through the intermediate states to the hard state.[1]

LF QPOs are seen at Fourier frequencies of $\sim$0.1 to tens of Hz in the X-ray light curves of accreting black hole X-ray binaries. Theories of the origins of LF QPO variability fall under the general categories of intrinsic variability, such as shocks in the accretion flow (Chakrabarti 1996), or geometric variability, like rigid-body Lense-Thirring precession of the Comptonizing region (Stella & Vietri 1998; Ingram et al. 2009). Systematic observational studies indicate that the LF QPO origin is geometric in nature, due to the correlation of the QPO amplitude (Schnittman et al. 2006; Motta et al. 2015; Heil et al. 2015a) and the sign of the energy-dependent lags (van den Eijnden et al. 2017) with the binary orbit inclination. The two most common types of LF QPOs seen in black hole X-ray binaries are called Type B and Type C. Type B LF QPOs are characterized by weak harmonics and broadband noise, and appear in the soft-intermediate state; Type C LF QPOs are characterized by strong harmonics and broadband noise, and appear in the hard-intermediate state (for more explanation of the types of LF QPOs, see Remillard et al. 2002; Casella et al. 2005; Motta et al. 2011). The two appear at similar frequencies but show different amplitude-inclination correlations, which suggests that they have different origins.

Spectral-timing analysis can determine the amplitudes of and causal relationships between variations of the different emission components seen in the X-ray spectrum (e.g., disk blackbody and Comptonized emission), opening up new avenues with which

[1] For up-to-date information on MAXI J1535–571, see `http://maxi.riken.jp/nakahira/1535monitor/`.

to analyze LF QPOs (e.g., Ingram & van der Klis 2015; Stevens & Uttley 2016). In this paper we present the discovery and initial spectral-timing analysis of a weak Type B QPO from MAXI J1535–571 observed with the *NICER* experiment, recently installed on the International Space Station. *NICER*'s large collecting area, soft X-ray response, CCD-quality energy resolution and high count-rate capability enable unprecedented spectral-timing studies of even weak QPOs, such as the one we report here. In Section 5.2, we give information about the data and data selection criteria. In Section 5.3, we present our results, including the average power spectrum, lag-energy spectrum, energy-dependent cross-correlation function, and a calibration-independent study of the QPO-phase-dependent spectral variability. The discussion and conclusions are presented in Section 5.4.

5.2 Data

The observations used in this paper were taken in September and October 2017 with the X-ray Timing Instrument (XTI) on *NICER* (Gendreau et al. 2012; Arzoumanian et al. 2014). During this time, *NICER* observed MAXI J1535–571 regularly, with many short continuous (few hundred s) exposures spread throughout each day. Each day's observations are packaged into 'segments'. The data used here to identify the Type B QPO corresponds to all useful exposures obtained during this period, which includes 31 segments, covering first the segment IDs 1050360104–1050360120 and then 1130360101–1130360114. Based on the source spectral and timing properties (see below), these observations correspond to the hard-intermediate and soft-intermediate states of accreting black holes. Over the 61.380 ks exposure of these observations combined, *NICER* recorded an extraordinary 936,910,259 photon counts spanning the full energy range of the detector.

In Figure 5.1 we show the coarse flux (count rate), spectral (hardness ratio) and timing (fractional rms) evolution of MAXI J1535–571, obtained from continuous 64 s sections of the *NICER* data. *NICER* XTI has a time resolution of ~ 40 ns, but for our timing studies, we bin the light curves to a time step of 2^{-8} s ($\simeq 3.9$ ms). For our 64 s light curve sections, this gives 16,384 time bins per section.

Detailed study of the power-spectral evolution throughout the observations showed that in the periods with high rms ($> 5\%$) and generally harder flux spectra, the power spectrum is dominated by strong Type C QPOs, with accompanying broadband noise and a strong harmonic (Uttley et al., in prep.). Therefore, we identify those data with the hard-intermediate state (see, e.g., Nowak 1995; Homan et al. 2001; Belloni et al. 2005). Here we focus on the observations showing low rms ($< 4\%$) and relatively soft spectra, which do not show strong variability, but still show larger fractional variability amplitudes than the canonical soft states observed in other transient BH XRBs (e.g., Belloni et al. 2005). We therefore identify these observations with the

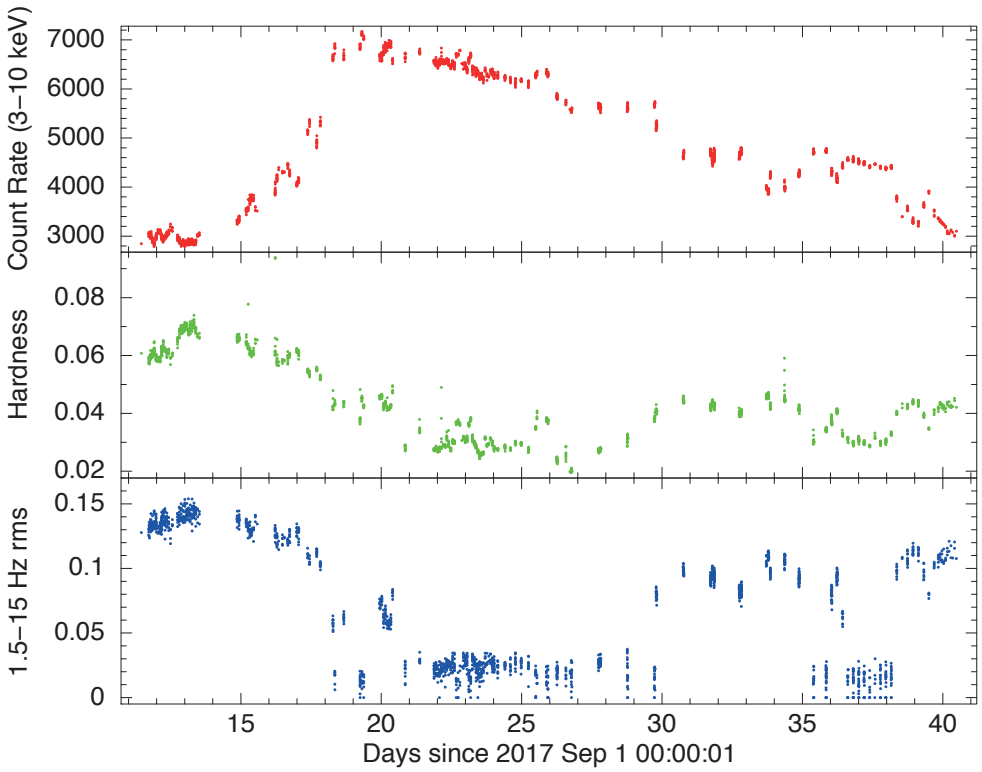

Figure 5.1: Evolution of MAXI J1535–571 coarse spectral and timing properties, for continuous 64 s sections of the data. Top: Total 3–10 keV count rate (count s^{-1}). Middle: Spectral hardness ratio (ratio of the count rates in 7–10 keV to 1–2 keV bands). Bottom: Fractional rms in the 3–10 keV band, integrated from 1.5–15 Hz.

soft-intermediate state.

Since QPOs cannot be readily identified in 64 s sections of the weakly-varying soft-intermediate state light curves, we searched for evidence of characteristic timing signatures by sorting the data according to hardness ratio (where the hardness ratio is the ratio of the 7–10 keV to 1–2 keV count rate). We then averaged together the 3–10 keV power spectra of 64 s sections contained in a sliding range of hardness ratios corresponding to a 30% change in hardness ratio. This analysis revealed the presence of a fairly broad QPO-like signal at a frequency of ~ 5.7 Hz, which is present for hardness ratios < 0.033 but not for harder spectra (see Figure 5.2). The QPO seems to shift in frequency and weaken significantly for the lowest hardness ratios, but appears to be relatively stable over the hardness ratio range 0.022–0.033. Therefore, for the rest of our study we combined the data from 64 s sections, selected in this hardness ratio range and with an integrated 1.5–15 Hz rms < 0.04.

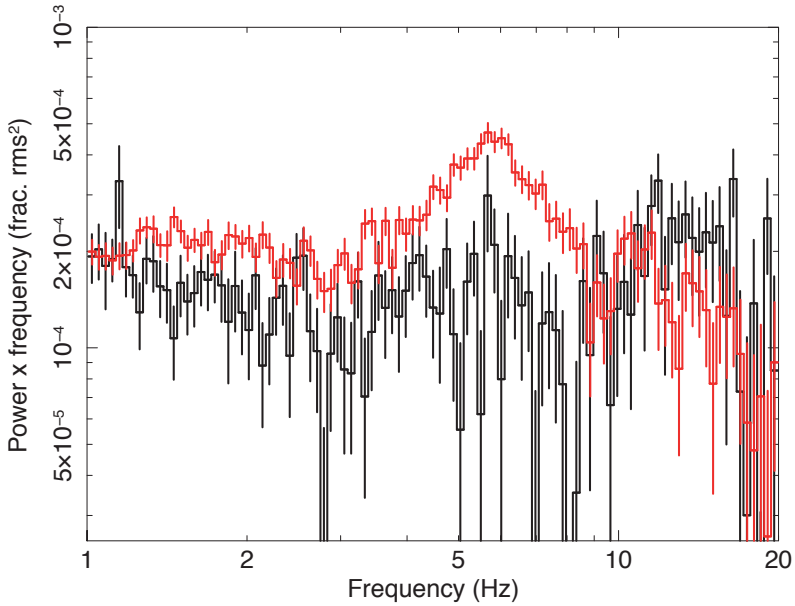

Figure 5.2: Poisson-noise-subtracted 3–10 keV power spectra averaged from 64 s sections of data with 1.5–15 Hz rms < 0.04, for two hardness ratio ranges: 0.0332–0.0432 (black) and 0.0256–0.0332 (red). A broad QPO-like signal emerges for the softer hardness ratio range.

There are 576 good 64 s continuous sections of data which fit our spectral hardness and rms criteria, which gives a total exposure of 36.864 ks. The total count rate in the 3–10 keV energy band is 5816 count s^{-1}. For the reference energy band (used in Sections 5.3.2 and 5.3.3), which covers 3–10 keV but uses detector chains 4–6, the count rate is 2542 count s^{-1}.

5.3 Results

5.3.1 Power Spectrum

We computed the average power spectrum (Figure 5.3) in the energy band 3–10 keV for all XTI detector chains using our own code[2] (see van der Klis 1989 for an overview of power spectra). The power spectrum was geometrically re-binned in frequency with a binning factor of 1.03. Using XSPEC v12.9.1 (Arnaud 1996), we fitted this power spectrum over the 0.03–128 Hz range with four Lorentzians: two of the Lorentzians model the broadband noise, one Lorentzian models the QPO fundamental, and one Lorentzian models the weak QPO harmonic. In units of power, the Poisson noise level

[2]The analysis scripts, processed data products, and explanatory Jupyter notebooks will be made publicly available after the public release of these observations and the publication of this paper.

Component	Parameter	Value	Notes
Poisson noise	($\times 10^{-4}$ Hz^{-1})	3.417 ± 0.003	
BBN_1	FWHM (Hz)	$0.6^{+0.5}_{-0.2}$	
BBN_1	$\nu_{\rm centroid}$ (Hz)	1×10^{-6}	[1]
BBN_1	norm. ($\times 10^{-4}$)	1.8 ± 0.2	
BBN_2	FWHM (Hz)	2.7 ± 0.2	
BBN_2	$\nu_{\rm centroid}$ (Hz)	0.5 ± 0.2	
BBN_2	norm. ($\times 10^{-4}$)	9.0 ± 0.2	
QPO	FWHM (Hz)	2.84	[2]
QPO	$\nu_{\rm centroid}$ (Hz)	$5.67^{+0.07}_{-0.08}$	
QPO	norm. ($\times 10^{-4}$)	3.5 ± 0.1	
Harmonic	FWHM (Hz)	3^{+3}_{-1}	
Harmonic	$\nu_{\rm centroid}$ (Hz)	11.0 ± 0.5	
Harmonic	norm. ($\times 10^{-5}$)	5 ± 1	

Table 5.1: Fit parameters to the power spectra shown in Figure 5.3. The Poisson noise was modeled with a power-law and the two broadband noise (BBN) components, QPO, and QPO harmonic were modeled with Lorentzians. The fit to the power spectrum yielded a χ^2 fit statistic of 188.98 for 179 degrees of freedom. The errors on the parameters represent the 90% confidence region computed from the XSPEC MCMC routine. Notes in the final column: [1]: This centroid was frozen at 1×10^{-6} to represent the low-frequency broadband noise with an approximately zero-frequency Lorentzian and eliminate degeneracy between the two broadband noise components. [2]: Forced to be half the QPO centroid value.

is well-modeled by a constant.

The total rms is low, but the power spectrum is dominated by a QPO-like feature with a centroid frequency of $\nu_{\rm centroid} = 5.67$ Hz. Since it is relatively broad, the width of this feature is degenerate with the other components in the power spectrum, so in the fits, we forced a quality factor $Q = 2$ ($Q = \nu_{\rm centroid} / {\rm FWHM}$). Thus, the assumed FWHM of the feature is 2.84 Hz by definition. The resulting parameters of the fit to the power spectrum are listed in Table 5.1. The QPO contributes a 3–10 keV rms of 1.8%. Although this feature is quite broad compared to more commonly studied QPOs, it is distinct from the broadband noise and possesses a clear — if weak — harmonic. Therefore we will refer to it as a QPO for the remainder of this paper.

It is interesting to note that the fitted noise level is only 0.6% lower than that expected given the count rate (2/5186, i.e., 3.439×10^{-4} in units of fractional rms^2 Hz^{-1}). Poisson noise levels are reduced by instrumental deadtime effects, which suppress the observed noise variance due to the resulting (anti-)correlations between successive photon counts. The data show that the fraction of photons lost to deadtime in the *NICER* detectors is indeed remarkably small for such a bright source.

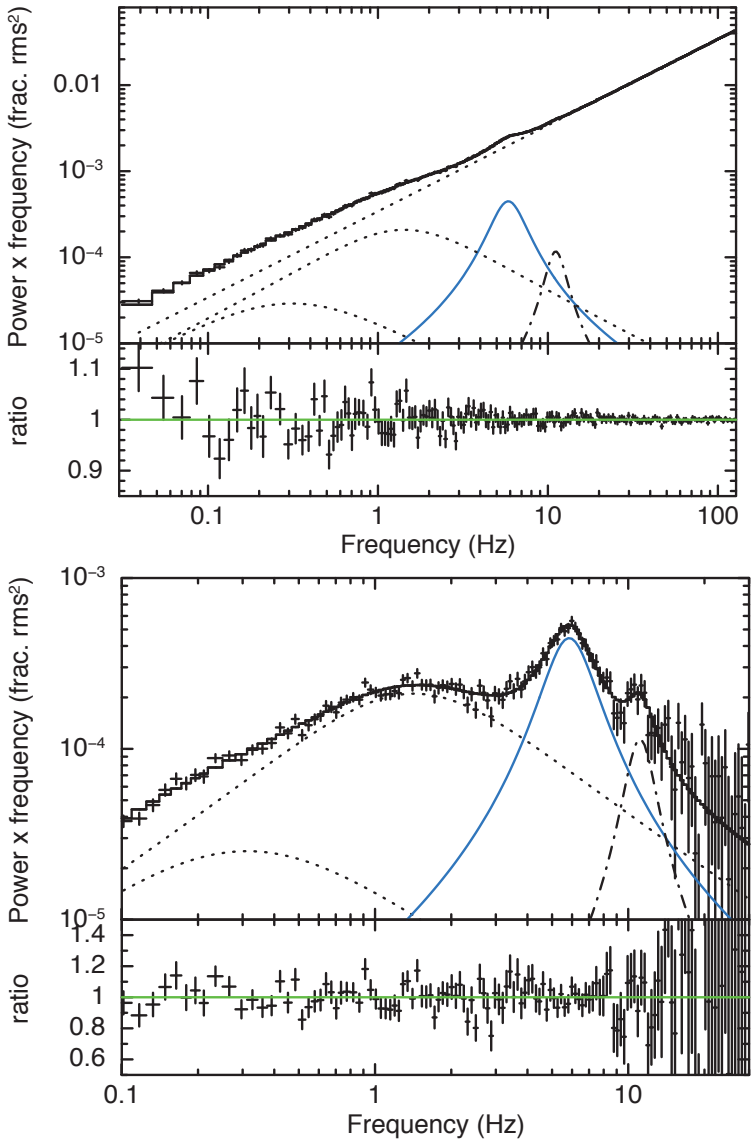

Figure 5.3: Power spectra in the energy band 3–10 keV, with and without Poisson noise (TOP and BOTTOM, respectively). The data are in black points, the Lorentzian fitted to the QPO fundamental is plotted in solid blue, the QPO harmonic is plotted in black dash-dotted, the two broadband noise Lorentzians and Poisson noise power-law are in dotted black lines, and the total model is plotted in a solid black line. Both panels also show a ratio of the total model (green line) with the data. The QPO (fundamental) has a centroid frequency of 5.67 Hz and a FWHM of 2.84 Hz (forcing $Q = 2$).

5.3.2 Lag-Energy Spectrum

The lag-energy spectrum, computed from the average cross spectrum, measures how much the variability in many narrow energy bands leads or lags the variability in a broad 'reference' energy band (see Uttley et al. 2014 for an overview of lag-energy spectra). To ensure that the narrow-band and broad reference band light curves are truly independent, we use the detector information stored in the event lists as a selection criterion. The *NICER* XTI is comprised of 56 concentrator-detector pairs, which are grouped into 7 chains (with 8 detectors on one chain) for read-out purposes. We used chains 0–3 (inclusive) for the narrow bands and chains 4–6 (inclusive) for the reference band. This division of chains was chosen to optimize the signal-to-noise in the narrow bands while maintaining enough signal in the reference band for fitting the power spectrum in Section 5.3.3. To further improve signal-to-noise, we bin up the intrinsic CCD energy resolution of the data so that there are 62 narrow-band 'channels of interest' across the energy range 0.1–11 keV. The energy resolution of the lag-energy spectrum (and subsequent analysis in Section 5.3.3) is thus limited by the width of the channels of interest, which are comparable to or broader than the instrumental resolution. Due to the strong absorption, there are few intrinsically low-energy source counts below ~ 1 keV and the spectrum there is dominated by photons redistributed from higher energies. Thus, 1 keV is chosen as the lower energy bound for the rest of the plots and analysis.

To compare the lag-energy spectra of the broadband noise and the QPO components, cross-spectra were averaged and resulting phase and time-lags calculated across two Fourier frequency ranges: 0.015625–1.0 Hz (intrinsic broadband noise), 4.25–7.09 Hz (QPO). The frequency range used for the QPO was chosen to be the FWHM centered on the centroid for the fitted Lorentzian model in the power spectrum (Figure 5.3). The lag-energy spectra are shown in Figure 5.4.

The broadband noise roughly follows a smooth 'soft lag' shape, in which the variability at softer energies lags the variability at harder energies. The QPO lags also show a clear soft lag shape, but with a clear flattening above 4 keV. The QPO soft lags are also very large: the 50 ms timescale of the lag at ~ 1.5 keV corresponds to 28% of a QPO cycle (where one QPO cycle is ~ 176 ms, computed from the centroid frequency of 5.67 Hz). The shape of the QPO lag-energy spectrum is reminiscent of the inflected lag-energy spectrum observed for the Type B QPO of GX 339-4 (Stevens & Uttley 2016 and Chapter 2 of this thesis), albeit with the *opposite* sign of lags. To better understand this behavior, we therefore follow the approach of Stevens & Uttley (2016) and study the phase-resolved spectral variability of the MAXI J1535–571 soft-intermediate state LF QPO.

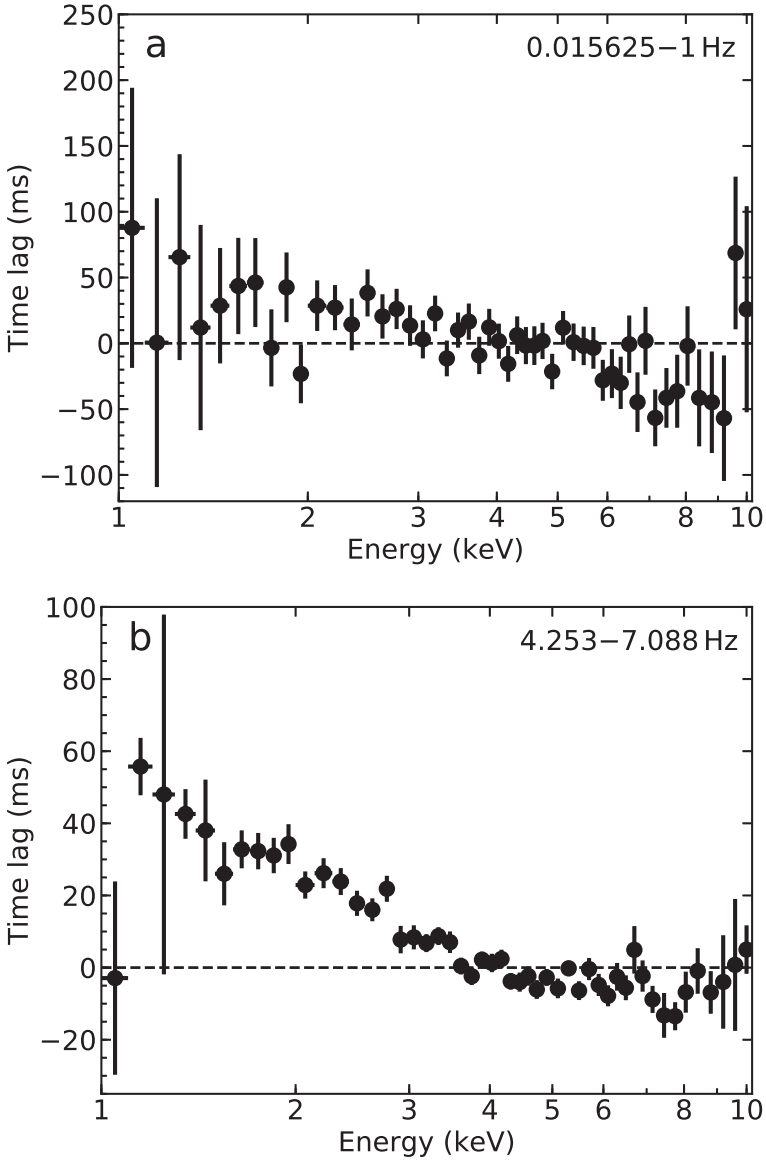

Figure 5.4: Lag-energy spectra computed over the frequency ranges 0.015625–1.0 Hz (**a**, broadband noise) and 4.253–7.088 Hz (**b**, LF QPO). A positive sign indicates that the variation in that energy channel arrived after the variation in the broad 3–10 keV reference band.

5.3.3 Phase-Resolved Spectral Variability

We followed the method of Stevens & Uttley (2016) (see Chapters 2 and 4 of this thesis) to investigate phase-resolved spectral variability on the time-scale of the LF QPO. Since the broadband noise variability amplitude is relatively large compared to the QPO-like signal, we filtered the average cross spectrum by applying a window function over the FWHM of the QPO (4.25–7.09 Hz). The QPO signal in the cross spectrum is weak (and therefore difficult to fit in all channels-of-interest to obtain an energy-dependent filter), but its amplitude is energy-dependent. By applying a window function (instead of an optimal filter constructed from the fitted cross spectrum amplitude as in Chapter 4), we make fewer assumptions about the shape and relative strength of the QPO-like signal (and how it changes in energy) albeit at the cost of some distortion of the signal waveform. However, given the relatively large width of the feature, we expect such distortion to be fairly weak. Hence, using a non-optimal filter but with minimal other assumptions is a reasonable trade-off.

We apply the inverse discrete Fourier transform to the filtered average cross spectrum of each channel-of-interest with the reference band, to obtain the energy-dependent cross-correlation function (CCF). As outlined in Chapters 2 and 4, we also apply a normalization of $2/N$, where N is the number of time steps in a segment of the light curve (16,384). This gives the CCF units of (count s^{-1})^2 in each channel of interest. We then divide the CCFs by the rms of the QPO in the reference band (in absolute rms-squared normalization, i.e units of count s^{-1}) so that the CCF only contains the amplitudes of count rate variation in the channels of interest. To compute the rms of the QPO in the reference band, we produced an averaged power spectrum of detector chains 4–6 in the energy range 3–10 keV, rebinned in frequency by a factor of 1.03, and fitted with the same power-law plus four Lorentzian model as was used in Section 5.3.1.

The CCF is shown in six energy channels in Figure 5.5. Here, the phase difference between the soft and hard energies is visible. To illustrate the relative amplitude of the CCF in each energy channel, we take the ratio of the CCF amplitude with the mean count rate in each energy channel (since the CCF has a larger absolute amplitude at lower energies, but the lower energies also have more counts). This 2-dimensional CCF showing both energy-dependence and time-delay-dependence is shown in Figure 5.6. Again, we see that the QPO in the hard energies leads the QPO in the soft energies (as was also indicated by the lag-energy spectrum).

Since the in-flight spectral calibration of *NICER* is still in progress and there are energy-dependent systematics in effective area of up to 5–10%, we cannot yet carry out full phase-resolved spectroscopy as in Chapters 2 and 4. Instead, we look at 'ratio spectra', ratios of the phase-resolved spectra to the mean spectrum (Figure 5.7). This approach accounts for the uncertainties in the instrument's effective area (which is simply divided out) and shows relative amplitude changes with energy. We use the

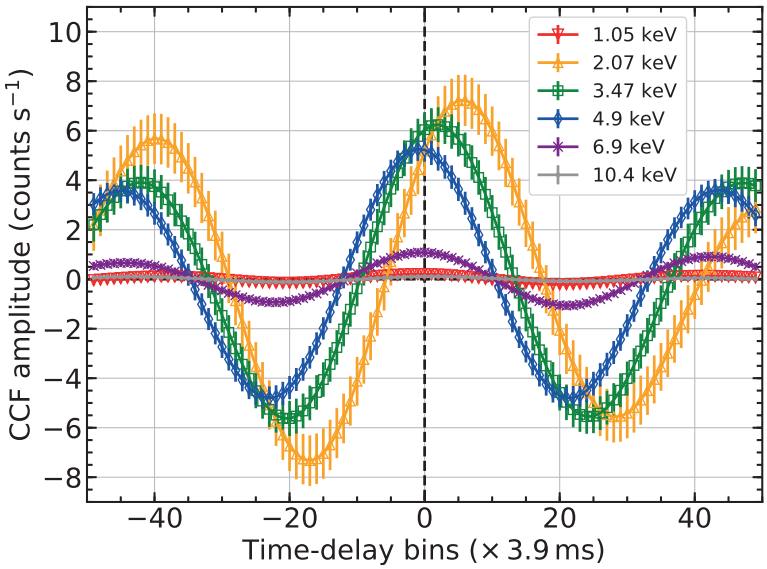

Figure 5.5: The CCF in six energy channels (corresponding to mean energies of 1.05, 2.07, 3.47, 4.90, 6.90, and 10.40 keV), computed from the filtered average cross spectrum. The time-delay bins have the same time resolution as the binned light curves ($1/256\,\mathrm{s} \sim 3.9\,\mathrm{ms}$). The CCF has a larger amplitude in the lower energy channels, but those lower energy channels also have higher mean count rates. Phase differences between the energies are already apparent.

mean count rate per binned channel of interest as the 'mean spectrum', and compute the error on the mean spectrum as the square root of the total counts in each energy bin, divided by the exposure time.

In Figure 5.7 we show ratio spectra at four phases of a QPO cycle (selected such that 0° is approximately at the minimum in the 4.9 keV channel). We see that the phase-dependent changes have larger fractional amplitudes at higher energies, as can also be inferred from Figure 5.6. The pink and green curves, indicating 90° and 270° respectively, correspond to the maximum and minimum flux phases at low energies ($\lesssim 2\,\mathrm{keV}$), while the purple and orange curves, indicating 0° and 180° respectively, are the minimum and maximum flux phases at higher energies ($\gtrsim 3\,\mathrm{keV}$).

To determine the approximate flux changes with phase, we can attempt to crudely fit the mean energy spectrum (Figure 5.8) with a simple continuum model appropriate to the soft-intermediate state. In so doing, we assume that the broadband flux distribution, which covers a relatively wide dynamic range in energy and flux, should be robust to the maximum 5–10% systematic errors in the current calibration. The mean spectrum was fitted in XSPEC with the model TBNEW_GAS × (SIMPLER * DISKBB), with 5% systematic error. The model corresponds to a multi-temperature accretion disk blackbody which contributes seed photons to a Comptonized component, with

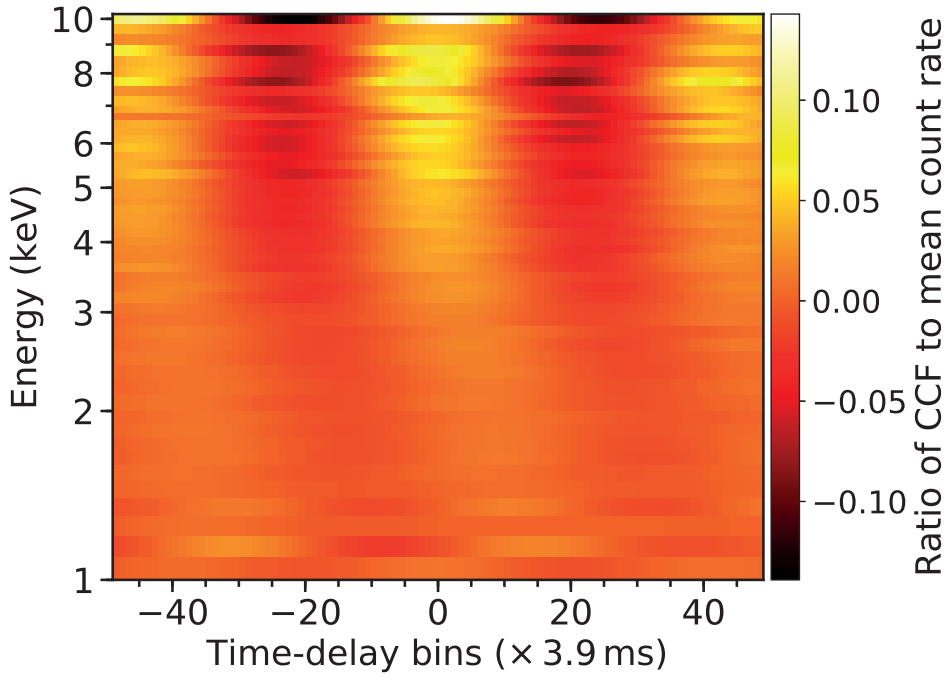

Figure 5.6: The energy-dependent cross-correlation function (CCF). The colour map shows the ratio of the CCF amplitude with the mean count rate in each energy channel. The CCF amplitude represents deviations from the mean count rate per energy channel (hence why it can be negative). One QPO cycle roughly corresponds to 45.15 time-delay bins (computed from the 5.67 Hz centroid frequency).

both components being absorbed by the Galactic absorbing column. The SIMPLER model, used also by Stevens & Uttley (2016) approximates a Compton-upscattered power-law with low-energy cut-off matching the shape of the seed spectrum, in this case the disk blackbody component (DISKBB). The power-law photon index was frozen at $\Gamma = 2$. The absorbing column density was frozen to $N_{\rm H} = 3.4 \times 10^{22}$ cm^{-2} (Negoro et al. 2017a). The best-fitting disk temperature was 1.3 keV and disk normalization was 1700. With $\Gamma = 2$, the best-fitting scattering fraction was 0.07. These model parameters are given for completeness, but should not be considered robust, given the current limitations of the calibration. The total spectrum was also multiplied by a constant of 1.75 to account for only using data from 4 out of the 7 detector chains, since the only response matrix currently available is for the combined effective area of all the detector chains (the disk normalization value quoted above has been re-scaled to account for this response factor).

We can combine our estimated model continuum spectrum with the robust phase-

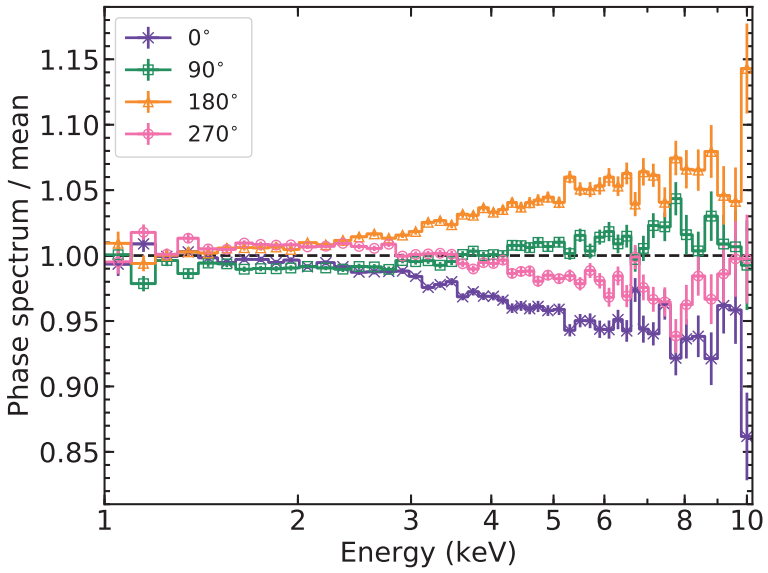

Figure 5.7: Ratios of the phase-resolved spectra with the mean spectrum at four phases of a QPO cycle. The black dashed line denotes no change from the mean spectrum. The purple X's are for a QPO phase of 0° (time-delay bin -23 in Figure 5.6), green squares are for 90° (time-delay bin -11), orange triangles are for 180° (time-delay bin of 1), and pink circles are for 270° (time-delay bin 13). The shape of the spectrum changes more significantly above 3 keV.

resolved ratio spectra obtained via the CCF, to estimate the flux variation as a function of energy. Figure 5.9 shows difference spectra obtained by multiplying the ratio spectra from Fig. 5.7 by the model used to fit the mean spectrum in Figure 5.8 (except with absorption removed so that the estimated variation of the disk blackbody component can be fully seen). The model spectrum is then subtracted from the resulting estimated phase-resolved spectra, to show a difference spectrum in each time-delay bin. The phase lag of the soft photons, likely that of the disk emission relative to the Comptonized component, is apparent. Furthermore it is clear that the disk variation is still relatively weak in flux compared to the power-law variation, even though the disk dominates the total flux in the mean spectrum. Integrating the flux in energy bands of 1–2 keV and 7–10 keV for each time-delay bin, we find that the variation in the hard band leads the variation in the soft band by 27% of a QPO cycle, consistent with the measured lag-energy spectrum (Figure 5.4).

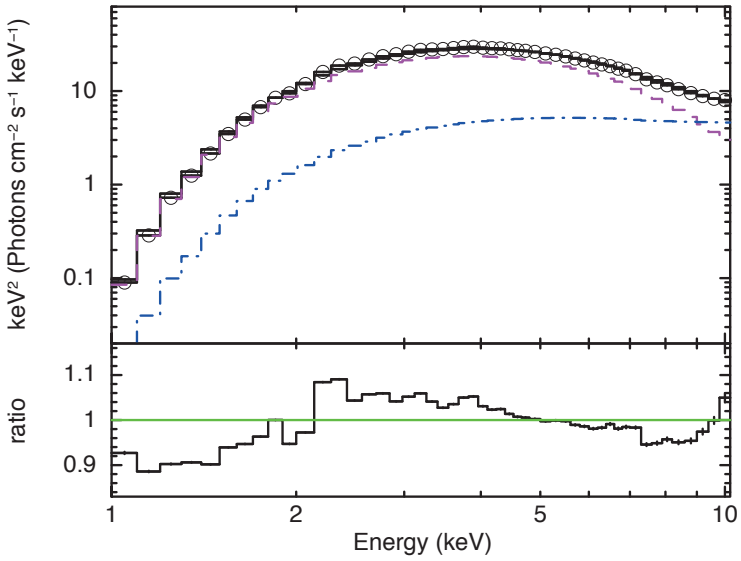

Figure 5.8: The mean energy spectrum of the data (black circles) shown with the total model (solid black), multi-colour disk blackbody (dashed magenta), and upscattered power-law (dot-dashed blue). In the lower panel, the residuals are shown as a ratio of the model to the data. The calibration systematics contribute to the residuals.

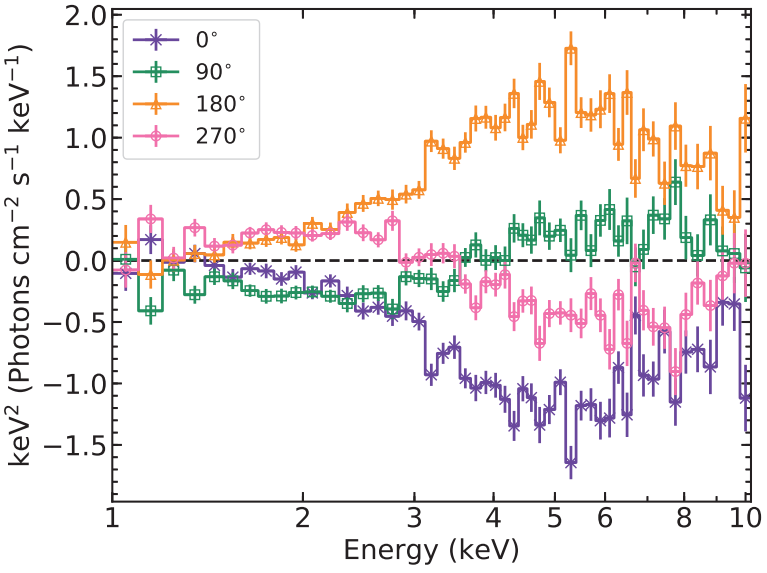

Figure 5.9: Deviations from the modelled mean flux spectrum at four parts of a QPO cycle. Absorption was removed from the assumed model spectrum. The colours have the same meanings as in Figure 5.7.

5.4 Discussion and Conclusion

We have reported the discovery and initial spectral-timing analysis of a QPO feature in the soft-intermediate state of the 2017–2018 outburst of the new candidate black hole transient MAXI J1535–571. This QPO has a centroid frequency of 5.67 Hz and is relatively broad, consistent with $Q \sim 2$. It is also weak, with an integrated fractional rms of 1.8% in 3–10 keV. The QPO shows strong soft lags with respect to the 3–10 keV reference band of about 50 ms, which is notably different from the weaker lags of the lower-frequency broadband noise, which also shows a smoother lag-energy spectrum compared to the inflected shape of the QPO lag-energy spectrum, which flattens above ~ 4 keV.

Given the QPO's appearance at times when the total rms is low, and centroid frequency between 5 and 6 Hz, it is tempting to identify this feature with the Type B QPO. In fact, the closest analogy appears to be the QPOs originally identified by Wijnands et al. (1999) as a Type A QPO in the black hole transient XTE J1550–564 (and Type A-I by Homan et al. 2001), and later identified by Casella et al. (2005) as a form of Type B QPO. Besides the similar frequencies (~ 5.7 Hz in MAXI J1535–571 versus ~ 5.9 Hz in XTE J1550–564), the shared characteristics of the MAXI J1535–571 QPO and the QPO in XTE J1550–564 are a low Q-value ($Q < 3$ in XTE J1550–564), a small rms (2–3% in XTE J1550–564) and occurrence during the soft-intermediate state (what Homan et al. 2001 referred to as the 'very high state' between the hard and soft states). Interestingly, soft lags are also observed for the similar QPO in XTE J1550–564 (Wijnands et al. 1999), albeit across the harder 3–50 keV energy range covered by *RXTE*.

The QPO analyzed in this chapter has some notable differences with the clear Type B QPO in Chapter 2. There, the QPO had a fractional rms of $\sim 14\%$ in 5–10 keV, and the broadband noise was much weaker. The QPO showed hard lags, and so the $\sim 25\%$ fractional rms variation in the power-law parameters *lagged* the $\sim 1.4\%$ variation in the soft blackbody flux. However, it is also notable that the inflected shapes of the lag-energy spectra observed in both sources are similar, but with opposite sign. It is possible that there is a similar mechanism for both QPOs, namely a large-scale-height precessing power-law emitting region, possibly the base of the jet (as we suggested for the Type B QPO in Chapter 2). The inflected shape of the lag-energy spectrum in MAXI J1535–571 is then linked to a systematic phase-offset between the peaks in power-law emission and the modulation of the disk blackbody spectrum. Our reconstruction of phase-dependent spectral variability (see Figure 5.9) suggests that the disk and power-law component variations are indeed shifted in phase, but unlike the case of GX 339–4, the disk component in MAXI J1535–571 *lags* rather than leads the power-law component by $\sim 30\%$ in phase.

The phase lead of disk relative to power-law variations in GX 339-4 was attributed to a geometry where the approaching side of the disk is illuminated by the jet base,

leading to enhanced (blue-shifted) disk emission, before the jet base points towards the observer leading to enhanced power-law emission (e.g., due to beaming effects). The flipped sign of the lag in MAXI J1535–571 could therefore mean that the jet precession is retrograde compared to the motion of gas in the disk. Alternatively, the maximum in jet base power-law emission might be seen when the jet is pointing away from the observer, e.g., if solid angle effects dominate over beaming effects. The latter scenario might be linked to the higher inclination of the MAXI J1535–571 system compared to GX 339-4, as suggested by measurements of the iron line in the hard-intermediate state (Xu et al. 2017).

Finally, we note that these observations highlight the enormous potential of *NICER* for making breakthroughs in our understanding of accreting compact objects via spectral-timing methods. At the peak flux levels of the soft-intermediate state of MAXI J1535–571, the total *NICER* count rate exceeded 17,000 count s^{-1}, with very little deadtime and hence minimal spectral distortion. The combination of such large count rates with CCD energy resolution and a soft X-ray response is a revolutionary capability, which can be fully realized once *NICER*'s calibration is completed. The mission will act as a pathfinder for potential future dedicated spectral-timing missions such as *eXTP* (Zhang et al. 2016) and *STROBE-X* (Wilson-Hodge et al. 2017), enabling new breakthroughs using the techniques presented in this thesis and pointing the way to a bright future for our understanding of the innermost regions of accreting compact objects.

Acknowledgements

A.L.S. acknowledges support from NOVA (Nederlandse Onderzoekschool voor Astronomie). A.L.S. and P.U. are grateful for the support and help of the *NICER* team (especially Keith Gendreau and Zaven Arzoumanian), data reduction pipeline development led by Craig Markwardt, and leadership of the observatory science program by Ron Remillard.

This research made use of data and software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC); NASA's Astrophysics Data System Bibliographic Services; NumPy v1.13.1 and Scipy v0.19.1 (Jones et al. 2001); Astropy v2.0 (Astropy Collaboration et al. 2013); Matplotlib v2.0.2 (Hunter 2007); and Jupyter Notebook v5.0.0 with iPython v6.1.0 (Perez & Granger 2007).

Bibliography

Akmal, A., Pandharipande, V. R., & Ravenhall, D. G. 1998, Phys. Rev. C, 58, 1804

AlGendy, M., & Morsink, S. M. 2014, ApJ, 791, 78

Altamirano, D., van der Klis, M., Méndez, M., et al. 2008, ApJ, 685, 436

Antia, H. M., Yadav, J. S., Agrawal, P. C., et al. 2017, ApJS, 231, 10

Arévalo, P., & Uttley, P. 2006, MNRAS, 367, 801

Armas Padilla, M., Ueda, Y., Hori, T., Shidatsu, M., & Muñoz-Darias, T. 2017, MNRAS, 467, 290

Arnaud, K. A. 1996, in Astronomical Society of the Pacific Conference Series, Vol. 101, Astronomical Data Analysis Software and Systems V, ed. G. H. Jacoby & J. Barnes, 17

Artigue, R., Barret, D., Lamb, F. K., Lo, K. H., & Miller, M. C. 2013, MNRAS, 433, L64

Arzoumanian, Z., Gendreau, K. C., Baker, C. L., et al. 2014, in Proc. SPIE, Vol. 9144, Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray, 914420

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33

Axelsson, M., & Done, C. 2016, MNRAS, 458, 1778

Axelsson, M., Done, C., & Hjalmarsdotter, L. 2014, MNRAS, 438, 657

Baker, T., Psaltis, D., & Skordis, C. 2015, ApJ, 802, 63

Barret, D. 2013, ApJ, 770, 9

Barret, D., Kluźniak, W., Olive, J. F., Paltani, S., & Skinner, G. K. 2005, MNRAS, 357, 1288

Barret, D., Olive, J.-F., & Miller, M. C. 2006, MNRAS, 370, 1140

Barret, D., & Vaughan, S. 2012, ApJ, 746, 131

Bartlett, M. S. 1955, An Introduction to Stochastic Processes (Cambridge, UK: Cambridge University Press)

Bauböck, M., Berti, E., Psaltis, D., & Özel, F. 2013, ApJ, 777, 68

Bauböck, M., Psaltis, D., & Özel, F. 2015, ApJ, 811, 144

Belloni, T. 2004, Nuclear Physics B Proceedings Supplements, 132, 337

Belloni, T., Homan, J., Casella, P., et al. 2005, A&A, 440, 207

Belloni, T. M. 2010, in Lecture Notes in Physics, Vol. 794, (Berlin: Springer-Verlag),

ed. T. Belloni, 53
Belloni, T. M., Sanna, A., & Méndez, M. 2012, MNRAS, 426, 1701
Beloborodov, A. M. 1999, in Astronomical Society of the Pacific Conference Series, Vol. 161, High Energy Processes in Accreting Black Holes, ed. J. Poutanen & R. Svensson, 295
Beloborodov, A. M. 2002, ApJ, 566, L85
Berger, M., van der Klis, M., van Paradijs, J., et al. 1996, ApJ, 469, L13
Bogdanov, S. 2013, ApJ, 762, 96
Box, G. E. P., & Jenkins, G. M. 1976, Time Series Analysis: Forecasting and Control, revised edn. (San Francisco: Holden-Day, Inc.)
Bradt, H. V., Rothschild, R. E., & Swank, J. H. 1993, A&AS, 97, 355
Britt, C. T., Bahramian, A., & Strader, J. 2017, ATel, 10816
Bult, P., van Doesburgh, M., & van der Klis, M. 2017, ApJ, 845, 124
Cackett, E. M. 2016, ApJ, 826, 103
Cadeau, C., Morsink, S. M., Leahy, D., & Campbell, S. S. 2007, ApJ, 654, 458
Camenzind, M. 2007, Compact Objects in Astrophysics: White Dwarfs, Neutron Stars, and Black Holes (Berlin: Springer-Verlag), doi:10.1007/978-3-540-49912-1
Casella, P., Belloni, T., & Stella, L. 2005, ApJ, 629, 403
Chakrabarti, S. K. 1996, ApJ, 471, 237
Chakrabarty, D., & Morgan, E. H. 1998a, IAU Circ., 6877
—. 1998b, Nature, 394, 346
Chandrasekhar, S. 1960, Radiative Transfer (New York: Dover)
Christian, D. J., & Swank, J. H. 1997, ApJS, 109, 177
Corbel, S., Kaaret, P., Jain, R. K., et al. 2001, ApJ, 554, 43
Dauser, T., Garcia, J., Wilms, J., et al. 2013, MNRAS, 430, 1694
de Avellar, M. G. B., Méndez, M., Altamirano, D., Sanna, A., & Zhang, G. 2016, MNRAS, 461, 79
de Avellar, M. G. B., Méndez, M., Sanna, A., & Horvath, J. E. 2013, MNRAS, 433, 3453
Di Salvo, T., Méndez, M., & van der Klis, M. 2003, A&A, 406, 177
Di Salvo, T., Méndez, M., van der Klis, M., Ford, E., & Robba, N. R. 2001, ApJ, 546, 1107
Dincer, T. 2017, ATel, 10716
Done, C., & Gierliński, M. 2003, MNRAS, 342, 1041
Done, C., Gierliński, M., & Kubota, A. 2007, A&A Rev., 15, 1
Elshamouty, K. G., Heinke, C. O., Morsink, S. M., Bogdanov, S., & Stevens, A. L. 2016, ApJ, 826, 162
Esin, A. A., McClintock, J. E., & Narayan, R. 1997, ApJ, 489, 865
Fabian, A. C. 2013, in IAU Symposium, Vol. 290, Feeding Compact Objects: Accretion on All Scales, ed. C. M. Zhang, T. Belloni, M. Méndez, & S. N. Zhang, 3–12

Feigelson, E. D., & Babu, G. J. 2012, Modern Statistical Methods for Astronomy With R Applications (Cambridge, UK: Cambridge University Press)
Fender, R. P., Belloni, T. M., & Gallo, E. 2004, MNRAS, 355, 1105
Fender, R. P., Homan, J., & Belloni, T. M. 2009, MNRAS, 396, 1370
Feroci, M., Stella, L., van der Klis, M., et al. 2012, Experimental Astronomy, 34, 415
Fiege, J., McCurdy, B., Potrebko, P., Champion, H., & Cull, A. 2011, Medical Physics, 38, 5217
Fiege, J. D. 2010, Qubist User's Guide: Optimization, Data-Modeling, and Visualization with the Qubist Global Optimization Toolbox for MATLAB, 4th edn., nQube Technical Computing Corporation
Fish, V., Akiyama, K., Bouman, K., et al. 2016, Galaxies, 4, 54
Fragile, P. C., & Anninos, P. 2005, ApJ, 623, 347
Fragile, P. C., Miller, W. A., & Vandernoot, E. 2005, ApJ, 635, 157
Fragile, P. C., Straub, O., & Blaes, O. 2016, MNRAS, 461, 1356
Frank, J., King, A., & Raine, D. J. 2002, Accretion Power in Astrophysics, 3rd edn. (Cambridge, UK: Cambridge University Press)
Franzmann, E. 2014, Master's thesis, University of Manitoba (Canada)
Gallo, E., Fender, R. P., & Pooley, G. G. 2003, MNRAS, 344, 60
Galloway, D. K., Muno, M. P., Hartman, J. M., Psaltis, D., & Chakrabarty, D. 2008, ApJS, 179, 360
Gao, H. Q., Qu, J. L., Zhang, Z., et al. 2014, MNRAS, 438, 341
García, J. A., McClintock, J. E., Steiner, J. F., Remillard, R. A., & Grinberg, V. 2014, ApJ, 794, 73
García, J. A., Steiner, J. F., McClintock, J. E., et al. 2015, ApJ, 813, 84
Gendreau, K., Arzoumanian, Z., Markwardt, C., et al. 2017, ATel, 10768
Gendreau, K. C., Arzoumanian, Z., & Okajima, T. 2012, in Proc. SPIE, Vol. 8443, Space Telescopes and Instrumentation 2012: Ultraviolet to Gamma Ray, 844313
Ghosh, P., & Lamb, F. K. 1979, ApJ, 232, 259
Giannios, D., Kylafis, N. D., & Psaltis, D. 2004, A&A, 425, 163
Gierliński, M., Done, C., & Page, K. 2008, MNRAS, 388, 753
Gierliński, M., Zdziarski, A. A., Poutanen, J., et al. 1999, MNRAS, 309, 496
Gilfanov, M. 2010, in Lecture Notes in Physics, Berlin Springer Verlag, Vol. 794, Lecture Notes in Physics, Berlin Springer Verlag, ed. T. Belloni, 17
Gilfanov, M., Churazov, E., & Revnivtsev, M. 2000, in Proceedings of 5th Sino-German Workshop on Astrophysics, 1999, SGSC Conference Series, ed. G. Zhao, J.-J. Wang, H.-M. Qiu, & G. Boerner, Vol. 1, 114–123
Gilfanov, M., Revnivtsev, M., & Molkov, S. 2003, A&A, 410, 217
Glendenning, N. K. 2000, Compact Stars: Nuclear Physics, Particle Physics, and General Relativity (New York: Springer)
Goldberg, D. E. 1989, Genetic Algorithms in Search, Optimization and Machine Learning (Reading: Addison-Wesley)

—. 2002, The Design of Innovation: Lessons from and for Competent Genetic Algorithms (Kluwer Academic Publishers)
Güver, T., Özel, F., Cabrera-Lavers, A., & Wroblewski, P. 2010, ApJ, 712, 964
Harding, A. K. 2013, Journal of Astronomy and Space Sciences, 30, 145
Hasinger, G., & van der Klis, M. 1989, A&A, 225, 79
Hawking, S. W. 1975, Communications in Mathematical Physics, 43, 199
Heil, L. M., Uttley, P., & Klein-Wolt, M. 2015a, MNRAS, 448, 3348
—. 2015b, MNRAS, 448, 3339
Heinke, C. O., Rybicki, G. B., Narayan, R., & Grindlay, J. E. 2006, ApJ, 644, 1090
Holland, J. H. 1975, Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence (Ann Arbor: University of Michigan Press)
Homan, J., & Belloni, T. 2005, Ap&SS, 300, 107
Homan, J., Wijnands, R., van der Klis, M., et al. 2001, ApJS, 132, 377
Hunter, J. D. 2007, Computing in Science Engineering, 9, 90
Hynes, R. I., Steeghs, D., Casares, J., Charles, P. A., & O'Brien, K. 2003, ApJ, 583, L95
—. 2004, ApJ, 609, 317
Ibrahim, A. I., Safi-Harb, S., Swank, J. H., et al. 2002, ApJ, 574, L51
Ingram, A., & Done, C. 2012, MNRAS, 427, 934
Ingram, A., Done, C., & Fragile, P. C. 2009, MNRAS, 397, L101
Ingram, A., & van der Klis, M. 2015, MNRAS, 446, 3516
Ingram, A., van der Klis, M., Middleton, M., Altamirano, D., & Uttley, P. 2017, MNRAS, 464, 2979
Ingram, A., van der Klis, M., Middleton, M., et al. 2016, MNRAS, 461, 1967
Inogamov, N. A., & Sunyaev, R. A. 1999, Astronomy Letters, 25, 269
Israel, W. 1968, Communications in Mathematical Physics, 8, 245
Jahoda, K., Swank, J. H., Giles, A. B., et al. 1996, in Proc. SPIE, Vol. 2808, EUV, X-Ray, and Gamma-Ray Instrumentation for Astronomy VII, ed. O. H. Siegmund & M. A. Gummin, 59–70
Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python, [Online]
Kaaret, P., Piraino, S., Ford, E. C., & Santangelo, A. 1999, ApJ, 514, L31
Kalamkar, M., Casella, P., Uttley, P., et al. 2016, MNRAS, 460, 3284
Kennea, J. A. 2017, ATel, 10731
Kennea, J. A., Evans, P. A., Beardmore, A. P., et al. 2017, ATel, 10700
Kluzniak, W., & Abramowicz, M. A. 2001, Acta Physica Polonica B, 32, 3605
Koljonen, K. I. I. 2015, MNRAS, 447, 2981
Kubota, A., Ebisawa, K., Makishima, K., & Nakazawa, K. 2005, ApJ, 631, 1062
Kubota, A., & Makishima, K. 2004, ApJ, 601, 428
Kumar, N., & Misra, R. 2014, MNRAS, 445, 2818

—. 2016, MNRAS, 461, 2580
Lamb, F. K., & Miller, M. C. 2001, ApJ, 554, 1210
Lasota, J.-P. 2001, New A Rev., 45, 449
Lattimer, J. M., & Prakash, M. 2016, Phys. Rep., 621, 127
Leahy, D. A. 2004, ApJ, 613, 517
Leahy, D. A., Darbro, W., Elsner, R. F., et al. 1983, ApJ, 266, 160
Leahy, D. A., Morsink, S. M., & Chou, Y. 2011, ApJ, 742, 17
Leahy, D. A., Morsink, S. M., Chung, Y.-Y., & Chou, Y. 2009, ApJ, 691, 1235
Lee, H. C., Misra, R., & Taam, R. E. 2001, ApJ, 549, L229
Lin, D., Remillard, R. A., & Homan, J. 2007, ApJ, 667, 1073
Liska, M., Hesp, C., Tchekhovskoy, A., et al. 2017, MNRAS, in press, arXiv:1707.06619
Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, ApJ, 776, 19
Ludlam, R. M., Miller, J. M., & Cackett, E. M. 2015, ApJ, 806, 262
Lyubarskii, Y. E. 1997, MNRAS, 292, 679
Maccarone, T. J. 2002, MNRAS, 336, 1371
—. 2015, MNRAS, 446, 3162
Mahmoodifar, S., & Strohmayer, T. 2016, ApJ, 818, 93
Markoff, S., Nowak, M. A., & Wilms, J. 2005, ApJ, 635, 1203
Markwardt, C. B., Burrows, D. N., Cummings, J. R., et al. 2017, GRB Coordinates Network, 21788
McClintock, J. E., & Remillard, R. A. 2006, in Compact Stellar X-ray Sources, ed. W. H. G. Lewin & M. van der Klis (Cambridge, UK: Cambridge University Press), 157–213
Méndez, M. 2006, MNRAS, 371, 1925
Méndez, M., & van der Klis, M. 1999, ApJ, 517, L51
Méndez, M., van der Klis, M., & Ford, E. C. 2001, ApJ, 561, 1016
Méndez, M., van der Klis, M., Wijnands, R., et al. 1998a, ApJ, 505, L23
Méndez, M., van der Klis, M., van Paradijs, J., et al. 1998b, ApJ, 494, L65
Mereminskiy, I. A., & Grebenev, S. A. 2017, ATel, 10734
Miller, J. M. 2007, ARA&A, 45, 441
Miller, J. M., & Homan, J. 2005, ApJ, 618, L107
Miller, M. C., & Lamb, F. K. 1998, ApJ, 499, L37
—. 2015, ApJ, 808, 31
Misner, C. W., Thorne, K. S., & Wheeler, J. A. 1973, Gravitation (San Francisco: W.H. Freeman and Co.)
Mitsuda, K., Inoue, H., Koyama, K., et al. 1984, PASJ, 36, 741
Miyamoto, S., Kimura, K., Kitamoto, S., Dotani, T., & Ebisawa, K. 1991, ApJ, 383, 784
Morsink, S. M., & Leahy, D. A. 2011, ApJ, 726, 56
Morsink, S. M., Leahy, D. A., Cadeau, C., & Braga, J. 2007, ApJ, 663, 1244

Motta, S., Muñoz-Darias, T., Casella, P., Belloni, T., & Homan, J. 2011, MNRAS, 418, 2292
Motta, S. E., Casella, P., Henze, M., et al. 2015, MNRAS, 447, 2059
Motta, S. E., Rouco Escorial, A., Kuulkers, E., Muñoz-Darias, T., & Sanna, A. 2017, MNRAS, 468, 2311
Muñoz-Darias, T., Casares, J., & Martínez-Pais, I. G. 2008, MNRAS, 385, 2205
Muñoz-Darias, T., Coriat, M., Plant, D. S., et al. 2013, MNRAS, 432, 1330
Murakami, T., Inoue, H., Koyama, K., et al. 1980, ApJ, 240, L143
Murphy, S. A., & Van Der Vaart, A. W. 2000, Journal of the American Statistical Association, 95, 449
Nakahira, S., Negoro, H., Mihara, T., et al. 2017, ATel, 10729
Nakamura, N., Dotani, T., Inoue, H., et al. 1989, PASJ, 41, 617
Nandi, A., Debnath, D., Mandal, S., & Chakrabarti, S. K. 2012, A&A, 542, A56
Negoro, H., Kawase, T., Sugizaki, M., et al. 2017a, ATel, 10708
Negoro, H., Ishikawa, M., Ueno, S., et al. 2017b, ATel, 10699
Nowak, M. A. 1995, PASP, 107, 1207
Nowak, M. A., & Wagoner, R. V. 1991, ApJ, 378, 656
Nowak, M. A., Wilms, J., & Dove, J. B. 2002, MNRAS, 332, 856
Özel, F., Psaltis, D., Arzoumanian, Z., Morsink, S., & Bauböck, M. 2016, ApJ, 832, 92
Palmer, D. M., Krimm, H. A., & Swift/BAT Team. 2017, ATel, 10733
Pechenick, K. R., Ftaclas, C., & Cohen, J. M. 1983, ApJ, 274, 846
Peille, P., Barret, D., & Uttley, P. 2015, ApJ, 811, 109
Perez, F., & Granger, B. E. 2007, Computing in Science Engineering, 9, 21
Pietrini, P., & Krolik, J. H. 1995, ApJ, 447, 526
Poutanen, J., & Gierliński, M. 2003, MNRAS, 343, 1301
Poutanen, J., Ibragimov, A., & Annala, M. 2009, ApJ, 706, L129
Psaltis, D., & Johannsen, T. 2012, ApJ, 745, 1
Psaltis, D., & Norman, C. 2000, ArXiv e-prints, arXiv:astro-ph/0001391
Psaltis, D., & Özel, F. 2014, ApJ, 792, 87
Psaltis, D., Özel, F., & Chakrabarty, D. 2014, ApJ, 787, 136
Rapisarda, S., Ingram, A., Kalamkar, M., & van der Klis, M. 2016, MNRAS, 462, 4078
Remillard, R. A., & McClintock, J. E. 2006, ARA&A, 44, 49
Remillard, R. A., McClintock, J. E., Sobczak, G. J., et al. 1999a, ApJ, 517, L127
Remillard, R. A., Morgan, E. H., McClintock, J. E., Bailyn, C. D., & Orosz, J. A. 1999b, ApJ, 522, 397
Remillard, R. A., Sobczak, G. J., Muno, M. P., & McClintock, J. E. 2002, ApJ, 564, 962
Revnivtsev, M., Gilfanov, M., & Churazov, E. 1999, A&A, 347, L23
Reynolds, C. S., & Nowak, M. A. 2003, Phys. Rep., 377, 389

Reynolds, M. T., & Miller, J. M. 2013, ApJ, 769, 16
Rogers, A., & Fiege, J. D. 2011, ApJ, 727, 80
—. 2012, ApJ, 759, 27
Rogers, A., Safi-Harb, S., & Fiege, J. 2015, IAU General Assembly, 22, 2257588
Ross, R. R., & Fabian, A. C. 2005, MNRAS, 358, 211
Russell, T. D., Altamirano, D., Tetarenko, A. J., et al. 2017a, ATel, 10899
Russell, T. D., Miller-Jones, J. C. A., Sivakoff, G. R., Tetarenko, A. J., & Jacpot Xrb Collaboration. 2017b, ATel, 10711
Scaringi, S., & ASTR211 Students. 2017a, ATel, 10702
—. 2017b, ATel, 10704
Schnittman, J. D., Homan, J., & Miller, J. M. 2006, ApJ, 642, 420
Shidatsu, M., Nakahira, S., Negoro, H., et al. 2017a, ATel, 11020
—. 2017b, ATel, 10761
Shimura, T., & Takahara, F. 1995, ApJ, 445, 780
Singh, K. P., Tandon, S. N., Agrawal, P. C., et al. 2014, in Proc. SPIE, Vol. 9144, Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray, 91441S
Skipper, C. J., & McHardy, I. M. 2016, MNRAS, 458, 1696
Sobolewska, M. A., & Życki, P. T. 2006, MNRAS, 370, 405
Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, ApJ, 722, 33
Steiner, J. F., Narayan, R., McClintock, J. E., & Ebisawa, K. 2009, PASP, 121, 1279
Steiner, J. F., Reis, R. C., McClintock, J. E., et al. 2011, MNRAS, 416, 941
Stella, L., & Vietri, M. 1998, ApJ, 492, L59
—. 1999, Physical Review Letters, 82, 17
Stevens, A. L. 2013, Master's thesis, University of Alberta (Canada)
Stevens, A. L., & Uttley, P. 2016, MNRAS, 460, 2796
Strohmayer, T. E. 2001, ApJ, 554, L169
Strohmayer, T. E., Zhang, W., Swank, J. H., et al. 1996a, IAU Circ., 6320
—. 1996b, ApJ, 469, L9
Strohmayer, T. E., Zhang, W., Swank, J. H., White, N. E., & Lapidus, I. 1998, ApJ, 498, L135
Su, Y.-H., Chou, Y., Hu, C.-P., & Yang, T.-C. 2015, ApJ, 815, 74
Suková, P., Charzyński, S., & Janiuk, A. 2017, MNRAS, 472, 4327
Sunyaev, R. A., & Titarchuk, L. G. 1980, A&A, 86, 121
Terrell, Jr., N. J. 1972, ApJ, 174, L35
Tetarenko, A. J., Russell, T. D., Miller-Jones, J. C. A., Sivakoff, G. R., & Jacpot Xrb Collaboration. 2017, ATel, 10745
Titarchuk, L. 2003, ApJ, 591, 354
Troyer, J. S., & Cackett, E. M. 2017, ApJ, 834, 131
Uttley, P., Cackett, E. M., Fabian, A. C., Kara, E., & Wilkins, D. R. 2014, A&A Rev., 22, 72
Uttley, P., Wilkinson, T., Cassatella, P., et al. 2011, MNRAS, 414, L60

van den Eijnden, J., Ingram, A., Uttley, P., et al. 2017, MNRAS, 464, 2643
van der Klis, M. 1989, in NATO Advanced Science Institutes (ASI) Series C, Vol. 262, NATO Advanced Science Institutes (ASI) Series C, ed. H. Ögelman & E. P. J. van den Heuvel, 27
van der Klis, M. 1998, Advances in Space Research, 22, 925
—. 2000, ARA&A, 38, 717
—. 2005, Ap&SS, 300, 149
—. 2006a, Advances in Space Research, 38, 2675
—. 2006b, in Compact Stellar X-ray Sources, ed. W. H. G. Lewin & M. van der Klis (Cambridge, UK: Cambridge University Press), 39–112
van der Klis, M., Jansen, F., van Paradijs, J., et al. 1985, Nature, 316, 225
van der Klis, M., Swank, J. H., Zhang, W., et al. 1996a, IAU Circ., 6319
—. 1996b, ApJ, 469, L1
van der Klis, M., Wijnands, R. A. D., Horne, K., & Chen, W. 1997, ApJ, 481, L97
van Straaten, S., van der Klis, M., & Méndez, M. 2003, ApJ, 596, 1155
Vaughan, B., van der Klis, M., Lewin, W. H. G., et al. 1994, ApJ, 421, 738
Vaughan, B. A., & Nowak, M. A. 1997, ApJ, 474, L43
Vaughan, B. A., van der Klis, M., Méndez, M., et al. 1997, ApJ, 483, L115
—. 1998, ApJ, 509, L145
Vaughan, S., Uttley, P., Pounds, K. A., Nandra, K., & Strohmayer, T. E. 2011, MNRAS, 413, 2489
Venter, C., Johnson, T. J., & Harding, A. K. 2012, ApJ, 744, 34
Verdhan Chauhan, J., Yadav, J. S., Misra, R., et al. 2017, ApJ, 841, 41
Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, ApJ, 465, 487
Wang, Z., Breton, R. P., Heinke, C. O., Deloye, C. J., & Zhong, J. 2013, ApJ, 765, 151
Watts, A. L. 2012, ARA&A, 50, 609
Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, Reviews of Modern Physics, 88, 021001
Wijnands, R., Homan, J., & van der Klis, M. 1999, ApJ, 526, L33
Wijnands, R., & van der Klis, M. 1998a, IAU Circ., 6876
—. 1998b, Nature, 394, 344
Wijnands, R., van der Klis, M., Homan, J., et al. 2003, Nature, 424, 44
Wilkinson, T., Patruno, A., Watts, A., & Uttley, P. 2011, MNRAS, 410, 1513
Wilkinson, T., & Uttley, P. 2009, MNRAS, 397, 666
Wilms, J., Allen, A., & McCray, R. 2000, ApJ, 542, 914
Wilson-Hodge, C. A., Ray, P. S., Gendreau, K., et al. 2017, Results in Physics, 7, 3704
Worpel, H., Galloway, D. K., & Price, D. J. 2013, ApJ, 772, 94
Worpel, H., & Schwope, A. D. 2015, A&A, 578, A80
Xu, Y., Harrison, F. A., Garcia, J. A., et al. 2017, ArXiv e-prints, arXiv:1711.01346

Yadav, J. S., Agrawal, P. C., Antia, H. M., et al. 2016, in Proc. SPIE, Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray, 99051D

Yamaoka, K., Sugizaki, M., Nakahira, S., et al. 2010, ATel, 2380

Zdziarski, A. A., Johnson, W. N., & Magdziarz, P. 1996, MNRAS, 283, 193

Zhang, G., Méndez, M., Sanna, A., Ribeiro, E. M., & Gelfand, J. D. 2017, MNRAS, 465, 5003

Zhang, S. N., Feroci, M., Santangelo, A., et al. 2016, in Proc. SPIE, Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray, 99051Q

Życki, P. T., Done, C., & Smith, D. A. 1999, MNRAS, 309, 561

Contribution from Co-Authors

Bibliographic references for Chapters 2–5. The position in the author list reflects the amount of contribution of each co-author.

Chapter 2:

Phase-resolved spectroscopy of Type B quasi-periodic oscillations in GX 339–4
A. L. Stevens and P. Uttley
Monthly Notices of the Royal Astronomical Society, vol. 460(3), p. 2796–2810

Chapter 3:

Neutron star mass-radius constraints using evolutionary optimization
A. L. Stevens, J. D. Fiege, D. A. Leahy, and S. M. Morsink
The Astrophysical Journal, 2016, vol. 833(2), p. 244–257

Chapter 4:

Phase-resolved spectroscopy of the lower kHz QPO in 4U 1608–52
A. L. Stevens, P. Uttley, and D. Altamirano
To be submitted to Monthly Notices of the Royal Astronomical Society

Chapter 5:

NICER spectral-timing studies of a low-frequency QPO in the soft-intermediate state of MAXI J1535–571
A. L. Stevens and P. Uttley
To be submitted to The Astrophysical Journal

Summary

Compact objects like stellar-mass black holes and neutron stars are the final remnants of core-collapse supernovae of massive stars. In the general relativistic framework, compact objects are dense enough to significantly warp spacetime. By studying emission from very close to the compact object, we can decipher the effects of strong-gravity on physical processes, and test general relativity in the strong-field limit. Improving our understanding of the strong gravitational regime will help to advance models of supermassive black holes in active galactic nuclei, compact object mergers, gamma-ray burst environments, and gravitational waves. Moreover, the inner cores of neutron stars are theorized to be denser than an atomic nucleus, and these supra-nuclear-density conditions cannot be re-created in a laboratory on Earth. As such, a major goal of the field is to determine the equation of state of neutron star matter. Modelling emission from the surfaces of neutron stars can place constraints on the equation of state, with implications for the extreme limits of nuclear physics, particle physics, and condensed matter physics.

One of the best laboratories to study compact objects is low-mass X-ray binaries (LMXBs), in which the compact object accretes from a low-mass stellar binary partner. For LMXBs in outburst, viscous heating in the accretion disk causes the accreting plasma to glow brightly in the X-rays as it falls down the potential well of the compact object. The X-rays are dominated by emission from the immediate vicinity of the compact object, in the strong-field gravity regime. The two primary analysis tools to study X-ray emission from LMXBs are spectroscopy and timing. Spectroscopy informs us of the physical processes and components present in the inner region of the LMXB, while timing informs us of dynamical changes in these systems.

There is a plethora of rapid sub-second variability in the X-ray light curves from accreting LMXBs, and a growing toolbox of analysis techniques and algorithms to apply to such phenomena. The two kinds of variability that are studied in this thesis are quasi-periodic oscillations (QPOs) and coherent X-ray burst oscillations. QPOs are a probe for studying physical processes in the inner regions of LMXBs, and X-ray burst oscillations are used to determine the masses and radii of neutron stars to constrain the neutron star equation of state. Of the novel analysis techniques, the one featured most

prominently in this thesis is X-ray spectral-timing, and specifically our method for QPO-phase-resolved spectroscopy. Spectral-timing analysis combines spectroscopy and timing simultaneously to take advantage of energy-dependent Fourier amplitudes and phases in the data. With this, we can deduce the temporal relationships between spectral components and break degeneracies in QPO emission models. The other new analysis technique featured in this thesis is an evolutionary optimization algorithm. When applied to models of X-ray burst oscillations, it is able to explore the degeneracies in the model parameter space more efficiently than other popular algorithms. With these new analysis techniques, we are able to get more out of the data and gain a better understanding of compact objects.

In Chapter 2, we present a novel spectral-timing technique to do phase-resolved spectroscopy of QPOs that tracks the variations of spectral parameters with QPO phase. We apply this new technique to the fleeting Type B low-frequency QPO from the black hole GX 339–4 during its 2010 outburst. We find that the blackbody emission has a small variation ($\sim 1.4\%$ fractional rms) that leads the large power-law variation ($\sim 25\%$ fractional rms) by ~ 0.3 in relative phase. We also find that the variable blackbody is cooler and larger than the inner edge of the accretion disk. From these results, we infer that a large-scale-height ('jet-like') Comptonizing region precesses during a QPO cycle, illuminating and heating azimuthal regions of the inner accretion disk. There is debate regarding the location of the Comptonizing region, and our results suggest that the Comptonizing region in this black hole is linked to the base of the radio jet. Furthermore, we find that the QPO lag-energy spectrum can be reproduced by *periodic* variations of the low- and high- energy spectral components that are shifted out of phase with each another.

In Chapter 3, we use ray-tracing to simulate pulse profiles of thermonuclear burst oscillations from an accreting neutron star. We then fit these with an evolutionary optimization algorithm (similar to a genetic algorithm, with more generalized features) to assess how well we could recover the input parameters. This application of an evolutionary optimization algorithm to burst oscillation pulse profile modelling is the first of its kind in the literature. For the regions of parameter space sampled by our tests, the mass and radius fits are accurate (average bias compared to the true value) to $\leq 5\%$, with an uncertainty (statistical error, including degeneracies among the fit parameters) of $\leq 7\%$ in mass and $\leq 10\%$ in radius. While the parameter spaces of all the models tested have some degeneracy, we find that synthetic pulse profiles with significant asymmetry and large amplitudes are the best-constrained. If these results could be obtained with real data from a large-area, high-throughput X-ray timing instrument like *eXTP* or *STROBE-X*, they will produce definitive tests that rule on the validity of equation of state models.

In Chapter 4, we apply our phase-resolved spectroscopy technique to a lower kilohertz (kHz) QPO from the neutron star 4U 1608–52 during its 1996 outburst. We find evidence for significant variations in the shape of the Comptonized blackbody

spectral component, like a nearly 10% rms variation in the photon index. The seed blackbody temperature changes by only a few percent, giving a total QPO spectrum that pivots at low energies with the dominant variation happening in the hard X-rays. The small time lags (few tens of microseconds) cannot be reproduced by the changing amplitude of the parameter variations alone, but by including small (few percent) phase lags in the parameter variations. These results are interpreted as a 'breathing' oscillation in the neutron star boundary layer (where the accretion disk meets the neutron star surface). Finer spectral and temporal resolution of the lower kHz QPO from, e.g., *NICER* or *STROBE-X*, would allow us to precisely determine the location of the seed blackbody and the internal spectral shape changes of the boundary layer, to better understand the physics of this 'breathing' oscillation.

Finally in Chapter 5, we carry out spectral-timing analysis of the Type B low-frequency QPO seen by *NICER* in the new black hole transient MAXI J1535–571 in the autumn of 2017. We discover a weak Type B QPO in the soft spectral state of the outburst and we compute its energy-dependent cross-spectrum amplitude, lag-energy spectrum, cross-correlation function, and phase-resolved ratio spectra. The lag-energy spectrum has a similar shape but opposite sign as compared to the Type B QPO lags in Chapter 2, and the ratio spectra indeed show that the soft X-rays *lag* the hard X-rays by $\sim$27% of a QPO cycle. *NICER*'s superb CCD energy resolution of the soft X-rays tracks the tiny blackbody variations with QPO phase, and the completed spectral calibration of *NICER* in the future will allow us to fit detailed models to the QPO-phase-resolved spectra. While we find no evidence of a changing iron line profile in the ratio spectra, a larger-area X-ray observatory like *STROBE-X* would place much stricter upper limits on the presence of this reflection component in the QPO emission.

Samenvatting

Compacte objecten zoals neutronensterren en zwarte gaten met massa's tot enkele tientallen zonsmassa's zijn overblijfselen van een supernova waarbij de kern van een zware ster ineenstort. Volgens de algemene relativiteitstheorie zijn compacte objecten zwaar en dicht genoeg om de ruimtetijd significant te krommen. Door de emissie van zeer dicht bij compacte objecten te bestuderen kunnen we de effecten van sterke zwaartekracht op fysische processen ontrafelen en de algemene relativiteitstheorie toetsen in het regime van sterke zwaartekracht. Een beter begrip van dit regime is cruciaal voor het verbeteren van modellen van superzware zwarte gaten in actieve sterrenstelsels, het samensmelten van compacte objecten, de omgeving van gammaflitsen en zwaartekrachtgolven. Bovendien heeft de kern van neutronensterren een theoretisch voorspelde dichtheid groter dan die van een atoomkern, en zulke supra-nucleaire dichtheden kunnen in aardse laboratoria niet worden gereproduceerd. Een belangrijk doel van het onderzoeksveld is daarom de toestandsvergelijking van deze neutronenstermaterie te bepalen. Door de emissie van het oppervlak van neutronensterren te modelleren kan de toestandsvergelijking worden begrensd, met consequenties voor de meest extreme gebieden van de kernfysica, deeltjesfysica, en vaste-stoffysica.

Een van de beste laboratoria om compacte objecten te bestuderen zijn lage massa röntgendubbelsterren (LMXBs) waarin een compacte object massa opneemt (accretie) van een begeleidende, lage-massa ster. Bij een LMXB in uitbarsting straalt het door wrijving verhitte plasma in de accretieschijf helder in het röntgengebied, terwijl het in de potentiaalput van het compacte object valt. De röntgenstraling wordt gedomineerd door de emissie uit de directe omgeving van het compacte object, in het sterke zwaartekracht-regime. De twee primaire hulpmiddelen om de röntgenstraling van LMXBs te analyseren zijn spectroscopie en analyse van tijdvariaties (*timing*). Spectroscopie vertelt ons over de fysische processen en de componenten in de binnenste delen van de LMXB, terwijl timing ons informeert over de dynamische veranderingen in deze systemen.

Er is allerlei snelle sub-seconde variabiliteit in de röntgenlichtkrommen van accreterende LMXBs, en een groeiende gereedschapskist met analysetechnieken en algoritmen om op deze verschijnselen toe te passen. De twee soorten variabiliteit die in

dit proefschrift zijn bestudeerd zijn quasi-periodieke oscillaties (QPOs) en coherente röntgenburst-oscillaties. Via QPOs worden de fysische processen in de binnendelen van LMXBs bestudeerd, en röntgenburst-oscillaties worden gebruikt om de massa en straal van neutronensterren te bepalen en daarmee de toestandsvergelijking te begrenzen. De nieuwe analysetechniek die in dit proefschrift het meest op de voorgrond treedt is timing van het röntgenspectrum (*spectral timing*), in het bijzonder onze methode voor QPO-fase afhankelijke spectroscopie. Spectral timing combineert spectroscopie en timing en gebruikt de energieafhankelijke Fourier-amplitudes en fases in de data. Hiermee kunnen relaties tussen spectrale componenten in het tijddomein worden bepaald en ontaarde parameters in QPO emissiemodellen worden ontrafeld. De andere nieuwe analysetechniek die wordt gebruikt in dit proefschrift is een evolutionair optimalisatie-algoritme. Met deze analysetechnieken kunnen we meer informatie uit de data verkrijgen en zo compacte objecten beter begrijpen.

In Hoofdstuk 2 presenteren we een nieuwe spectral timing techniek die QPO-fase afhankelijke spectroscopie mogelijk maakt en de variaties in de spectrale parameters als functie van de QPO fase meet. We passen deze nieuwe techniek toe op de kortdurende Type B lage-frequentie QPO van het zwarte gat GX 339–4 tijdens zijn uitbarsting in 2010. We vinden een kleine variatie in de blackbody emissie ($\sim 1.4\%$ fractional rms) die ~ 0.3 in fase voorloopt op een grote ($\sim 25\%$ fractional rms) variatie in de power-law component. Ook is deze variabele blackbody koeler en groter dan de binnenrand van de accretieschijf. Uit deze resultaten leiden wij af dat een Comptonverstrooiend gebied van grote ('jet-achtige') schaalhoogte gedurende de QPO cyclus precedeert en daarbij achtereenvolgens verschillende azimutale locaties in de binnenste accretieschijf aanstraalt en verhit. Binnen het vakgebied zijn er discussies gaande betreffende de locatie van het Comptonverstrooiende gebied en onze resultaten suggereren dat dit gebied, in dit zwarte gat, is geassocieerd met de basis van de radiojet. Verder concluderen we dat de energieafhankelijke faseverschuiving van de QPO verklaard kan worden door *periodieke* variaties in laag- en hoog- energetische spectrale componenten die uit fase zijn.

In Hoofdstuk 3 gebruiken we ray-tracing om pulsprofielen te simuleren van thermonucleaire burst-oscillaties van een accreterende neutronenster. Vervolgens fitten wij deze profielen met een evolutionair optimalisatie-algoritme (een gegeneraliseerd genetisch algoritme) om te bepalen hoe goed we de ingevoerde parameters kunnen reproduceren. Dit is de eerste toepassing in de literatuur van een evolutionair optimalisatie-algoritme om burst-oscillatie profielen te modelleren. In het gebied van de parameter-ruimte bestreken door onze tests, zijn de massa en straal fits nauwkeurig (d.w.z. voor wat betreft systematische afwijkingen t.o.v. de werkelijke waarde) tot $\leq 5\%$, met een onzekerheid (statistische fout inclusief de effecten van ontaarding tussen de fit parameters) van $\leq 7\%$ in de massa en van $\leq 10\%$ in de straal. In alle geteste modellen is er een zekere mate van ontaarding tussen de parameters, maar synthetische pulsprofielen met een significante asymmetrie en grote amplitude begrenzen

de parameters het best. Als vergelijkbare resultaten kunnen worden verkregen met echte data van een groot röntgen-timing instrument dat grote fotonenfluxen aankan, zoals *eXTP* of *STROBE-X*, zal dat definitieve tests opleveren van de geldigheid van model-toestandvergelijkingen.

In Hoofdstuk 4 passen we onze fase-afhankelijke spectroscopie methode toe op een lage kilohertz (kHz) QPO van de neutronenster 4U 1608–52 tijdens zijn uitbarsting in 1996. We vinden aanwijzingen voor significante variaties in het spectrum van de Comptonized blackbody component, zoals een variatie in fotonindex van bijna 10% rms. De temperatuur van de seed blackbody varieert met maar een paar procent, wat bij elkaar leidt tot een QPO spectrum dat heen en weer draait rond een punt bij lage energie en waar de sterkste variaties plaatsvinden in het harde röntgengebied. De kleine tijdverschuivingen (van enkele tientallen microseconden) kunnen niet verklaard worden door alleen de amplitude van de spectrale variaties, maar vereisen kleine faseverschuivingen (enkele procenten) tussen de parametervariaties. Deze resultaten worden geïnterpreteerd als een zg. 'breathing' oscillatie in de boundary layer van de neutronenster (waar de accretieschijf bij het oppervlak van de neutronenster komt). Waarnemingen van de lage-kHz QPO met hogere spectrale- en tijdsresolutie, bijvoorbeeld van *NICER* of *STROBE-X*, zouden ons in staat stellen om de precieze locatie van de seed blackbody en de interne spectrale veranderingen van de boundary layer te bepalen, en zo de fysica van deze 'ademhalingsoscillatie' te begrijpen.

In Hoofdstuk 5 tenslotte, doen we een spectral timing analyse van de Type B, lage-frequentie QPO in de pas ontdekte black-hole transient MAXI J1535–571, zoals waargenomen met *NICER* in het najaar van 2017. Wij ontdekken een zwakke Type-B QPO in de soft spectral state van de uitbarsting en bepalen zijn energie-afhankelijke cross-spectral amplitude, tijdverschuiving als functie van energie, cross-correlatie functie en fase-afhankelijke ratio spectra. De tijdverschuiving als functie van energie heeft een vergelijkbare vorm, maar omgekeerd teken, vergeleken met die van de Type-B QPO in Hoofdstuk 2. De ratio spectra laten zien dat de zachte röntgenstraling inderdaad *achterloopt* op de harde röntgenstraling met $\sim 27\%$ van een QPO cyclus. *NICER*'s grandioze CCD energieresolutie in het zachte röntgenspectrum volgt de kleine blackbody variaties als functie van QPO fase op de voet, en als de spectrale calibratie van *NICER* voltooid is zullen we gedetailleerde modellen kunnen fitten aan de QPO-fase afhankelijke spectra. We vinden in de ratio spectra geen aanwijzingen voor veranderingen in het ijzerlijnprofiel, maar een groter röntgen-observatorium, zoals *STROBE-X* zal veel striktere bovenlimieten plaatsen op de aanwezigheid van deze reflectiecomponent in de QPO emissie.

Acknowledgements

While a PhD is an individual undertaking, it takes a village to get through it, and I would not have succeeded without an incredible network of support.

First, I would like to thank Phil for the opportunity to do a PhD and for teaching me so much about X-ray astronomy, Fourier techniques, and statistics. Thank you for seeing through my nerves at the interview and choosing me! Thank you also to Michiel for your additional guidance throughout the PhD. I would also like to thank my committee members, Chris, Ed, Anna, Rudy, and Sera, for evaluating my thesis and being part of the defense. Additionally, thank you to Sharon for your help and guidance in the research in Chapter 3.

Huge thanks to Lucy, Adam, and Anne, for your invaluable guidance and support as I navigated my PhD. Lucy, thank you for your patience in helping me learn the ropes as I jumped in the deep end of X-ray data analysis as a total beginner. Adam, thank you for sitting down and working through cross spectral analysis details when I kept getting stuck. Anne, thank you for our numerous illuminating chats about harmonics and cross-correlation functions. I can say with absolute certainty that I wouldn't have finished without you three. I hope to mentor others as a postdoc as you did for me!

Thank you to the X-ray group at API. I greatly appreciated both the larger X-ray timing sandwich meetings and the smaller Timing Club meetings. Thanks also to the radio pulsar/FRB group for letting me crash some of your meetings in my last year. Having a group atmosphere definitely enhanced my research experience at API.

More huge thanks to Cátia and Marieke, my paranymphs. Cátia, my PhD sister, thank you for frantic research advice, dinners, code checks, tea breaks, fashion nights out, sewing afternoons, and occasional haircuts. Was gezellig!! Marie(ke), my longtime office buddy and Nerd Nite co-conspirator, thank you for chats, tea breaks, birthday decorations, and research advice on my postdoc applications. You both infused so much fun and work-life balance into my PhD journey, and I'm so glad to have you by my side at the defense.

Thank you so much to Daniela, Lucy, and Caroline, my first close group of friends in Amsterdam, for your advice and so much tea that kept me on track to finish this

PhD. Our dinners and get-togethers helped me feel at home here. Daniela, thank you also for helping me with my postdoc applications! Thank you to Arnim for the pizza and movies. Thank you to Samayra, Amruta, Susan, and Cátia, the "goat girls", for brainstorming sessions, more advice, even more tea, lunches, dinners, and trips to the geitenboerderij for baby goat snuggles. Susan, thank you also for our monthly bagel brunches. Thank you to Anne for so much chocolate and tea, and many long chats. And of course, thank you very much to Samayra for translating the samenvatting.

Thank you to my wonderful friends and colleagues from Nerd Nite (Shanna, Lisa, Marieke, Milena, Eva, and Brechje) and the inaugural API PhD & PD Council (Samayra, Macla, and Justin) for broadening my horizons. I'm so glad to have been a part of those with you all! Thank you to my amazing Stingray buddies, Daniela and Matteo, for being the best software collaborators. I love how much Stingray has grown!! Thank you also to the broader python in astronomy folks – I'm so glad to be involved with building an inclusive and welcoming community for astronomy research software!

Thank you to my lovely officemates over the years in the big office and the little office at the API. I've loved our many conversations and tea breaks and vocabulary-searching sessions over the years. Thank you to the API secretariat and manager (Milena, Susan, Vivian, and Annemarie) for being the driving force of the institute and for many lovely chats. Thank you also to the APIs I've had the pleasure of working alongside over my 4+ years in Amsterdam! Our institute is a gem and I'm so glad to have been here. Additionally, thank you to the Huber group in Regensburg for your hospitality when I was in town.

Thank you to everyone in the broader astronomy community who I've had the joy of interacting with. Having such a great network of collaborators and conference buddies really makes this job fun. Thank you also to the online astro community on Twitter for adding a whole new dimension to the PhD experience!

Finally, thank you to my family for your love and support through these last 4+ years, and thank you to Tyler for being my rock through it all. I've had so much fun exploring Europe with you, and I'm so excited for us to build a life together in Michigan. The long-distance sacrifice paid off.

It's hard to succinctly write down magic words that can somehow encompass everything I'm thankful for and everyone I'm thankful to. I'm just immeasurably glad that you all were a part of my PhD.

Abbie