UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL AND APPLIED SCIENCES

Physics & Astronomy

Novel Spectral-Timing Methods for X-ray Binary Variability Studies

by

Pablo M. Cassatella

Thesis for the degree of Doctor of Philosophy

October 17, 2012

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
Physics & Astronomy

Doctor of Philosophy

NOVEL SPECTRAL-TIMING METHODS FOR X-RAY BINARY VARIABILITY STUDIES

by Pablo M. Cassatella

This work explores the connection between the spectral properties of the X-ray emission produced by black hole X-ray binaries (BHXRBs) and the variability properties of these sources that are embedded in lightcurves. While BHXRBs are mostly studied using 'static observables' such as the average energy spectrum, a variability timescale-resolved study (in terms of Fourier-frequencies) provides a much stronger and more reliable tool for understanding the intricate physics and geometrical properties of accretion flows around stellar-mass black holes.

In the present Thesis, we first show how the geometrical properties of the accretion disc can be mapped by combining spectra and frequency-resolved time-lags. A time-scale-dependent reflection model is developed in terms of Fourier-frequencies and compared to observational hard-state data of GX 339–4 obtained with *XMM-Newton*. Secondly, we explore the variability properties of the accretion disc in SWIFT J1753.5–0127 using a combination of novel spectral-timing methods and the soft X-ray coverage of *XMM-Newton*. Finally, we show the development of a novel technique for fitting correlated signals in the Fourier domain as a function of energy and Fourier-frequency that can prove particularly useful with high signal-to-noise datasets from future X-ray missions.

Contents

List of Figures

List of Tables

To my grandmother, for her love, and to Rosa, for her courage.

Declaration of Authorship

I, Pablo M. Cassatella, declare that the thesis entitled *Novel Spectral-Timing Methods for X-ray Binary Variability Studies* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
  - Joint spectral-timing modelling of the hard lags in GX 339-4: constraints on reflection models, 2012, MNRAS, 422, 2407
  - Accretion flow diagnostics with X-ray spectral-timing: the hard state of SWIFT J1753.5–0127, accepted for publication in MNRAS (arXiv:1208.6236)
  - The causal connection between disc and power-law variability in hard state black hole X-ray binaries, 2011, MNRAS 414, L60

Signed: .....

Date: .....

Acknowledgements

I would like to thank my supervisors Phil Uttley and Tom Maccarone. Phil, thanks for the time that you spent teaching me how to learn new information about the surroundings of black hole binaries from apparently boring lightcurves and for your patience correcting me whenever I was wrong. I feel so lucky having been able to learn such amazing techniques from you. Thanks also for sending me to 23 conferences, schools and meetings in the last 3 years, 3 months and 3 days (the Hawaii meeting was remarkably interesting), and for allowing me to organise the variability meeting in Southampton and co-organise the timing school in Sardinia. I gained very valuable transferable skills during the whole PhD process! **Thanks Tom** for your encyclopædic knowledge of accretion and your patience in many aspects during the last 10 months, which were long and difficult at times. **Thanks to Jörn Wilms and Sera Markoff** for organising the ITN 215212 'Black Hole Universe" network that funded my studies, travel, conferences, schools equipment and the variability workshop, and thanks to the EU tax-payers that made this possible. Jörn, thanks for your hospitality in Bamberg during two weeks, and for all your help. Tomaso Belloni, together with Jörn, was my other co-supervisor. **Thanks Tomaso** for teaching me about RXTE and all your knowledge of the variability properties of probably every source on Earth (well, on the sky). The Network ESRs and ERs made life more amenable, and I am very grateful for their support (Refiz, el niño está corriendo). **I am extremely thankful to John Davis and John Houck** for their sincere help when I visited them at MIT, and their help with ISIS and S-Lang. I spent three of the best weeks of my life working there (the coffee shop in the corner and the chai latte in the kitchen provided enough calories during this time). Mike Nowak's humour while I was there is something I won't forget either. **The help I got from Tony Wilkinson, my big-brother academically-wise, is invaluable.** He made the first two years of my PhD much more fun, and I learned much from our discussions over afternoon-coffee. Tony, coffee, 2pm? No sorry, I can't, 3pm? He would, nearly always, say yes. I am sorry that most of my ideas were completely silly. Thanks for teaching me some important vocabulary in life, such as *tart*, the meaning of which is not completely clear to me yet. **My deep gratitude goes to Manfred Hanke,** who helped me with my first steps into the world of S-Lang, ISIS, and SLxfig, and who helped me to correct some inaccuracies in this Thesis. Communicating via email with you is always a pleasure, and often a lot of fun. **Thanks to my (other) friends spread around the world** (UK, Germany, US, Italy, Spain...), for their support, especially to my Southampton friend Massoud Mokhtarpour, who is always there if I need anything, and who enlightened me with his knowledge of

real life and stuff outside of Physics. **And to my housemates Tinesh and Hang** for their funny disputes about ice cream and chocolate, and the intensive food-centred weekends. The time I spent at the *XMM-Newton* Science Operations Centre working with Matteo Guainazzi prior to starting my PhD was exceptionally useful and challenging. **Thanks Matteo, and thanks to all of the SOC staff and the trainees!** **Gracias a mi familia por su apoyo incondicional** y por estar siempre pendientes de cualquier cosa que pueda necesitar. Y a mi hermano Giovanni por enseñarme tanto de la vida. And finally, 謝謝美君. You relentlessly supported me every time my motivation was bounded by ϵ (especially when $\epsilon \rightarrow 0$), which resulted in my moaning and your taking care of me, which was all I wanted.

CHAPTER 1

Introduction

This introductory chapter is aimed at presenting a review of the current understanding of black hole X-ray binaries, what they are, why outbursts in these systems occur and how their emission is affected. I also introduce the properties of their spectral emission, a common model for explaining the emission variations and discuss their time variability. Finally, I introduce the Fourier methods that will be used throughout the present Thesis.

1.1 Black Hole X-ray Binaries

1.1.1 Accretion

Black hole X-ray binaries are formed by a stellar-mass black hole ($M \sim 3-10\ M_{\odot}$) that lives in a gravitationally-bound system together with a companion star. The gravitational (Roche) potential in binary systems was derived by Edouard Roche in the 19th Century and is defined as (Frank et al., 2002):

$$\Phi_R(\mathbf{r}) = -\frac{GM}{|\mathbf{r}-\mathbf{r}_1|} - \frac{GM_2}{|\mathbf{r}-\mathbf{r}_2|} - \frac{1}{2}(\omega \mathbf{e} \times \mathbf{r})^2 \quad , \tag{1.1}$$

where $\mathbf{r}_1$ and $\mathbf{r}_2$ are the vectors towards the centre of the two objects, $\mathbf{e}$ is a unity vector normal to the orbital plane, and ω is the angular velocity of the binary system with respect to an

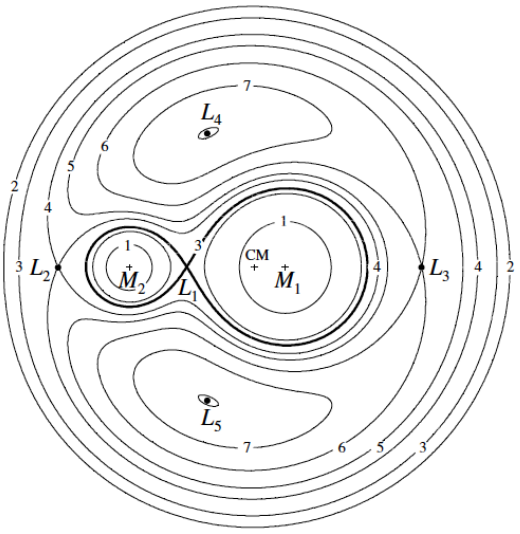

Figure 1.1: Equipotential surfaces and Lagrangian points in the case of a Roche potential with a mass ratio $M_2/M = 0.25$, from Frank et al. (2002).

inertial frame, $\omega = \left(\frac{G(M+M_2)}{a^3}\right)^{1/2}$ (a is the binary separation, M and M_2 are the masses of the black hole and companion star).

The production of X-rays in these systems is caused by accretion of matter that leaves the companion star and is attracted towards the black hole, as we shall see. The Roche potential has two imaginary volumes centred at the position of the two objects called Roche lobes, where matter is gravitationally bound to the respective object. Linking the two Roche lobes there is the Lagrangian point L1 that corresponds to a saddle point, or point of meta-stable equilibrium, in the Roche potential. In total, there are five Lagrangian points that correspond to saddle points in the Roche potential (see Fig. 1.1).

The potential in Eqn. 1.1 follows the Euler equation for conservation of momentum for each gas element in the co-rotating frame, taking into account the centrifugal and Coriolis forces:

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v} = -\nabla\phi_R - 2\omega\mathbf{e}\times\mathbf{v} - \frac{1}{\rho}\nabla\mathbf{P} \quad . \tag{1.2}$$

In this equation, $\mathbf{v}$ is the velocity field of the accreting gas of density ρ and $\nabla\mathbf{P}$ is the pressure gradient that produces the transfer of momentum from the companion star to the accretion disc that forms around the black hole.

Mass transfer from the companion star to the black hole can be produced via two main mechanisms (Frank et al., 2002):

- Wind accretion, in which case the companion star is an early-type, red giant or AGB star with a strong stellar wind that is transferred to the black hole via L1 or L2. The typical mass loss is $\sim 10^{-6} - 10^{-5} M_\odot$/yr.

- Roche overflow. This happens when the companion star fills the Roche lobe, transferring mass via L1, with a typical mass transfer rate $\sim 10^{-11} - 10^{-8} M_\odot$/yr.

The mass that enters the Roche lobe of the black hole has too high an angular momentum to be accreted directly onto the black hole. Instead, it spirals around it until its trajectory becomes quasi-circular and a flattened structure, the accretion disc, is formed. Angular momentum has to be conserved, hence there has to be a transfer of angular momentum outwards caused by internal torques. The disc can then be approximated as the sum of concentric disc annuli, each of which has a Keplerian angular velocity

$$\Omega_K(R) = \left(\frac{GM}{R^3}\right)^{1/2} \tag{1.3}$$

at a radius R for a black hole mass M. The gravitational constant $G \approx 6.674 \times 10^{-8}$ cm^3 g^{-1} s^{-2}.

The gravitational potential energy of a test mass ΔM orbiting around the black hole is

$$E_{GR}(R) = -\frac{GM\Delta M}{R} \quad . \tag{1.4}$$

When the mass travels from infinity to a distance R from the black hole, it releases an energy $\Delta E_{GR} = \frac{GM\Delta M}{R}$, half of which is converted to rotational kinetic energy, the other half being radiated away in the form of a black-body (the gas is optically thick). If the black hole accretes steadily at a rate $\dot{M}$, the luminosity produced by accretion is:

$$L_{acc} = \frac{1}{2}\frac{GM\dot{M}}{R} \quad . \tag{1.5}$$

L_{acc} can be parameterised as $L_{acc} = \eta \dot{M} c^2$ (c speed of light), in which case it is easy to show that accretion is the second most efficient way of converting mass into energy (after annihilation), even beyond nuclear fusion of hydrogen into helium ($\eta \sim 0.057 - 0.42$ for the first case, depending on the spin of the black hole, compared with a poor $\eta \sim 0.007$ for the second).

In the case (and rather common approximation) of spherical accretion of fully ionised hydrogen, the maximum luminosity that can be reached is given by the balance between radiation pressure and gravity:

$$\frac{GMm_p}{r^2} = L_{Edd}\frac{\sigma_T}{4\pi c r^2} \quad . \tag{1.6}$$

Solving for L_{Edd} gives

$$\frac{4\pi GMm_p}{\sigma_T} \approx 1.3 \times 10^{38}\frac{M}{M_\odot}\text{erg/s} \quad , \tag{1.7}$$

where $\sigma_T = 6.652 \times 10^{-25}$ cm^2 is the Thomson cross-section and $m_p = 1.67 \times 10^{-24}$ g $= 938.27\ MeV/c^2$ is the proton mass. The corresponding Eddington mass-accretion rate can be written as $\dot{M}_{Edd} = \frac{L_{Edd}}{\eta c^2}$.

Under the predictions of General Relativity, black holes modify the metric of the space-time around them. The event horizon around the black hole is called the Schwarzschild radius R_s:

$$R_s = 2\frac{GM}{c^2} \approx 2953 \frac{M}{M_\odot} \text{ m} = 295300 \frac{M}{M_\odot} \text{ cm} \quad . \tag{1.8}$$

(At the speed of light, a photon can travel a Schwarzschild radius in $9.9 M/M_\odot\ \mu s$.) Accreting matter inside the Schwarzschild radius for non-rotating black holes (and inside half a Schwarzschild radius for rotating, Kerr black holes) cannot interact with the outside, therefore radiating matter beyond this point cannot be observed. The GR metric starts deviating from Newtonian at a few R_s, and $3R_s$ is last radius where accretion can be circular (this radius is also called the Innermost Stable Circular Orbit, or ISCO) for a non-rotating Schwarzschild black hole. Matter at radii below $3R_s$ reaches R_s after only $\sim (R_s^3/GM)^{1/2} \sim 10^{-4}$ s (Frank et al., 2002).

The emitted luminosity per unit area of the optically-thick accretion disc equals the energy lost via black-body radiation, giving:

$$T(R) \approx \left(\frac{GM\dot{M}}{8\pi\sigma R^3}\right)^{1/4} \propto \left(\frac{\dot{M}}{\dot{M}_E}\right)^{1/4} \left(\frac{M}{M_\odot}\right)^{-1/4} \left(\frac{R}{R_s}\right)^{-3/4} \quad . \tag{1.9}$$

The internal torques that conserve the total angular momentum while keeping the steady accretion responsible for the production of luminosity as in Eqn. 1.5 are called viscous torques. Eqn. 1.3 implies that inner radii rotate faster than outer radii. This will inevitably produce shear viscosity, resulting in heat –radiation– being produced, while driving angular momentum outwards.

The origin of viscosity is unclear, yet it should allow the following conditions to be satisfied (King, 2012):

$$\frac{\partial}{\partial R}(R^2\Omega) > 0, \qquad \text{and} \quad \frac{\partial \Omega}{\partial R} < 0 \quad . \tag{1.10}$$

(The increase of angular momentum for larger radii corresponds to a decrease of angular velocity.) A potential candidate model for explaining shear viscosity that obeys Eqn. 1.10 and is stable against axisymmetric perturbations resides in the magnetic dragging model proposed by Balbus and Hawley (1991). In this model, the disc is threaded by a weak vertical magnetic field that simultaneously acts towards reaching rigid rotation and towards returning elements of the disc to their original locations. Much progress has been made in recent

years thanks to 3D magnetohydrodynamic simulations (both local and global), and convergence tests are now being done (see e.g. Shiokawa et al. 2012; Sorathia et al. 2012 for more information).

Historically, the uncertainty in the exact physics of viscous stresses has been condensed in the parameter α that is widely present in contemporary literature. Shakura and Sunyaev 1973 (see also Novikov, I. D. and Thorne, K. S. 1973) proposed the α-prescription,

$$\nu_{\rm visc} = \alpha H c_s \quad , \tag{1.11}$$

where c_s is the sound speed, H is the disc thickness, $\nu_{\rm visc}$ is the kinematic viscosity, and α is a unitless parameter, typically $\alpha < 1$ (Frank et al., 2002). In the case of Keplerian motion, the torque exerted by an outer ring on an immediately inner ring by viscous stress in a disc of constant surface density Σ equals (Kolb, 2010):

$$\tau_{visc} = -3\pi\nu_{visc}\Sigma(GMR)^{1/2} \quad . \tag{1.12}$$

The torque of the inner ring on the outer ring has the same magnitude but opposite sign.

The disc is often assumed to be geometrically thin (i.e. its height z at a radius R satisfies $z/R \ll 1$). The vertical gravitational acceleration at a radius R is

$$g = \frac{GM}{r^2}\cos\theta = \frac{GM}{r^2}\frac{z}{r} \approx \frac{GM}{R^3}z \quad (\text{z} \ll \text{R}) \quad , \tag{1.13}$$

where r is the distance between the black hole and the element of the disc at radius R and height z above the midplane, and $\cos\theta = z/r$. The vertical pressure gradient $\frac{\mathrm{d}P}{\mathrm{d}z}$ in a disc of density ρ balances the gravitational force in the vertical direction,

$$\frac{\mathrm{d}P}{\mathrm{d}z} = -\rho g \quad . \tag{1.14}$$

If the disc is isothermal in the vertical direction, then the pressure is given by $P = \rho c_s^2$ (c_s is the sound speed). In the approximation of a thin disc, it can be found that $\frac{h}{R} \approx \frac{c_s}{v_\phi}$ (where h is the disc scale-height). Thin discs are therefore highly supersonic.

1.1.2 Outbursts and states

Black hole X-ray binaries show outbursts that last weeks, months or even years (Chen et al., 1997) (see Fig. 1.2 for the lightcurve showing some of the most recent outburst of GX 339–4).

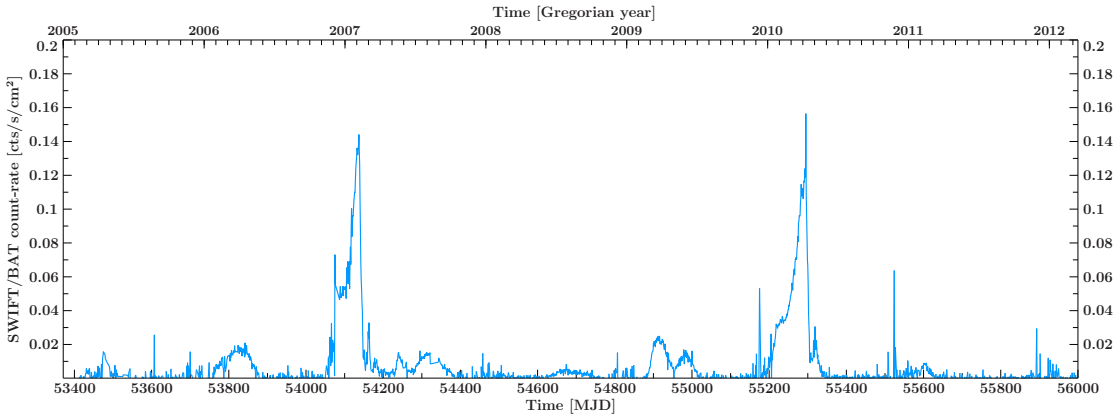

Figure 1.2: Historical lightcurve of GX 339–4 taken with *Swift*/BAT showing the most recent outbursts of the source.

These cycles show a rapid increase of X-ray luminosity produced by the trigger of the hydrogen ionisation instability (Smak, 1984; Dubus et al., 2001; Lasota, 2001). When the process of accretion builds up a sufficient mass, the disc reaches a temperature of 10^4 – 10^5 K at a local radius in the disc. At these temperatures[1], the hydrogen in the disc starts ionising, increasing sharply an otherwise low opacity. The high opacity means that the photons are less likely to escape the gas. The higher-energy photons (in the Wien tail of the black-body distribution) ionise the medium even further, producing a runaway ionisation, until most of the hydrogen is ionised. As the viscous stress is linearly proportional to temperature in the α-prescription, the increase in temperature triggers an increase in the mass-accretion rate, and the material is eaten away. This process continues until the outer disc temperature drops below the hydrogen ionisation temperature, and the disc is cooled inwards, returning back into quiescence.

The rise and later decay of an outburst imply a rise and a fall in luminosity, yet the overall X-ray luminosity is unevenly distributed among all energies, but redistributed among different spectral components that dominate the soft and hard bands at each stage of the outburst. Generally-speaking in the *XMM-Newton* era, these bands correspond roughly to 0.5 – 2.0 keV and 2.0 – 10.0 keV, respectively. This behaviour gives rise to a well-known hysteresis cycle in the hardness-intensity plane, that can be plotted in a hardness-intensity diagram (HID) (Fig. 1.3).

For such a hysteresis to take place, there needs to be (at least) one other spectral component. This additional component is in fact a power-law with a cut-off at high energies (at least in the hard state) $E \sim 50 - 150$ keV whose relative intensity and slope vary throughout the outburst (Motta et al., 2009). At low fractions of the Eddington luminosity, a hard power-law ($\Gamma \sim 1.4 - 2.0$) dominates the spectrum and the thermal component may be absent or weak in the X-rays. At higher luminosities, as the disc emission becomes stronger, the power-law becomes

[1] If the disc is in thermal equilibrium at a given radius, a rough estimate of the temperature at which a certain fraction of ionised hydrogen can be inferred from Saha's equation for a given hydrogen number density.

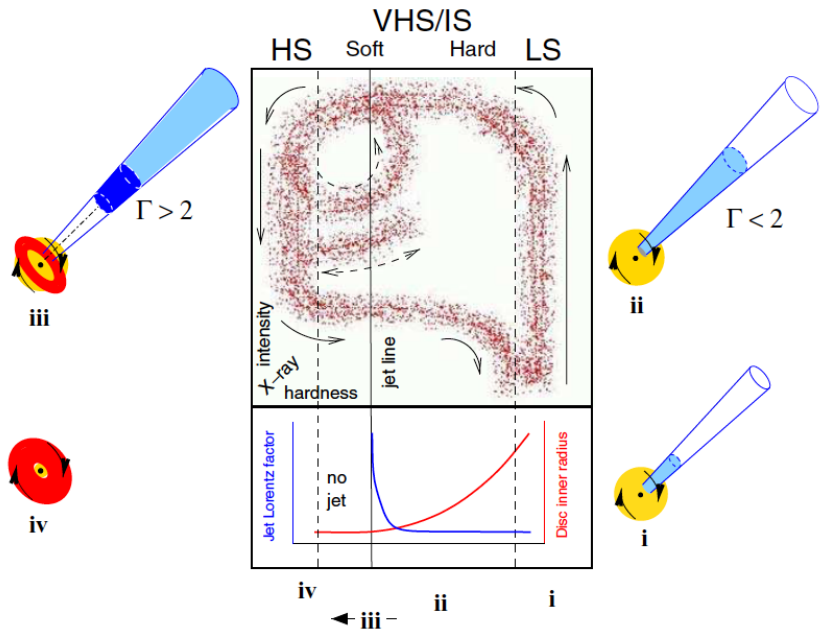

Figure 1.3: Schematic example of a hardness-intensity diagram (LS = low/hard state, HS = high/soft state, VHS = very high state, IS = intermediate state), shown in Fender et al. (2004), and its connection with the appearance of jets.

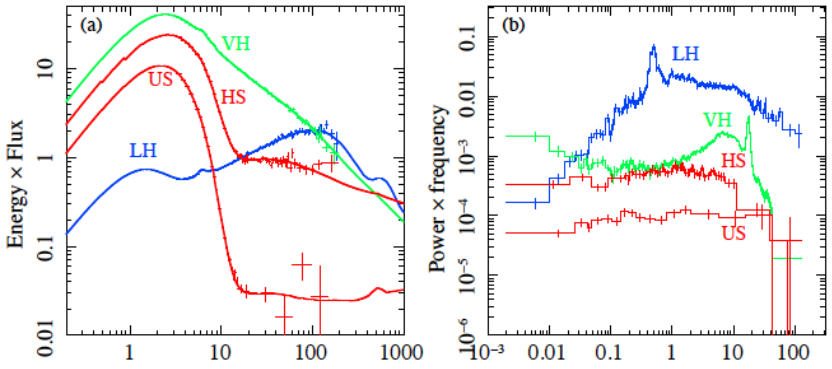

Figure 1.4: Spectra of hard, intermediate and soft states of GRO 1655–40 (left) and their relative power spectral densities (right), from Done et al. (2007). The soft state (HS, or high/soft) is characterised by a strong disc black-body component and a soft power-law at higher energies, whereas the hard state (LH, or low/hard) shows little disc emission, and a much harder power-law. Intermediate states, not discussed in this Chapter, are also shown. An important characteristic of the transition from the hard to the soft state is a reduction of the variability amplitude and a shift of the characteristic frequencies of variability. The latter effect is seen better in Fig. 1.9.

softer ($\Gamma \gtrsim 2$), see Fig. 1.4. The relative amplitudes of the spectral components give rise to different states (e.g. Homan et al. 2001).

The change in relative aplitude of these two spectral components can be understood in the context of two different flows, an outer optically-thick accretion flow (the accretion disc) and an inner optically-thin inner flow. This optically-thin, geometrically-thick inner flow ($H/R \sim$

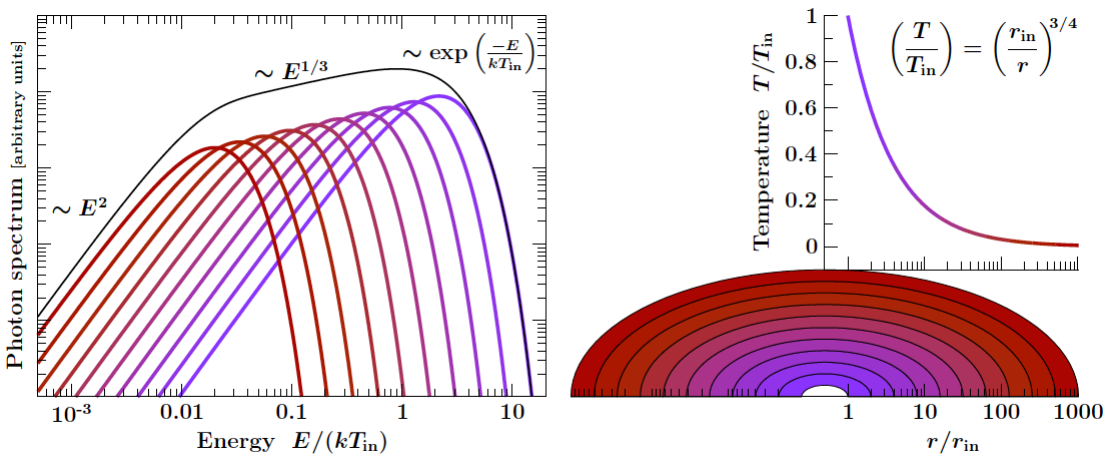

Figure 1.5: Spectrum resulting from the superposition of black-body components originating from individual rings in the optically-thick accretion disc, following a $T(R) \propto R^{-3/4}$ relation (see Eqn. 1.9). Taken from Hanke (2011).

0.3 – 0.4) appears only at low luminosities and may responsible for the power-law emission (see Done et al. 2007 and references therein).

1.1.3 Spectral components

1.1.3.1 Thermal emission

An optically-thick medium where photons are in thermodynamic equilibrium at a temperature T Subsec. 1.1.1 emits a black-body spectrum. The intensity of the emission produced by a black-body follows Planck's law:[2]

$$B(\nu, T) = \frac{2h\nu^3}{c^2} \frac{1}{e^{\frac{h\nu}{kT}} - 1} \quad , \tag{1.15}$$

where $k = 1.38 \times 10^{-16}$ erg/K is Boltzmann's constant. The optically-thick accretion disc can be approximated as a sum of concentric rings, each of which has a local equilibrium temperature $T(R)$ and an area $2\pi R\mathrm{d}R$. The sum of the black-body distributions at each ring produces a characteristic distribution, often referred to as the 'multicolour disc black-body', see Lynden-Bell (1969) for its derivation. The resulting spectral shape can be found in Fig. 1.5.

1.1.3.2 Power-law emission

The power-law component is likely produced by Compton up-scattering by hot electrons of seed photons that originate in the accretion disc. In the case of hard states ($\Gamma \lesssim 2.0$), the

[2] The Cosmic Microwave Background is a perfect example of such dependence.

Comptonising medium is a plasma of electrons in thermal equilibrium with a temperature kT_e ~50 – 150 keV, often referred to as the 'corona' (for softer states, the lack of a high energy cut-off suggests a non-thermal, power-law distribution of electrons).

If the electrons were at rest, the relative change in photon energy after one scattering, in the isotropic case, would be given by:

$$\frac{\Delta E}{E} \approx -\frac{E}{m_e c^2} \quad , \tag{1.16}$$

where $m_e = 9.10 \times 10^{-28}$g = 0.511MeV/c^2 is the mass of the electron. In the case of non-relativistic thermal Comptonisation expected in the hard state, the electrons are expected to follow a Maxwellian distribution with a temperature much smaller than their rest-mass energy, $kT_e \ll m_e c^2$, or $\Theta = \frac{kT_e}{m_e c^2} \ll 1$. In this case, the relative energy change of the electrons is:

$$\frac{\Delta E}{E} \approx -\frac{4kT_e - E}{m_e c^2} \quad , \tag{1.17}$$

(Rybicki and Lightman, 1979), and the average amplification of the photon energy is $A = 1 + 4\Theta$. The probability of scattering depends on the number of electrons in the scattering volume (optical depth), $\tau = nR\sigma_T$, where n is the electron number density, σ_T is the Thomson scattering cross-section and $R\sigma_T$ is the volume where the scattering may take place, with a probability $e^{-\tau}$. The average energy increase of a photon travelling in a cloud of electrons is given by the total number of scatterings $\max(\tau, \tau^2)$ (which depends on the random walk that the photons follow in the medium) times the average energy increase per scattering $A = 1 + 4\Theta$, and is given by the Compton parameter y:

$$y = \frac{4kT_e}{m_e c^2} \max(\tau, \tau^2) \quad . \tag{1.18}$$

When the photon energy is much smaller than the electron temperature and occurs in an optically thin cloud, Compton upscattering gives an intensity with a power-law slope that derives from multiplying the seed photon intensity times the probability of the photons reaching energy E after $k = \log(E/E_0)/\log(A)$ scatterings.

1.1.3.3 Reflection spectrum

A good fraction of the photons that undergo Compton up-scattering in the 'corona' are intercepted by the much colder accretion disc. The exact fraction is difficult to quantify, but in an Euclidean metric and assuming isotropic scattering, the number of photons that reach the disc is proportional to the solid angle subtended by the disc as seen by the 'corona'. While

about 60 – 70 per cent of the illuminating flux is absorbed and contributes to heating the disc (Wilkinson and Uttley, 2009), the rest is observed in the form of a characteristic reflection spectrum. For photon energies below ~ 10 keV, the cross-section for photo-absorption ($\propto Z^5 E^{-3}$) is larger than the cross-section for Comptonisation, hence most of the photons at these energies are absorbed in the disc (Done, 2010). Photons with higher energies are Compton down-scattered a number of times until they leave the medium or reach the above energies below which photo-absorption becomes much more probable. Compton down-scattering produces the 'Compton hump', first predicted by Lightman and White (1988), usually between 20 and 50 keV.

The photo-absorbed photons can be re-emitted in the form of fluorescent lines. Iron Kα lines are often observed, and range in energy between 6.4 (FeI) and 6.97 keV (FeXXVI) depending on the ionisation state[3]. Fluorescent lines are produced when the incoming photon is absorbed by an electron in an inner shell (e.g. K). The vacant 'Pauli slot' can be filled by an electron from an outer shell (e.g. L), emitting a photon. The emission of a photon following the transition $L \rightarrow K$ happens with a probability given by the 'fluorescent yield' which depends on the ionisation state of the iron (for Kα, it is almost constant up to FeXVII). The combination of photo-electric absorption, Compton down-scattering and fluorescent line emission gives rise to the so-called reflection spectrum.

The ionisation state of the disc can be parameterised with $\xi = \frac{4\pi F}{n}$ where F is the incident flux, and n is the hydrogen number density. For high ξ, the reflection spectrum resembles a power-law, not easily distinguishable from the power-law responsible for the primary Comptonised emission. Collisional ionisation can also contribute to an increased ionisation state (Masai, 1997).

1.2 Fourier methods applied to BHXRBs

Fourier[4] methods are the fundamental tool to study variability of black hole X-ray binaries as a function of variability frequency, providing an important advantage over the study of variability in the time domain, especially when two or more sources in the same state have to be compared, or e.g. in the case where a new variability model has to be validated by comparing it to observational data. In both cases, it will be considerably easier to compare variability amplitudes and phases as a function of frequency than matching noisy lightcurves with arbitrary phases.

[3] Kβ lines may be observed too, although these may blend with Kα lines due to relativistic blurring that is discussed in this Subsection.

[4] Named after the French mathematician and physicist Jean Baptiste Joseph Fourier (1768 – 1930).

A very important assumption for lightcurves in the context of BHXRBs is that of stationarity. This means that, given two arbitrary lightcurve segments, the statistical moments of the individual segments are equal. As shown in this Introduction, BHXRBs can have very different spectral properties, hence this condition can only be applied when considering lightcurves of a length comparable to the typical observation time (~ few hours).

In the following, I shall introduce the Discrete Fourier Transform, its properties and applications.

1.2.1 The Discrete Fourier Transform

In the X-ray domain, a lightcurve $x_{h,k}$ is a discrete quantity that equals the number of counts $c_{h,k}$ that are recorded by a detector within the semi-open time interval $[t_k, t_{k+1})$ whose energy falls within the detector channel h,[5] divided by the element of time resolution that is used for the analysis in seconds, hence $x_{h,k} = c_{h,k}/\Delta t$. For an ideal detector[6] and in absence of background, a detector count in a channel h corresponds to a source photon hitting the detector with an energy in the range $[E_h^{\min}, E_h^{\max}]$, where $E_h^{\min}$, $E_h^{\max}$ are the lower and upper boundaries of channel h, respectively, defined in the instrumental Ancillary Response File provided with each observation.

The Discrete Fourier Transform (DFT) $\mathscr{F}(\cdot)$ of the signal $x_{h,k}$ can be computed thus (Press, 1992):[7]

$$\hat{X}_{h,l} = \mathscr{F}[x_{h,k}]_l = \sum_{k=0}^{N-1} e^{i\frac{2\pi kl}{N}} x_{h,k} = \sum_{k=0}^{N-1} x_{h,k}\left[\cos\left(\frac{2\pi kl}{N}\right) + i\sin\left(\frac{2\pi kl}{N}\right)\right] \quad , \tag{1.19}$$

where $k \in [0, N-1]$ is the partition of the time interval T used for the analysis, such that $T = N\Delta t$. In addition, $l \in [-\frac{N}{2}, \frac{N}{2}]$ and therefore the corresponding discrete Fourier frequency $\nu_l = \frac{l}{N\Delta t}$ with a step $\delta\nu = \frac{1}{N\Delta t}$. The value $\nu_{N/2} = \nu_{\rm Nyq}$ is called the Nyquist frequency, and corresponds to the maximum frequency that can be sampled for the above values of N and Δt.

A property of the DFT with many applications is linearity: given the constants a and b and two signals u and v,

$$\mathscr{F}(au + bv) = a\mathscr{F}(u) + b\mathscr{F}(v) \quad . \tag{1.20}$$

[5] This discussion can be generalised to a binned set of channels, see Chapter 4 for more details.

[6] In an ideal world, detectors would have a response $R(E,h) \propto \delta(E - (E_h^{\min} + E_h^{\max})/2)$

[7] A very detailed and comprehensive overview of Fourier transforms by Alan V. Oppenheim can be found as a collection of MIT video lectures at: `http://ocw.mit.edu/resources/res-6-008-digital-signal-processing-spring-2011/`.

The important Parseval's theorem states that the power contained in a signal is invariant to the Fourier transform. In discrete terms,

$$\sum_{k=0}^{N-1} |x_{h,k}|^2 = \frac{1}{N} \sum_{l=-N/2}^{N/2} |\hat{X}_{h,l}|^2 \quad . \tag{1.21}$$

An appropriate choice of the time bin size Δt is necessary to avoid leakage/aliasing; in particular Δt must be an integer multiple of the detector time resolution or frame time.

The DFT then allows discrete, time-dependent signals to be studied as a function of the inverse of time, frequency, and is therefore very important for the analysis of stationary signals where signal frequencies are directly related to the physics of the objects under study and the variability time-scales.

In general, because of the nature of X-ray detectors, an observed signal $x_{h,k}$ will be formed by the 'true', underlying physical signal $s_{h,k}$ plus a counting noise Poissonian deviation $n_{h,k}$ such that

$$x_{h,k} = s_{h,k} + n_{h,k} \quad . \tag{1.22}$$

Thanks to the linearity property in Eqn. 1.20, the separation between signal and noise is also possible in the Fourier-frequency domain: $\hat{X}_{h,l} = \hat{S}_{h,l} + \hat{N}_{h,l}$.

1.2.2 Power spectral density

For any given detector channel, the variability amplitude carried by astronomical signals at each Fourier frequency is often quantified in the form of the power spectral density

$$P(\nu_l) = A\frac{2\Delta t}{N}|\hat{X}(\nu_l)|^2 \quad , \tag{1.23}$$

where A is a normalisation constant, as we shall see. Examples of power spectral densities extracted in this Thesis can be found in Figs. 2.9, 3.4 and 4.8. Because astronomical lightcurves are real-valued, their Fourier transform is complex and antisymmetric around 0 ($\hat{X}_l = \hat{X}_{-l}$). Hence values corresponding to negative Fourier frequencies provide no information and can be dropped. The zero-frequency transform $\hat{X}_0$ is often referred to as the DC component as it represents the non-variable part of the signal, and is of no practical interest for variability studies. In addition, the presence of a constant signal at 0 frequency may leak to neighbouring frequency bins (ν_1, ν_2, etc). Therefore, the mean $\langle x_k \rangle$ is substracted from the signal before the DFT is computed to avoid this effect.

The factor 2 above therefore ensures that all the original signal power is conserved ($|\hat{X}(\nu_l)|^2 = |\hat{X}(\nu_{-l})|^2$) and A is a normalisation constant. If A = $1/\langle x_k \rangle$, the sum $R_{a-b} = \sum_{l=a}^{b} P(\nu_l)\Delta\nu$ gives the total rms amplitude between ν_a and ν_b (the Leahy normalisation), whereas A = $1/\langle x_k \rangle^2$, R_{a-b} gives the *fractional* rms amplitude for the same range. The latter normalisation is used throughout this thesis and is referred to as the Miyamoto or Belloni normalisation.

For a noise process, each of the real and imaginary parts of the Fourier transform in Eqn. 1.19 is independent and identically distributed following a Normal distribution centred at 0. The power density $P(\nu_l)$ is then χ^2-distributed with 2 degrees of freedom (an exponential distribution with $\lambda = 1/2$) and therefore, for a given mean M, the variance is M^2.

The above signifies that the standard deviation of the power spectrum is not sensitive to the length of the observation, hence it is not a reliable measure of the 'true' underlying power spectrum caused by the physical process that is taking place (i.e. the observed power spectrum is not a consistent estimator of the 'true' power spectrum). In order to overcome this limitation and decrease the variance, lightcurves are usually split into m consecutive segments of the same length, and then averaged, so that the new standard deviation will decrease as $1/\sqrt{m}$. This is called the Bartlett method (Bartlett, 1948). An additional method for decreasing the variance is to rebin $\hat{X}_{h,l}$ in frequency. This can be done e.g. logarithmically in frequency (see e.g. Nowak et al. (1999)).

The presence of Poisson counting noise in the Fourier transform of a lightcurve was mentioned in Sec. 1.2. The variance of the noise contribution due to the counting nature of a lightcurve is simply $c_k/\Delta t^2$ (where $c_k = x_k \Delta t$) and is power is independent of frequency, and equals $2/\langle x_k \rangle_k$ in the case of the Miyamoto normalisation.

1.2.3 Coherence

Given two ideal signals $u(t)$ and $v(t)$ observed in different energy channels or ranges of channels and assuming that they are continuous, it is possible to find a function $h(t)$ that relates the two signals via (Vaughan and Nowak, 1997):

$$v(t) = \int_{-\infty}^{+\infty} h(t-\zeta)\, u(\zeta)\, \mathrm{d}\zeta \quad . \tag{1.24}$$

By virtue of the convolution theorem, the Fourier transform of the convolution of two functions equals the product of the Fourier transforms, hence

$$\hat{V}(\nu) = \hat{H}(\nu)\hat{U}(\nu) \quad . \tag{1.25}$$

If the quantity $\hat{H}(\nu)$ is independent of the segment used to sample it, then the process signals $\hat{U}(t)$ and $\hat{V}(t)$ are said to be fully coherent at frequency ν (Vaughan and Nowak, 1997).

The coherence can be quantified in the Fourier-frequency domain thus:

$$\gamma^2(\nu) = \frac{|\langle C(\nu)\rangle|^2}{\langle|U(\nu)|^2\rangle\langle|V(\nu)|^2\rangle} \quad , \tag{1.26}$$

where the averages are to be taken over a number of segments m. The quantity $C(\nu) = U^*(\nu)V(\nu)$ is the cross-spectrum that is used below, and has an uncertainty $\delta\gamma^2(\nu) = \sqrt{2/m}\gamma^2(\nu)(1-\gamma^2(\nu))/|\gamma(\nu)|$ if no observational noise is present (Bendat and Piersol, 2010).

When $\gamma^2(\nu) = 1$, the two signals are said to be fully coherent at frequency ν, in which case $H(\nu)$ is a constant complex number independent of segment (i.e. the two signals are linearly correlated).

In real applications, signals are noisy due to Poisson additive noise. Hence, the coherence function in Eqn. 1.26 has to be corrected for the extra noise term.

The observed signal in Fourier domain in the bands h and h' will be:[8]

$$\hat{X}_{h,l} = \hat{S}_{h,l} + \hat{N}_{h,l} \tag{1.27}$$

and

$$\hat{X}_{h',l} = \hat{S}_{h',l} + \hat{N}_{h',l} \quad . \tag{1.28}$$

Using this notation, the term at the numerator in Eqn. 1.26 can be rewritten thus:

$$\begin{aligned} |\langle C_{hh',l}(\nu)\rangle_m|^2 &= |\langle (S_{h,l} + N_{h,l})^* (S_{h',l} + N_{h',l})\rangle_m|^2 \\ &= |\langle S^*_{h,l}S_{h',l}\rangle_m + \langle S^*_{h,l}N_{h',l}\rangle_m + \langle N^*_{h,l}S_{h',l}\rangle_m + \langle N^*_{h,l}N_{h',l}\rangle_m|^2 \quad , \end{aligned} \tag{1.29}$$

where the averages $\langle\cdot\rangle_m$ are taken over m contiguous segments. The subscript m will be omitted in the following.

In this treatment, the quantity $s_{hh',l} = \langle S^*_{h,l}S_{h',l}\rangle$ can be understood as the intrinsic signal component to the observed cross-spectrum $\langle C_{hh',l}(\nu)\rangle$. The remaining terms can be considered together as the noise contribution to the true underlying signal, $n_{hh',l} = \langle S^*_{h,l}N_{h',l} + N^*_{h,l}S_{h',l} + N^*_{h,l}N_{h',l}\rangle$. Given that the signal in each band is uncorrelated with the noise in either band, the averages $\langle\Re(n_{hh',l})\rangle = \langle\Im(n_{hh',l})\rangle = 0$ for sufficiently large m, and its variance will be given by

[8] Thanks to the linearity property of the Fourier transform, for a set of channels $h'' \in S$, $\mathscr{F}(\sum_{h''\in S} x_{h'',l}) = \sum_{h''\in S}\mathscr{F}(x_{h'',l})$ and therefore the concept of energy channels and bands can be used interchangeably in this context.

$$\langle |n_{hh',l}|^2 \rangle = \frac{1}{m} \left(\langle |S_{h,l}|^2 \rangle_m \langle |N_{h',l}|^2 \rangle_m + \langle |N_{h,l}|^2 \rangle_m \langle |S_{h',l}|^2 \rangle_m + \langle |N_{h,l}|^2 \rangle_m \langle |N_{h',l}|^2 \rangle_m \right) \quad (1.30)$$

equally distributed between its real and imaginary components.

In the general (and not restricted to cross-spectra) case where a weak signal of amplitude s is observed in the presence of observational noise with rms-amplitude d, it it known from the literature (Thomas, 1969; Goodman, 1985) that the observed squared amplitude a^2 follows a probability density function given by (Vaughan and Nowak, 1997; Groth, 1975):

$$p(a^2|s^2,d^2) = \frac{1}{d^2} \exp\left(-\frac{a^2+s^2}{d^2}\right) I_0\left(\frac{2as}{d^2}\right) \quad , \quad (1.31)$$

where $I_0(\cdot)$ is the zero-th order modified Bessel function of the first kind (Abramowitz and Stegun, 1965). This relation can be inverted to give the probability of s^2 given a^2 and d^2 (Chakrabarty, 1996):

$$p(s^2|a^2,d^2) = \frac{1}{ds\sqrt{\pi}} \exp\left(-\frac{a^2+2s^2}{2d^2}\right) \frac{I_0(2as/d^2)}{I_0(a^2/2d^2)} \quad . \quad (1.32)$$

In the case where $a \gg d$, Eqn. 1.32 can be approximated by a Gaussian, and follows that $\langle a^2 \rangle = \langle s^2 \rangle + \langle d^2 \rangle$.

In our treatment, the squared amplitude of the cross-spectrum, in the case of high measured coherence and high powers can be found to be (Vaughan and Nowak, 1997): $|\langle C_{hh',l} \rangle|^2 = |\langle S^*_{h,l} S_{h',l} \rangle|^2 + n^2_{hh',l}$. Taking Eqn. 1.32 into account, the noise-corrected coherence becomes (Vaughan and Nowak, 1997):

$$\gamma^2_{cor,hh',l} = \frac{|\langle C_{hh',l} \rangle|^2 - n^2_{hh',l}}{\langle |S_{h,l}|^2 - |N_{h,l}|^2 \rangle \langle |S_{h',l}|^2 - |N_{h',l}|^2 \rangle} \quad . \quad (1.33)$$

1.2.4 Lags

The phase-delay (or phase-lag) between the signal in two bands $U(\nu)$ and $V(\nu)$ can be obtained from the phase of the cross-spectrum $C(\nu)$ that is averaged over an ensemble of measurements (in our case, segments), and can be therefore obtained as:

$$\varphi(\nu) = \arctan \frac{\Im(\langle C(\nu) \rangle)}{\Re(\langle C(\nu) \rangle)} \quad . \quad (1.34)$$

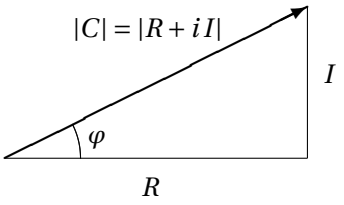

Figure 1.6: Amplitude and real and imaginary parts of the cross-spectrum $C = R + iI = \sqrt{R^2 + I^2}e^{i\varphi}$.

The time-lag then equals:

$$\tau(\nu) = \frac{\varphi(\nu)}{2\pi\nu} \quad . \tag{1.35}$$

The uncertainty in the estimation of the phase-lag φ from a cross-spectrum C (assuming $C = R + iI$ and dropping the dependence on frequency) can be obtained from $\tan\varphi = \frac{I}{R}$ (Fig. 1.6) by taking differential increments :

$$\sec^2(\varphi)\Delta\varphi \approx \frac{R\Delta I - I\Delta R}{R^2} \tag{1.36}$$

where

$$\sec^2(\varphi) = \frac{|C|^2}{R^2} \quad , \tag{1.37}$$

so that

$$\Delta\varphi \approx \frac{R\Delta I - I\Delta R}{|C|^2} \quad . \tag{1.38}$$

The variance of the phase-estimates can be obtained from Eqn. 1.38 (see Bendat and Piersol 2010, Eqn. 9.51), from which follows that the error on the phase-lag equals the standard deviation of the best estimates of the phase-lag after carrying out m measurements (Eqn. 1.39) like so:

$$\delta\varphi = \sigma_\varphi = \frac{(1 - \gamma_{\text{raw}}^2)^{1/2}}{\gamma_{\text{raw}}\sqrt{2m}} \tag{1.39}$$

where

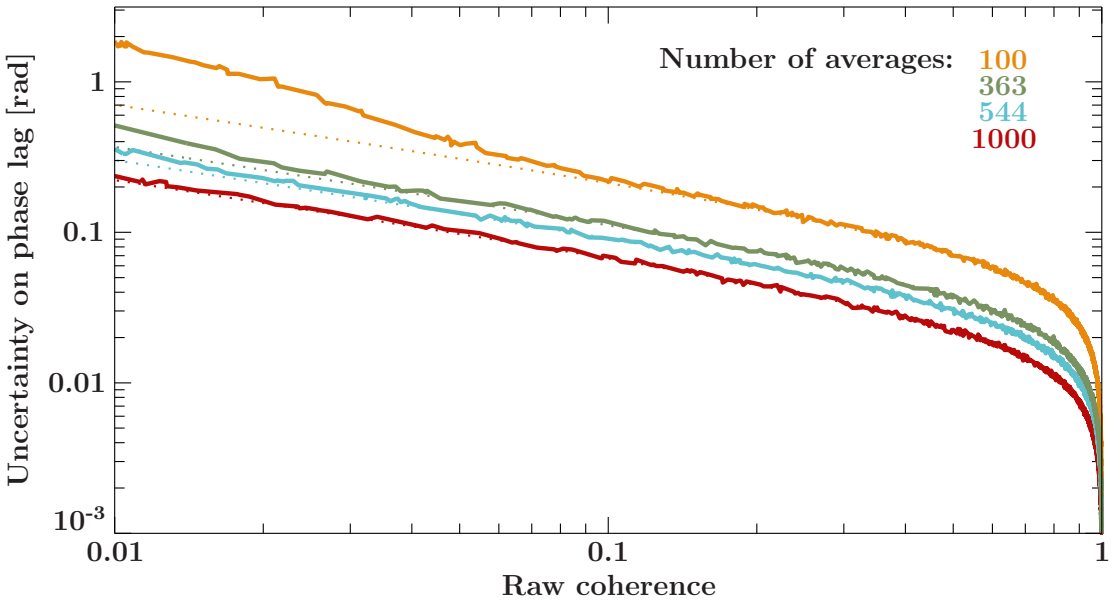

Figure 1.7: Dependence of the uncertainty of the phase lag on raw (i.e. non-noise-corrected) coherence, using both the method described in Bendat and Piersol (2010) (see Eqn. 1.39, dotted lines) and simulated data using different numbers of segments to estimate the cross-spectra (see text for details). The colours orange, green, cyan and red correspond to 100, 363, 544 and 1000 segments used to average the cross-spectra.

$$\gamma^2_{\rm raw} = \frac{|\langle C(\nu)\rangle|^2}{\langle |U(\nu)|^2\rangle \langle |V(\nu)|^2\rangle} \quad . \tag{1.40}$$

The quantity $\gamma^2_{\rm raw}$ in Eqn. 1.40 differs from $\gamma^2_{\rm cor}$ in Eqn. 1.33 in that no noise subtraction is performed, and follows directly from the mathematical manipulation in Bendat and Piersol (2010).

It follows naturally that the error on the time-lag, for a given frequency ν, equals

$$\delta\tau(\nu) = \frac{\delta\varphi(\nu)}{2\pi\nu} \quad . \tag{1.41}$$

The validity of Eqn. 1.39 to estimate the uncertainty on the observed phase-lags is tested here with simulations. 1000 combinations of signal and noise amplitudes are extracted (see Eqn. 1.27 and Eqn. 1.28) for both energy bands, given different intrinsic signal phase-lags between the two bands. Each combination of amplitudes is used to average the cross-spectrum m times, where m =100, 363, 544 and 1000. The values m = 363 and m = 544 correspond to the number of segments used to extract the phase-lags for the 2006 and 2009 observations of SWIFT J1753.5–0127 in Chapter 3, respectively. The phase-lag is then extracted from the averaged cross-spectrum. Finally, the standard deviation of the phase-lag best-estimates as a function of resulting coherence is plotted in Fig. 1.7 and compared to the theoretical estimates in Bendat and Piersol (2010). As it can be seen, Eqn. 1.39 the theoretical estimates match the simulated data perfectly for $m > 100$ for values of the raw coherence above 0.1.

For $m = 100$, the simulated errors diverge clearly from the expected curve. This behaviour is caused by the fact that, at low coherence, the phase-lags are more likely to wrap around and produce additional scatter. In this Thesis, the errors on the phase-lags and time-lags have been derived from Eqns. 1.39 and 1.41. The amplitude resulting from the above equations matches well the simulated amplitude in the case of the observations of SWIFT J1753.5–0127 and GX 339-4 in the ranges studied, where the raw coherence > 0.015 depending on the combination of bands (the analysis of GX 339–4 has been performed with a much larger number of segments). It is worth noting that the error-bars in Fig. 1.7 are for a single frequency, while the phase-lags shown in the following chapters have been rebinned in frequency, thus the values in Fig. 1.7 should be decreased by a factor $1/\sqrt{k_l}$, where k_l is the number of frequencies that contribute to a given frequency bin.

1.2.5 Variability of black hole X-ray binaries

Variability in black hole X-ray binaries is most easily studied using power spectral densities, extracted using the method outlined above.

In the hard state, power spectral densities were initially modelled as a doubly-broken power-law, showing a flat-top (in νP_ν), or band-limited flicker noise[9], with $P(\nu) \propto \nu^{-1}$ (Done et al., 2007). This band-limited noise is limited by two characteristic frequencies, ν_b and ν_l, that correspond to two breaks in the power-law distribution, $P(\nu < \nu_b) \propto \nu^0$ and $P(\nu > \nu_l) \propto \nu^{-2}$.

Variability in power density spectra is observed over a broad range of frequencies. Additive shot models where shots of different widths appear independently of one another are able to reproduce virtually any noise continuum, yet it has been difficult to ascribe them to a physical process that takes place in the source. An explanation for the large dynamic range of frequencies where variability power is observed was proposed by Lyubarskii (1997). The model proposed in this work assumes that the variations in luminosity are produced by variations in mass-accretion rate at inner radii, where most of the energy is released. These variations would result from the contribution of independent fluctuations in mass-accretion rate at larger radii on the local viscous time-scale[10], via a multiplicative process. The resulting power spectrum would have a ν^{-1} power-law shape. From this model, it follows that the frequencies between which variability is observed could be inversely proportional to the viscous time-scale at the inner and outer radii of the accretion flow. Uttley and McHardy (2001)

[9] In the following, we will refer to 'noise' as the effect of a stochastic process that, for infinite realisations, has a particular average distribution of powers as a function of frequency or higher-order properties, e.g. in terms of expected cross-correlations between different bands, as opposed to simple shot noise models, e.g. Doi (1978); Terrell (1972).

[10] The viscous time-scale at a radius R in an accretion flow with a viscosity parameter α and disc half-thickness H is $t_{\rm visc} = \frac{2}{3}\frac{1}{(H/R)^2\alpha}\frac{1}{\Omega(R)}$ (Shakura and Sunyaev, 1973), where $\Omega(R) = \left(\frac{GM}{R^3}\right)^{1/2}$ is the Keplerian angular velocity for a black hole mass M. The viscous time-scale corresponds to the time-scale on which the material (of surface density Σ) drifts inwards with a radial drift velocity $v_r(R)$, giving rise to a local mass-accretion rate $\dot{M} = 2\pi R v_r(R)\Sigma$.

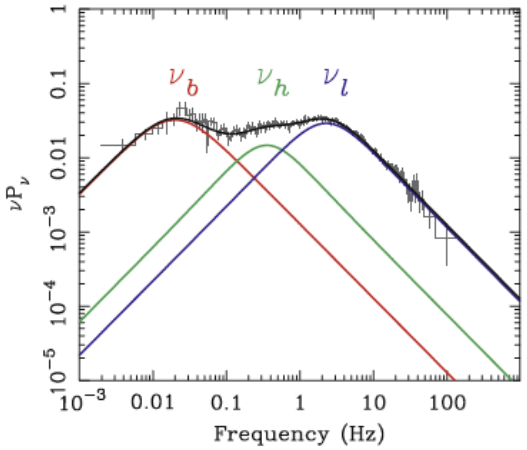

Figure 1.8: Hard-state PSD of GX 339–4 in νP_ν units, showing the Lorentzian components described in the text. From Done et al. (2007).

showed that the absolute rms amplitude of small segments of a BHXRB lightcurve scales linearly with the segment mean flux, and predicted that Lyubarskii's model would reproduce the data. Uttley et al. (2005) defined the non-linear underpinnings of the processes that produce the rms-flux relation and confirmed that the multiplicative nature of Lyubarskii's model would reproduce the log-normality seen in the observed lightcurves. The spectral-timing implications of the multiplicative nature of Lyubarskii's model in Kotov et al. (2001) and Arévalo and Uttley (2006) show that mass-accretion rate fluctuations can, in principle, reproduce the observed data beyond power spectra alone.

With better data, the PSDs show deviations away from simple broken power-laws and can be better fitted by a superposition of several (~4–5) Lorentzian components, where 2–3 typically form the band-limited noise[11] (Done et al. 2007 and references therein). Between ν_b and ν_l there is another Lorentzian component peaking at ν_h that is linked to a quasi-periodic oscillation (QPO) at $\nu_{\rm LF}$ and may replace it when the QPO is not clear.

As the transition between the hard and soft states takes place, the Lorentzian components are seen to move together (Figs. 1.4 and 1.9), leading to important correlations, $\nu_{\rm LF} \sim 5\nu_b$ (Wijnands and van der Klis, 1999) and $\nu_l \sim 10\nu_{\rm LF}$ (Psaltis et al., 1999; Belloni et al., 2002). The amplitude and width of these components show additional characteristics as they shift in frequency (see Done et al. 2007 and references therein). The low-frequency break ν_b is consistent with being produced at the inner edge of the geometrically-thin accretion disc, the disc damping most of the variability produced on time-scales shorter than the local viscous time-scale, while the variability at higher frequencies could be produced in an inner flow (see e.g. Churazov et al. 2001; Wilkinson and Uttley 2009). If ν_b moves towards higher frequencies

[11] A Lorentzian profile has the form $L(\nu|\nu_0) = \frac{1}{\pi}\frac{2R_0^2 Q_0 \nu_0}{\nu_0^2+4Q_0^2(\nu-\nu_0)^2}$, where ν_0 is the centroid frequency, $Q_0 \sim \nu_0/\Delta\nu_{0,\rm FWHM}$ is the quality factor, and R_0 is the normalisation (Belloni et al., 2002; Pottschmidt et al., 2003). See Pottschmidt et al. (2003) for more details.

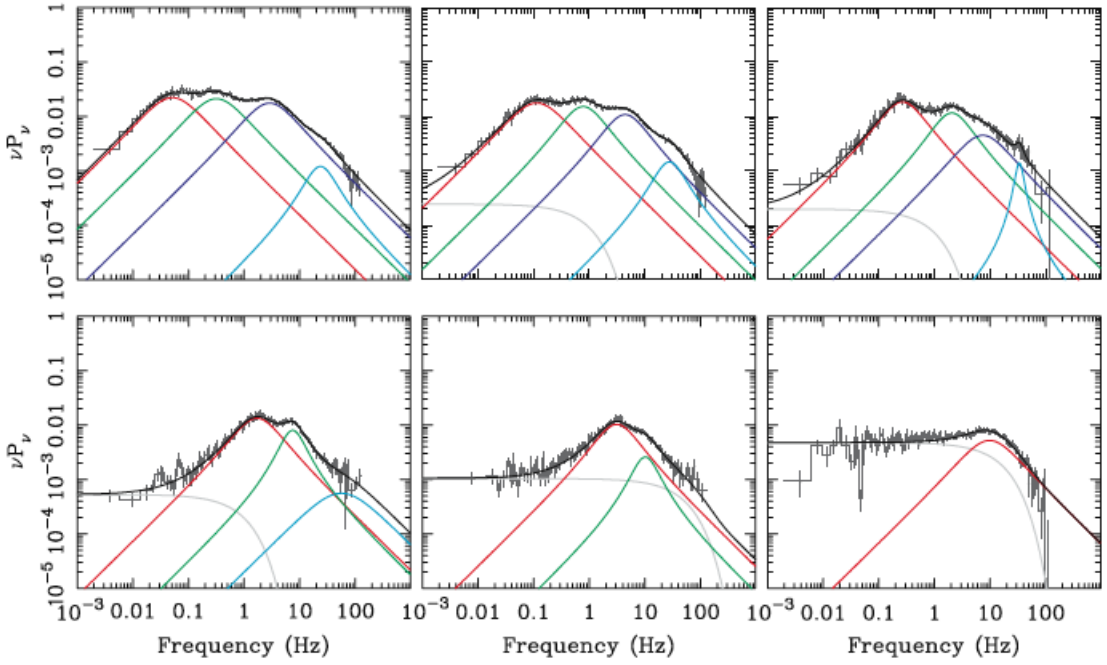

Figure 1.9: Evolution of power spectral densities as Cygnus X–1 transitions from the hard (top left) to the soft (bottom right) state, from Axelsson et al. (2005). It is worth noting that the soft state variability in persistent sources such as Cygnus X–1 (e.g. bottom right panel) is stronger than in transient sources.

as the source transitions from the soft to the hard state, then this suggests that the disc may be less truncated for increasing ν_b, supporting the disc truncation scenario.

There are independent arguments that support the picture of a centrally-located source of power-law photons. These include the amount of reflected emission argument. If the source of power-law photons were more extended, too large a solid angle would produce too large a fraction of reflected emission compared to observed spectra. From spectral-timing studies with *RXTE*, it has been shown that the strength of the reflection spectrum depends on the Fourier-frequencies probed in frequency-resolved rms spectra, thus on the radial extent of the accretion disc (Revnivtsev et al., 1999). In this case, the illuminating power-law component must be central, too. Some authors argue that these Comptonised photons must originate from an optically-thin inner region (Thorne and Price, 1975; Esin et al., 1997), perhaps Advection-Dominated Accretion Flow (Narayan and Yi, 1994) that results from evaporation of the accretion disc at low luminosities. The low density of this gas results in the electrons not Comptonising much of the seed disc photons. The ADAF automatically results in the standard accretion disc being truncated, the truncation radius depending on the mass-accretion rate (higher accretion rate, as expected, results in less truncation). Notwithstanding the actual model for the inner flow, there exist additional correlations in the timing domain between the source hardness, reflection fraction and break frequency in power density spectra (Gilfanov et al., 1999), supporting this picture. In addition, it is also possible that, closer to the black hole, the optically-thin accretion flow condenses in the form of an 'inner' disc (Liu et al., 2007; Chiang et al., 2010). Other sources of power-law emission that have been studied

in the literature include a hot corona (Malzac and Belmont, 2009) or even the base of a jet (Markoff et al., 2005).

The understanding of variability power as a function of source luminosity and spectral shape has resulted in tighter constraints on the possible geometries and physics of the accreting flows, with a formidable contribution from the observations obtained with the timing satellite *RXTE*. Most of the variability studies done were based on power density spectra to infer physical parameters from the distribution of variability power as a function of Fourier frequency. Yet variability 'signals' in different bands can be used to map the geometry of these systems via the cross-spectrum and its phase.

Using *RXTE*, which did not cover energies below ~ 2 keV, it was thought that the variability observed originated in the optically-thin inner flow, while the disc remained stable (Churazov et al., 2001). This evidence was compatible with the anti-correlation of rms variability amplitude during the hard-soft transition, hence the disc was seen as stable (Belloni et al., 2005).

However, recent work with *XMM-Newton* using covariance spectra and energy-dependent time-lags (Wilkinson and Uttley, 2009; Uttley et al., 2011) shows that the disc is intrinsically variable on viscous time-scales and leads the harder, power-law emission with a delay that is consistent with the scenario described above where mass-accretion rate fluctuations produce the observed variability. Fluctuations produced in the accretion disc may be propagated towards smaller radii via a multiplicative process until they reach the source of power-law photons. In fact, at low frequencies, the energy-dependent time-lags show that the emission where the power-law dominates lags the soft emission. At higher frequencies ($\nu \gtrsim 1$ Hz), an up-turn at low energies appears, consistent with the power-law now leading the disc emission via reprocessing of power-law photons in the disc. This model underlies most of the work in this Thesis, and will be discussed further in the following Chapters.

In any case, the above results show that joint spectral-timing studies provide a much broader picture than spectral or timing studies separately. This Thesis will address some of the open questions in the study of black hole X-ray binaries as follows.

1.3 Purpose of this work

The work presented in this Thesis is aimed towards understanding some of the basic properties of the geometry of accretion discs in black hole X-ray binaries and its emission, via a powerful combination of spectral and timing technniques that have resulted in interesting discoveries in the past regarding disc variability and causality links between spectral components (Wilkinson and Uttley, 2009; Uttley et al., 2011). In the present Thesis, it is our scope

to study stellar-mass black hole systems with a special focus on the link between emission variability properties in the Fourier domain and their multi-channel cross-correlation properties. This is because each emission component may vary with different amplitudes, and their variable emission be delayed with respect to the emission from other components. A careful examination and modelling of such relationships between different variable spectral components is able to put very strong constraints on the geometry of black hole systems as well as the location of the X-ray emitting regions and relative solid angles. These constraints are much stronger than the best-fit parameters that can result from 'standard' spectral fitting alone, where energy resolution and calibration, as well as model degeneracy make it almost impossible to confirm many of the scenarios predicted so far.

In Chapter 2, we explore the possibility that the hard-to-intermediate energy (4.0 – 10.0 keV vs 2.0 – 3.5 keV) time-lags in black hole X-ray binaries be produced by reflection of coronal photons off a flared accretion disc. We model a parameter-dependent transfer function for reflection that results in two different outputs, a resulting reflection spectrum, and a time-lag vs. Fourier-frequency dependence. Due to the limited energy resolution of the instrument, energy spectra are poorly sensitive to large-scale reflection to constrain the outer radius of the disc, yet they can constrain the inner radius of the accretion disc (where Doppler shifts are most important). Time-lags, on the other hand, increase for increasing size-scale of the reflector, but are very small for small-scale reflection. We fit the model simultaneously to energy spectra and time-lags of GX 339–4 in the 2004 and 2009 in order to set limits to the solid angle subtended by the accretion disc and understand the cause of the hard-to-intermediate time-lags.

Chapter 3 is devoted to the understanding of the variability properties of the accretion disc in high-luminosity hard states, as compared to fainter hard states. Previous observations of hard-state black hole X-ray binaries with soft X-ray coverage below ~ 2 keV showed that the disc variability leads the power-law variability on time-scales comparable to the disc viscous time-scale, and likely drives the power-law variability (Uttley et al., 2011). We compare the 2009 hard-state observation of SWIFT J1753.5–0127 to the 2006 hard state, and compare interesting different properties in the disc amplitude of variability, coherence of its fluctuations, as well as time-lags.

A novel two-dimensional fitting technique is shown in Chapter 4 for constraining the link between spectral and variability properties of emission using a self-consistent approach. Best-fit model values obtained with this method should reproduce observational variability amplitudes as well as phase-/time-lags dependencies across the energy and frequency domains. This technique opens up an exploratory space in a deeper understanding of the connection between different spectral components and the geometry of black hole systems.

CHAPTER 2

Joint spectral-timing modelling of the hard lags in GX 339–4: constraints on reflection models

The X-ray variations of hard state black hole X-ray binaries above 2 keV show 'hard lags', in that the variations at harder energies follow variations at softer energies, with a time-lag τ depending on Fourier-frequency ν approximately as $\tau \propto \nu^{-0.7}$. Several models have so far been proposed to explain this time delay, including fluctuations propagating through an accretion flow, spectral variations during coronal flares, Comptonisation in the extended hot corona or a jet, or time-delays due to large-scale reflection from the accretion disc. In principle these models can be used to predict the shape of the energy spectrum as well as the frequency-dependence of the time-lags, through the construction of energy-dependent response functions which map the emission as a function of time-delay in the system. Here we use this approach to test a simple reflection model for the frequency-dependent lags seen in the hard state of GX 339–4, by simultaneously fitting the model to the frequency-dependent lags and energy spectrum measured by *XMM-Newton* in 2004 and 2009. Our model cannot simultaneously fit both the lag and spectral data, since the relatively large lags require an extremely flared disc which subtends a large solid angle to the continuum at large radii, in disagreement with the observed Fe Kα emission. Therefore, we consider it more likely that the lags > 2 keV are caused by propagation effects in the accretion flow, possibly related to the accretion disc fluctuations which have been observed previously.

2.1 Introduction

Black hole X-ray binaries (BHXRBs) have extensively been studied both in terms of their spectra (where a soft, multi-colour disc black-body and a harder power-law component play the main role) and in terms of their variability, which can be straightforwardly quantified using the common approach of computing the Power Spectral Density (PSD) of their signal. In combination with the time-averaged X-ray spectrum, this approach has enabled the variability to be studied as a function of spectral state, so that the hard power-law-dominated states can be associated with large variability amplitudes (tens of per cent fractional rms) and band-limited PSD shapes, while the soft disc-dominated states show much weaker variability and broadband power-law like PSD shapes.

Although they provide information on the global connection between the emitting components and the variability process, the PSD and mean energy spectra alone carry no information about the complex pattern of interlinkage between diverse regions in the accreting system that lead to transfer of variability from one physical component to another over a range of time-scales. In order to understand the causal connection between emission processes that show different relative strengths in two different energy bands, the time-lags between these two bands can be extracted as a function of Fourier frequency. To date, frequency-dependent lags have been studied mostly in the hard state of BHXRBs. There the lags (τ) are 'hard', in the sense that variations in harder bands lag behind variations in softer bands and depend on frequency as $\tau \propto \nu^{-0.7}$, albeit with some sharper 'steps' in the lag-frequency relation (Miyamoto et al., 1988; Cui et al., 1997; Nowak et al., 1999). The general form of these lags has been explained by a variety of models, including Comptonisation in the extended corona (Kazanas et al., 1997) or a jet (Reig et al., 2003; Kylafis et al., 2008), spectral variability during coronal flares (Poutanen and Fabian, 1999), accretion fluctuations propagating through a power-law emitting region (e.g. a corona) with an energy-dependent radial emissivity profile (Kotov et al., 2001; Arévalo and Uttley, 2006), or light-travel times to an extended reflecting region (Kotov et al., 2001; Poutanen, 2002). Models invoking Comptonisation in the extended regions predict broader auto-correlation function for photons at higher energies, which suffer more scattering, which is opposite to what is observed (Maccarone et al., 2000; Poutanen, 2001), although more complex models of Comptonisation by the jet may be considered (see e.g. Giannios (2005)).

Recent measurements of the lag (Uttley et al., 2011) between variations of the accretion disc blackbody emission and the power-law component strongly indicate that the variations are driven by fluctuations propagating through the disc (Lyubarskii, 1997) to the power-law emitting region, however this model cannot simply explain the hard lags between bands where the power-law dominates.

Models to generate frequency-dependent lags work by determining the response to an input signal of an emitting region which produces a hard spectral component, e.g. a compact coronal region sandwiching the disc, the upscattering region in the jet or an extended reflector. The input signal may be changes in the geometry, a fluctuation in accretion rate, seed photon illumination or illuminating primary continuum. The hard emitting region can only respond after a delay which, broadly speaking, is set by the time taken for the signal to propagate to and across the region. The delay time is determined by the signal speed (viscous time-scale or light-travel time) and the size scale of the hard emitting region. If the delay is large compared to the variability time-scale, the variations of the hard emitting region are smeared out and so the amplitude of variations of the lagging component is reduced. Thus the observed drop in lags with Fourier frequency can be reproduced.

In principle, models for the emitting regions which can reproduce the lags should also be able to reproduce the energy spectrum. Therefore, the combination of lag information with information on the X-ray spectral shape should provide much greater constraints on models for the emitting regions than either the commonly used spectral-fitting methods or the rarely-attempted fits to timing data. As a good proof-of-principle, the simplest models to attempt this joint lag and spectral-fitting are those where the lags are produced by reflection, i.e. so-called 'reverberation' lags. Evidence for small (few to tens of $R_{\rm g}$ light-crossing time) reverberation lags generated by reflection close to the black hole has been seen in Active Galactic Nuclei (Fabian and Ross, 2010; Zoghbi et al., 2010; de Marco et al., 2011; Emmanoulopoulos et al., 2011) and BHXRBs (Uttley et al., 2011), however the much larger lags we consider here require a larger scale reflector, perhaps from a warped or flaring outer-disc, which was originally considered by Poutanen (2002). Reflection of hard photons from a central corona off an accretion disc should produce both a reflection signature in the spectrum and a frequency-dependent time-lag, the shapes of which should depend on the exact geometry of the reflector as well as on the location of the corona itself (for a central, point-like corona, this would be the height above the disc). It has to be stressed that only non-flat disc geometries are expected to contribute a significant lag at low frequencies (Poutanen, 2002).

The CCD technology of the EPIC-pn camera onboard *XMM-Newton* provides a time resolution of 5.965 ms (frame readout time) in Timing mode with a 99.7% livetime and an energy resolution ~ 125 eV at 6 keV that combined, offer great potential for developing models that take into account energy and variability information together down to millisecond timescales. Our aim in the present Chapter is to fit the spectra and time lags from an EPIC-pn Timing mode observation of a BHXRB simultaneously to discern whether reflection is the main driver of the observed hard-to-medium lags. In order to achieve this, we have developed REFLAGS, a reflection model that assumes a flared accretion disc acting as reflector (assuming the constant density ionised disc reflection spectrum of Ballantyne et al. 2001), and is able to output for a given geometry either the resulting spectrum, or the expected lags as a function of frequency. We use this model to fit simultaneously the spectra and time-lags of the low-mass BHXRB GX 339–4 in the *XMM-Newton* observations of the hard state obtained

in 2004 and 2009, and so determine whether the reflection model can explain the hard lags observed in these observations while also remaining consistent with the X-ray spectrum. In principle, this combined approach could yield much greater constraints on the outer disc geometry than can be obtained with spectral-fitting alone, since low velocities of the outer regions of the disc cannot be resolved with CCD detectors, although these regions should contribute significantly to the lags if they subtend a large solid angle as seen from the source.

We describe our model in Section 2.2 and the data reduction and extraction of spectra and time-lags in Section 2.3. In Section 2.4 we discuss the results from fitting the energy spectra and frequency-dependent time-lags of GX 339–4 together with our model, and also compare the expected optical/UV reprocessing signature of the inferred geometries with data from the *Swift* satellite. We discuss our results and the wider implications of our combined spectral-timing model-fitting approach in Section 2.5.

2.2 A flared accretion disc model

2.2.1 Model parameters

We consider a simple reflection model where the hard lags are produced by the light travel times from a variable power-law continuum source to the surface of a flared accretion disc which absorbs and reprocesses the incident radiation or scatters it producing a hard reflection spectrum (George and Fabian, 1991). Following the geometry described by Poutanen (2002), we have written REFLAGS, an XSPEC model that can describe both the lags as a function of frequency and the mean spectrum, allowing simultaneous fitting by tying together the parameters from the spectral and lag fits. We assume that the accretion disc is axially symmetric but has a height of the disc surface above the mid-plane which depends on radius r as a power-law

$$z(r) = H_{\rm out}\left(\frac{r}{R_{\rm out}}\right)^{\gamma} \quad , \tag{2.1}$$

where $H_{\rm out}$ is the height of the disc surface at the outermost radius $R_{\rm out}$ and γ is the 'flaring index'.

We place a point-like source of Comptonised photons at a certain height $H_{\rm src}$ above the disc, located in the axis of symmetry of the system, and assume that the disc is truncated at some inner radius $R_{\rm in}$. Because the lags are proportional to the distances, while the spectral distortions due to rapid rotations are a function of radii in units of gravitational radius $R_{\rm g} = GM/c^2$, we have to specify the black hole mass, which we assume to be $M_{\rm BH} = 10{\rm M}_{\odot}$.

The amount of reflection of Comptonised photons that is expected to come from each region of the disc strongly depends on the value γ, and the only case where the contribution of outer radii to spectra and lags can be significant is for concave ($\gamma > 1$) geometries (Poutanen, 2002).

The model also depends on the two parameters that describe the spectral shape: the incident power-law photon index Γ and the ionisation parameter ξ. Due to computational constraints in the model evaluation, the latter is approximated to be constant throughout the disc. For computational purposes, the disc is also divided into 100 equally spaced azimuthal angles and 300 logarithmically spaced radii, to produce of a total of 30 000 cells.

The total energy-dependent power-law plus reflection luminosity per unit solid angle at a time t that is emitted by a system whose geometry is described by the parameters above (represented by the set α) and seen at an inclination angle i with respect to the axis of symmetry, equals:

$$L_{\rm tot}(E,t|\alpha,i) = L_{\rm PL}(E,t) + a(E)\frac{\Omega_{\rm eff}}{2\pi}\sum_l L_{\rm PL}(E,t-\tau_l)\kappa_{\alpha,l}(E,i) \quad , \tag{2.2}$$

where the sum is performed over the index l that represents a single disc cell, $L_{\rm PL}(E,t)$ is the luminosity per unit solid angle produced by power-law emission and $a(E)$ is the albedo function.

The factor $\kappa_{\alpha,l}(E,i)$ contains the several projection terms and solid angle corrections required for the l-th cell for the geometry described by α (see Poutanen 2002), as well as relativistic Doppler and gravitational redshift corrections (following the same simplified approach as the DISKLINE model, Fabian et al. 1989), and is normalised so that $\sum_l \kappa_{\alpha,l}(E,i) = 1$. For a given geometry α and inclination angle i, $\Omega_{\rm eff}$ then equals the solid angle subtended by the disc, corrected for the source inclination as seen by the observer. Finally, τ_l represents the time delay between observed direct power-law emission and reflected emission coming from the l-th cell.

2.2.2 System response and time-lags

For a direct power-law emission pulse (a delta-function) equation (2.2) can be rewritten as:

$$\begin{aligned} L_{\rm tot}(E,t|\alpha,i) &= L_{\rm PL}(E) \\ &\times \left(\delta(t) + a(E)\frac{\Omega_{\rm eff}}{2\pi}\sum_l \kappa_{\alpha,l}(E,i)\delta(t-\tau_l)\right) \quad , \end{aligned} \tag{2.3}$$

and can be factorised as

$$L_{\rm tot}(E,t|\alpha,i) = L_{\rm PL}(E)\,T(E,t|\alpha,i) \quad , \tag{2.4}$$

where

$$T(E,t|\alpha,i) = \delta(t) + a(E)\frac{\Omega_{\rm eff}}{2\pi}\sum_l \kappa_{\alpha,l}(E,i)\delta(t-\tau_l) \quad . \tag{2.5}$$

The system response function $T(E,t|\alpha,i)$ contains the energy-dependent time redistribution of an input power-law luminosity as a function of time in a geometry α, as observed at an inclination i. Provided that the energy spectrum is averaged over a time significantly longer than the largest τ_l, $\tau_{\rm max}$ (which in this case is of order 10 s), the spectrum is given simply given by: const. $\times \int_0^{\tau_{\rm max}} T(E,t'|\alpha,i)\mathrm{d}t'$. By taking the Fourier transform of equation (2.4):

$$\begin{aligned}\tilde{L}_{\rm tot}(E,\nu,i) &= \mathscr{F}\left\{L_{\rm PL}(E)\,T(E,t|\alpha,i)\right\} \\ &= L_{\rm PL}(E)\,\tilde{T}(E,\nu|\alpha,i) \quad .\end{aligned} \tag{2.6}$$

In general, the time-lags between two light curves $s(t)$ and $h(t)$ in the soft and hard bands can be calculated by computing their Fourier transforms $\tilde{S}(\nu) = \mathscr{F}[s(t)]$ and $\tilde{H}(\nu) = \mathscr{F}[h(t)]$ and forming the complex-valued quantity $C(\nu) = \tilde{S}^*(\nu)\tilde{H}(\nu)$, called the cross-spectrum (where the asterisk denotes complex conjugation). Its argument is the phase lag or phase difference between $\tilde{S}(\nu)$ and $\tilde{H}(\nu)$ (Nowak et al., 1999):

$$\phi(\nu) = \arg\left[C(\nu)\right] = \arg\left[\langle\tilde{S}^*(\nu)\tilde{H}(\nu)\rangle\right] \tag{2.7}$$

and therefore

$$\tau(\nu) = \frac{\phi(\nu)}{2\pi\nu} \tag{2.8}$$

equals the frequency-dependent time-lag.

The transfer function $\tilde{T}(E,\nu|\alpha,i)$ is equivalent to the response function $T(E,t|\alpha,i)$ in the Fourier-frequency domain. It is then possible to take the expression above to compute the cross-spectrum that will give the time-lags caused by reflection between two broad energy bands s and h and obtain the lags as above (the notation has been simplified):

$$C(\nu) = L_{\rm PL}(E_{\rm s})L_{\rm PL}(E_{\rm h})\tilde{T}^*(E_{\rm s},\nu)\tilde{T}(E_{\rm h},\nu) \quad , \tag{2.9}$$

$$\tau(\nu) = \frac{\arg\left[\tilde{T}^*(E_{\rm s},\nu)\tilde{T}(E_{\rm h},\nu)\right]}{2\pi\nu} \quad . \tag{2.10}$$

2.3 Data reduction

2.3.1 Extraction of spectra, event files and instrumental response files

GX 339–4 was observed using the *XMM-Newton* X-ray satellite on 2004 March 16 and 2009 March 26 (ObsIds 0204730201 and 0605610201, respectively). In the present work, we will concentrate on the data taken using the EPIC-pn camera in Timing Mode, which allows for a fast readout speed, as opposed to imaging modes. This is done by collapsing all the positional information into one dimension and shifting the electrons towards the readout nodes, one macropixel row at a time, thus allowing a fast readout with a frame time of 5.965 ms.[1]

The raw Observation Data Files (ODFs) were obtained from the *XMM-Newton* Science Archive (XSA) and reduced using the *XMM-Newton* Science Analysis System version 10.0.0 tool `epproc` using the most recent Current Calibration Files (CCFs).

With the aid of EVSELECT, the source event information was then filtered by column (`RAWX in [31:45]`), and pattern (`PATTERN <= 4`, corresponding to only single and double events), and only time intervals with a low, quiet background (where `PI in [10000:12000]`) were selected for the subsequent data analysis. The total exposure times are ~ 164 ks and ~ 39 ks for the 2004 and 2009 observations, respectively. The background spectrum was obtained from ObsId 0085680501 between columns 10 – 18 when the source was fainter so that the background is not contaminated by the source (see Done and Diaz Trigo (2010) for further details). `rmfgen` and `arfgen` were used to obtain the instrumental response files. To account for the effects of systematic uncertainties in the instrumentation, an additional 1% was added to the error bars.

In the present work, uncertainties in the estimation of the parameters are quoted at the 90% confidence level for one parameter of interest.

2.3.2 Extraction of the time-lags

When measuring the lags for real, noisy data one needs to first take the average of the cross-spectrum over many independent light curve segments (and also adjacent frequency bins).

[1] The frame time is the time required to read all the events detected in the 200 RAWY rows of the CCD. In this observing mode, the spatial information is preserved in only one dimension (RAWX columns), while the RAWY rows act as a de facto time axis. Every 29.56 μs the charge deposited by the events in the CCD is transferred from one RAWY row to the next towards the reading nodes, which are then read and time-tagged every 5.965 ms. The use of a 29.56 μs time resolution for the analysis could lead to deadtime effects and the detection of a spurious feature at 166.67 Hz (1/0.005965 Hz) caused by a bright row in RAWY = 19. See Kuster et al. (2002) and the XMM-Newton Users Handbook.

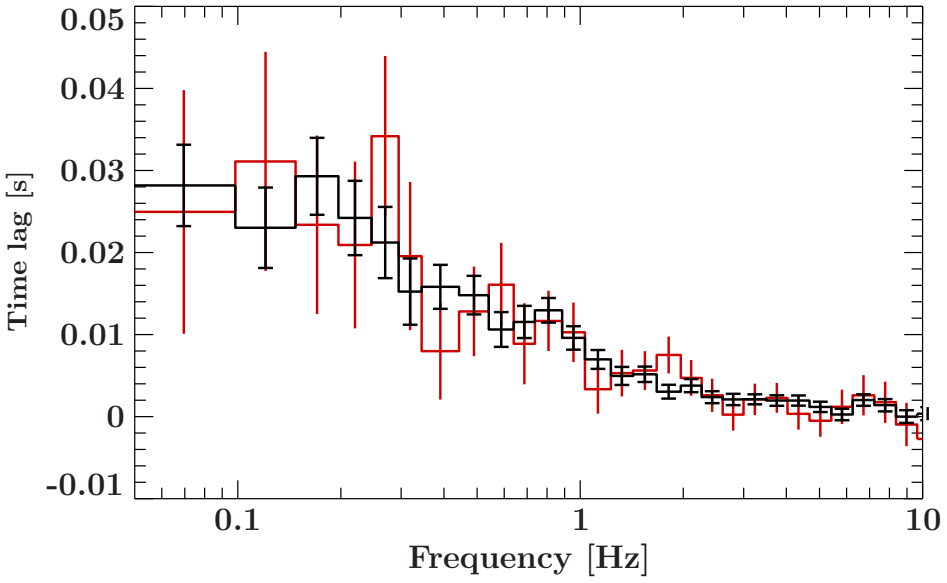

Figure 2.1: Time-lags between 4.0–10.0 and 2.0–3.5 keV energy bands as a function of frequency for the 2004 (black) and 2009 (red) observations.

This is because observational noise adds a component to the cross-spectrum which has a phase randomly drawn from a uniform distribution (see e.g. Nowak et al. 1999). By averaging over many independent measures of the cross-spectrum, the contribution of these noise components can be largely cancelled out (the residual error provides the uncertainty in the lag). Therefore, the time-lags $\tau(\nu)$ were extracted from the argument of the cross spectrum averaged over many segments of the light curve as the exposure time and the dropouts due to telemetry saturation permit. Throughout this work we will take as the soft band the interval 2.0–3.5 keV and as a hard band the interval 4.0–10.0 keV.

We choose to have 8192 bins/segment for a total duration of ~ 24.43 seconds per segment, giving 6690 segments for 2004 and 1582 segments for 2009. This will also constrain the frequency ranges over which lags can be obtained. These frequencies have also been rebinned geometrically with a binning factor 1.15 in order to improve the signal-to-noise per frequency bin.

The uncertainties on the lag measurements follow from the technique in Nowak et al. (1999) (see also Bendat and Piersol 2010). These scale as $(NM)^{-1/2}$ where N is the number of segments used to average the cross-spectrum and M is the number of frequencies averaged per bin. The lags are plotted on Fig. 2.1. Given the fact that, within the errors, the time-lag dependence on Fourier frequency is consistent between the two observations, we will sometimes assume that the time-lags for the two observations are equivalent and do not vary between the two observations as their shapes cannot be distinguished within the errors. This 'substitution' should in principle give us tighter constraints for our study in Section 2.4.

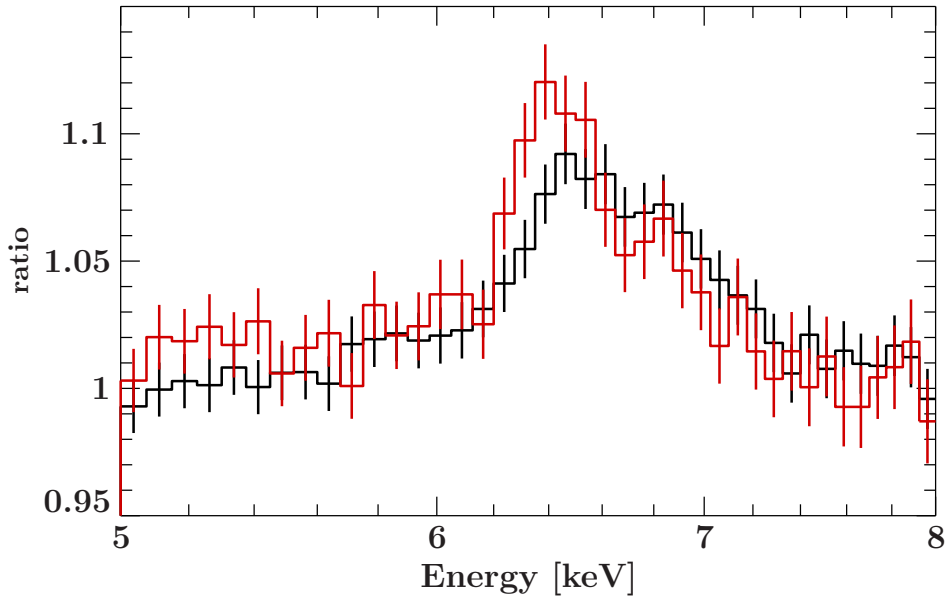

Figure 2.2: Ratio to a powerlaw after fitting the 2.0–5.0 keV and 7.0–10.0 keV continuum (black: 2004, red: 2009). The presence of the iron line is clear.

2.4 Results

We test the model described in Section 2.2 using the 2004 and 2009 observations of GX 339–4, in order to understand the validity of the model and infer the geometrical parameters of the reflecting disc that is required to explain both the spectrum and the lags. To do this, the same model is concurrently fitted to the spectrum as well as the lag versus frequency data. Using this approach, it is possible to discern the importance of reflection to explain the observed lags, and whether or not an extra variability component to produce the lags is required.

In Fig. 2.2 we plot the ratio of the 2004 and 2009 spectra to a power-law fitted to the 2.0–10.0 keV interval excluding the region 5.0–7.0 keV. The iron line shape appears clearly different, with the 2004 dataset showing a more broadened profile that is skewed towards higher energies, while the 2009 line appears narrower and peaked at about 6.4 keV, the value expected for neutral or weakly-ionised emission.

We perform fits using the 2004 and 2009 spectra in the 3.0–10.0 keV band as well as their respective lags (we name these model fits A and B, respectively). These energy intervals are chosen in order to avoid contamination from the band dominated by the disc blackbody emission. Due to the small uncertainties on the spectral data, the weak steepening in spectral shape seen below 3 keV can skew the fit to the spectrum, hence we cut off the spectral fit at 3 keV. However, as shown by Uttley et al. (2011), the lags are not significantly affected by the disc at energies down to 2 keV, hence we include photons down to this energy in the lag determination, in order to increase signal-to-noise.

Table 2.1: Spectral and lag fit parameters using the model REFLAGS in the band 3–10 keV. Hereinafter, we quote the total χ^2 as well as the contribution of the χ^2 that corresponds to both the spectrum and lags, separately.

Parameter	Obs. A	Obs. B
$R_{\rm in}/R_{\rm g}$	<130	<220
$R_{\rm out}/R_{\rm g}$	$(16.8^{+2.4}_{-2.7})\times10^3$	$(20^{+50}_{-8})\times10^3$
$H_{\rm out}/R_{\rm out}$	>0.83	>0.46
$H_{\rm src}/R_{\rm g}$	>380	>330
γ	$2.18^{+0.19}_{-0.24}$	>1.69
i	$39.7^{+0.4}_{-0.5}$	$29.8^{+1.8}_{-4.3}$
Γ	$1.481^{+0.003}_{-0.002}$	$1.468^{+0.004}_{-0.006}$
$\log\xi$	$2.05^{+0.03}_{-0.02}$	$2.03^{+0.02}_{-0.01}$
$\Omega^{\rm eff}/(2\pi)$	1.13	1.28
χ^2 (spectrum)	1405	1305
χ^2 (time-lags)	152	32
χ^2/dof	1557 / 1417	1337 / 1416

The best-fitting parameters for fits A and B are shown in Table 2.1. The corresponding model comparisons to the data, including the data-to-model ratios are plotted in Figs 2.3 and 2.4. These initial fits suggest that the model can represent the 2009 (B) data better. The upper limit on the disc inner radius is below 220 $R_{\rm g}$ in both cases. This value does not give us any information on whether the accretion disc in GX 339–4 is truncated or not (for discussions about disc truncation, see e.g. Tomsick et al. 2009).

As for the values of the outer radius, its value could reach up to $\sim 8\times10^4$ $R_{\rm g}$ in B, much larger than what is found in A, indicating a larger contribution to the narrow iron line in this case. This is in agreement with the stronger core of the line in B as seen in the ratio plot (Fig. 2.2) as well as the difference between the solid angles. An extremely high value for the source height $H_{\rm src}$ and a high value for the $H_{\rm out}/R_{\rm out}$ ratio suggest that the fits are being driven by the lags, whose amplitude is strongly dependent on the distances to the furthest regions of the reflector, even if their contribution is small. The ionisation parameter remains consistent between the observations, although a visual inspection of the residuals and a clear difference in χ^2 are indicative of a disc that can be described with a line of ionised iron, and whose contribution would likely come from the inner disc regions of the accretion disc where the ionisation could plausibly be larger. The blue wing of a relativistically broadened disc line could conceivably also contribute to the residuals which we fit with a narrow line, however since the inner disc radius is not strongly constrained by the lags, the model is already relatively free to fit this feature with relativistic emission, by being driven by the spectral data alone. The fact that it does not fit these residuals suggests that a more complex ionisation structure of the disc is likely (also see Wilkinson (2011)).

However, as seen in the right panel of Fig. 2.3, the model is clearly under-predicting the lags that are observed by a factor of ~ 2, mainly at low frequencies that would correspond to a light-crossing time expected from distant reflection. Given that our focus is on the lags and our aim is to understand whether they are compatible with being caused by reflection, we

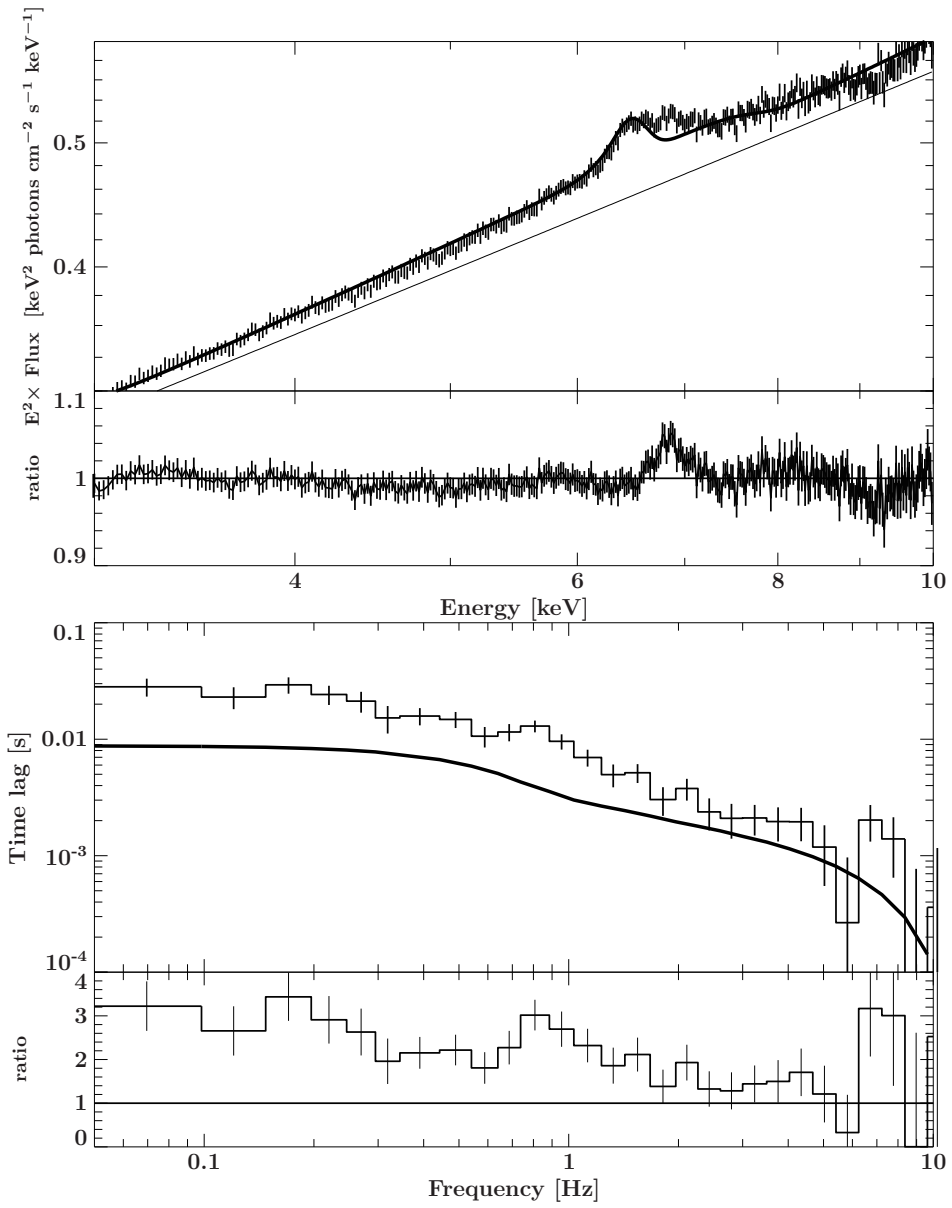

Figure 2.3: Obs. A (2004) spectrum (top) and lags vs frequency (bottom). The thick solid line represents the model REFLAGS for the best-fitting parameters (see Table 2.1, left column). In the spectrum, the thin solid line represents the direct continuum power-law. The lower panels show the ratio data/model.

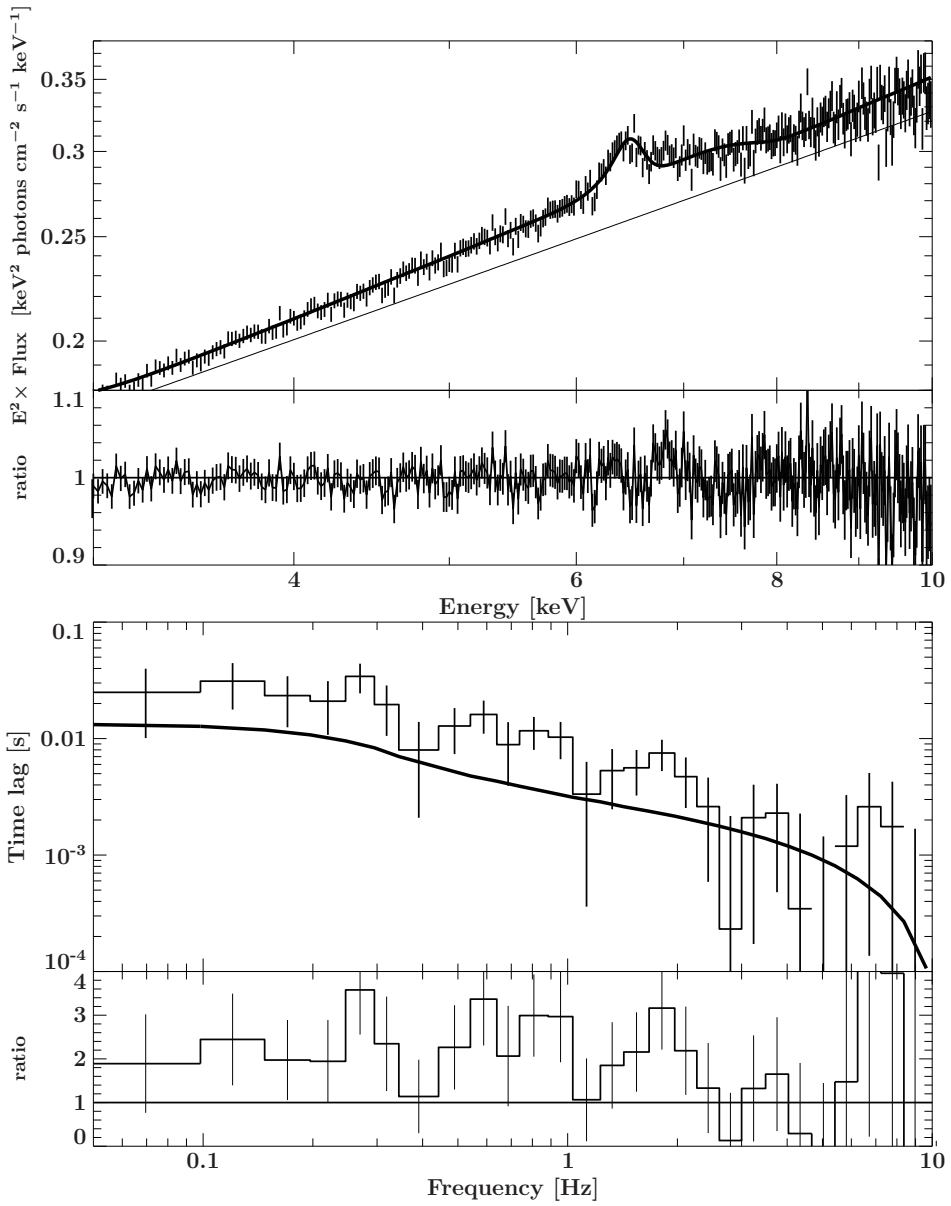

Figure 2.4: Obs. B (2009) spectrum (top) and lags vs frequency (bottom). The thick solid line represent the model REFLAGS for the best-fitting parameters (see Table 2.1, right column).

test whether the lags are being constrained by the line shape by adding an extra Gaussian component of variable width to the spectral model for the 2004 data only, that corresponds to the dataset with higher signal-to-noise. In this way, we improve the fit 'artificially' so that the model can find the best-fitting contribution to the narrow 6.4 keV core which might originate at larger radii, and thus we allow the model more freedom to produce larger lags which better fit the data.

The best-fitting parameters after this procedure can be found in the first column of Table 2.2, and an improvement in χ^2 is clear. The centroid energy for the additional emission line does not correspond to any ionised iron fluorescence transition, hence its value may come from a weighted mean of lines with a complex ionisation.

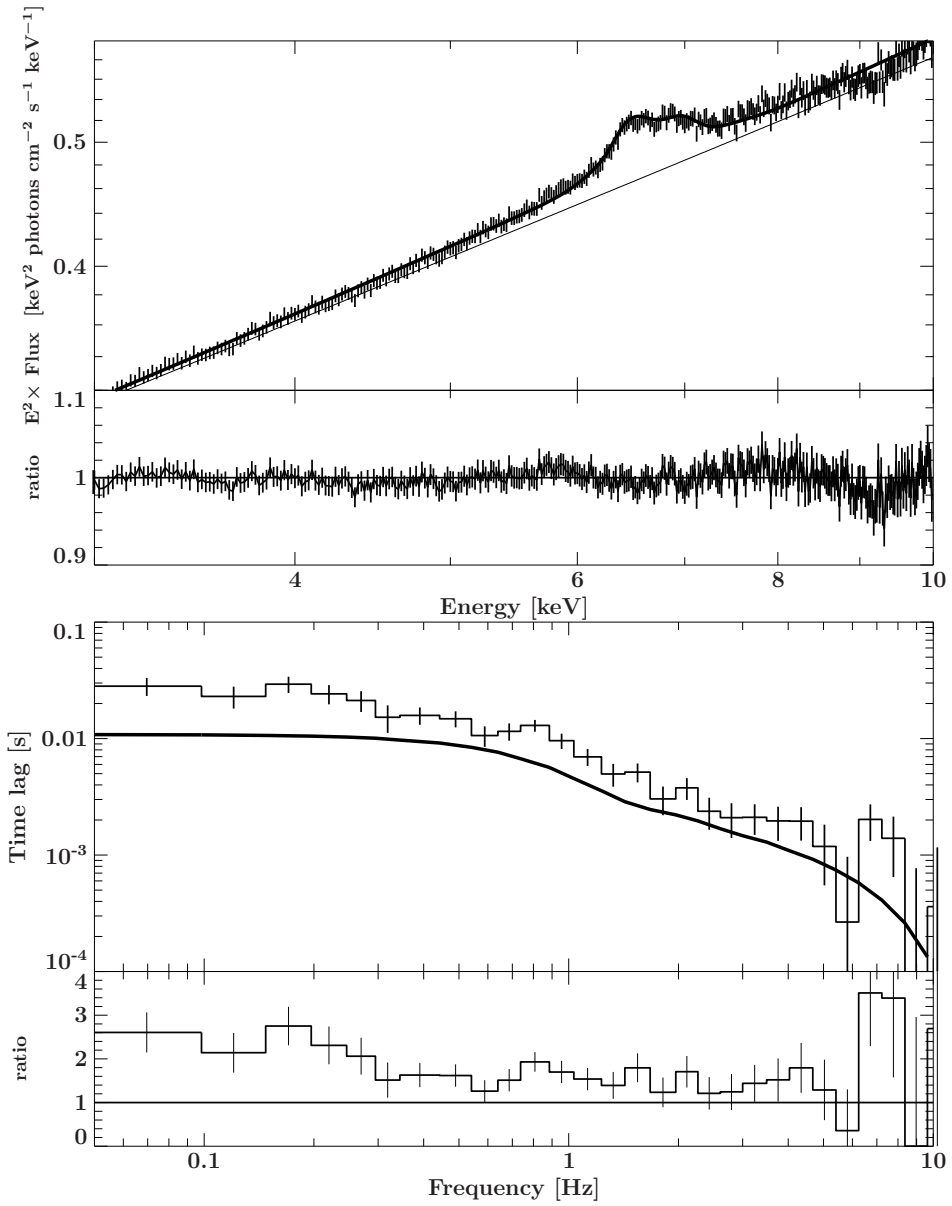

Figure 2.5: Obs. A (2004) spectrum (top) and lags vs frequency (bottom). The thick solid lines represent the model REFLAGS (including a Gaussian component in the case of the spectrum) for the best-fitting parameters (see Table 2.2, left column).

The decrease in the REFLAGS solid angle is compatible with the addition of the extra line, i.e. the main model does not need to fit it anymore. Despite the attempt to artificially improve the spectral fit to allow the model to better fit the large scale reflection, the lags from the model are still too low (Fig. 2.5).

Given the fact that the lags in 2004 and 2009 are very similar despite a factor of 2 smaller error bars in 2004, we also try to fit the 2009 spectrum swapping 2009 and 2004 lags (B', hereinafter). This worsens the fit with respect to fit B (Fig. 2.6, second column in Table 2.2), as the smaller error bars in the 2004 lags push the geometrical parameters to more extreme values, resulting in an increase of the residuals around the iron line and a slightly larger width of the line

Table 2.2: Spectral and lag fit parameters using the model REFLAGS plus an extra Gaussian fitted to 2004 spectrum and lags (obs. A), and mixing the 2009 observation spectrum with the 2004 observation lags (obs. B').

Parameter	Obs. A + extra line	Obs. B'
$R_{\rm in}/R_{\rm g}$	< 90	< 120
$R_{\rm out}/R_{\rm g}$	$(8.8^{+3.8}_{-1.6}) \times 10^3$	$(11.0^{+2.6}_{-1.9}) \times 10^3$
$H_{\rm out}/R_{\rm out}$	0.84 ± 0.04	$0.95^{+0.08}_{-0.14}$
$H_{\rm src}/R_{\rm g}$	430^{+90}_{-80}	> 380
γ	$2.22^{+0.21}_{-0.33}$	$2.62^{+0.36}_{-0.43}$
i	$48.9^{+1.0}_{-0.6}$	$42.3^{+1.3}_{-0.8}$
Γ	1.479 ± 0.003	1.46 ± 0.05
$\log\xi$	$1.50^{+0.12}_{-0.24}$	$1.65^{+0.08}_{-0.17}$
$E_{\rm c}$ [keV]	6.92 ± 0.02	–
$\sigma_{\rm c}$ [keV]	$0.17^{+0.03}_{-0.02}$	–
$\Omega^{\rm eff}/(2\pi)$	0.78	1.03
χ^2 (spectrum)	1132	1335
χ^2 (time-lags)	107	77
χ^2/dof	1239 / 1414	1412 / 1417

Table 2.3: Spectral and lag fit parameters using the model REFLAGS plus an extra Gaussian, and mixing the 2009 observation spectrum with the 2004 observation lags in the range 1–10 Hz

Parameter	Obs. A + extra line	Obs. B'
$R_{\rm in}/R_{\rm g}$	< 60	< 22
$R_{\rm out}/R_{\rm g}$	$(8.2^{+14.2}_{-3.6}) \times 10^3$	$(9.1^{+12.2}_{-3.5}) \times 10^3$
$H_{\rm out}/R_{\rm out}$	> 0.06	< 0.05
$H_{\rm src}/R_{\rm g}$	> 80	> 80
γ	> 1.1	> 1.0
i	< 31.0	< 30.2
Γ	1.459 ± 0.003	$1.4398^{+0.003}_{-0.009}$
$\log\xi$	< 1.34	< 1.31
$E_{\rm c}$ [keV]	6.85 ± 0.02	–
$\sigma_{\rm c}$ [keV]	0.25 ± 0.03	–
$\Omega^{\rm eff}/(2\pi)$	0.81	0.86
χ^2 (spectrum)	1108	1449
χ^2 (time-lags)	74	78
χ^2/dof	1182 / 1402	1527 / 1405

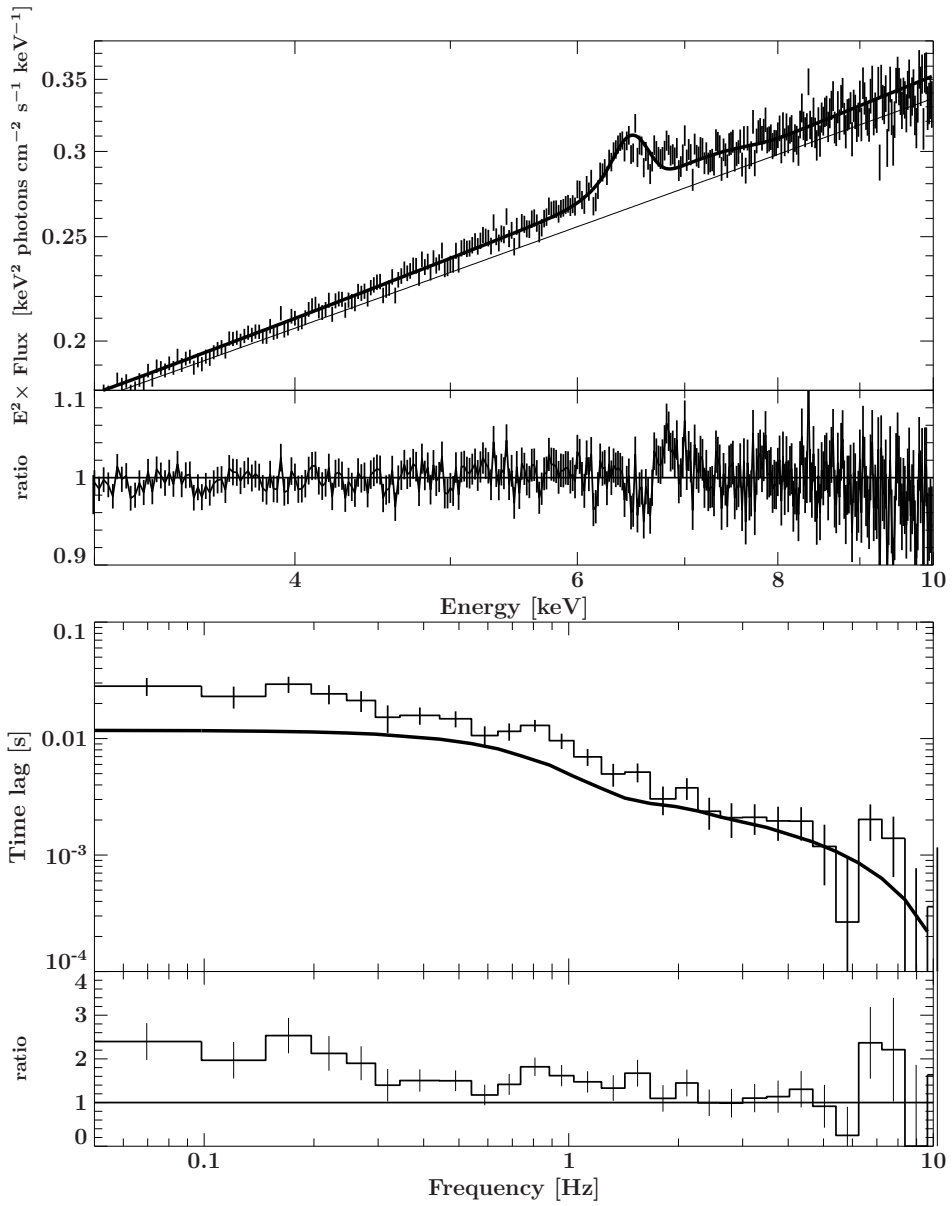

Figure 2.6: Obs. B' (2009 spectrum, 2004 lags) spectrum (top) and lags vs frequency (bottom). The thick solid lines represent the model REFLAGS for the best-fitting parameters (see Table 2.2, right column).

due to the increase of the flaring parameter accompanied by a decrease of the outer radius. These changes also increases the frequency at which the lag-frequency dependence becomes constant, which corresponds to the longest time-scale of variations in the response function).

The disc parameters in all these fits lead to a geometry where the source height, disc flaring as well as the $H_{\rm out}/R_{\rm out}$ ratio at the outer radius are extreme. This is due to the fact that the lags are driving the fits towards increased outer reflection to increase their amplitude, however the spectral model is not necessarily sensitive to small X-ray flux contributions from outer radii or the change in iron line shape that results (since the different line widths contributed by radii beyond a few thousand $R_{\rm g}$ cannot be resolved by the EPIC-pn detector).

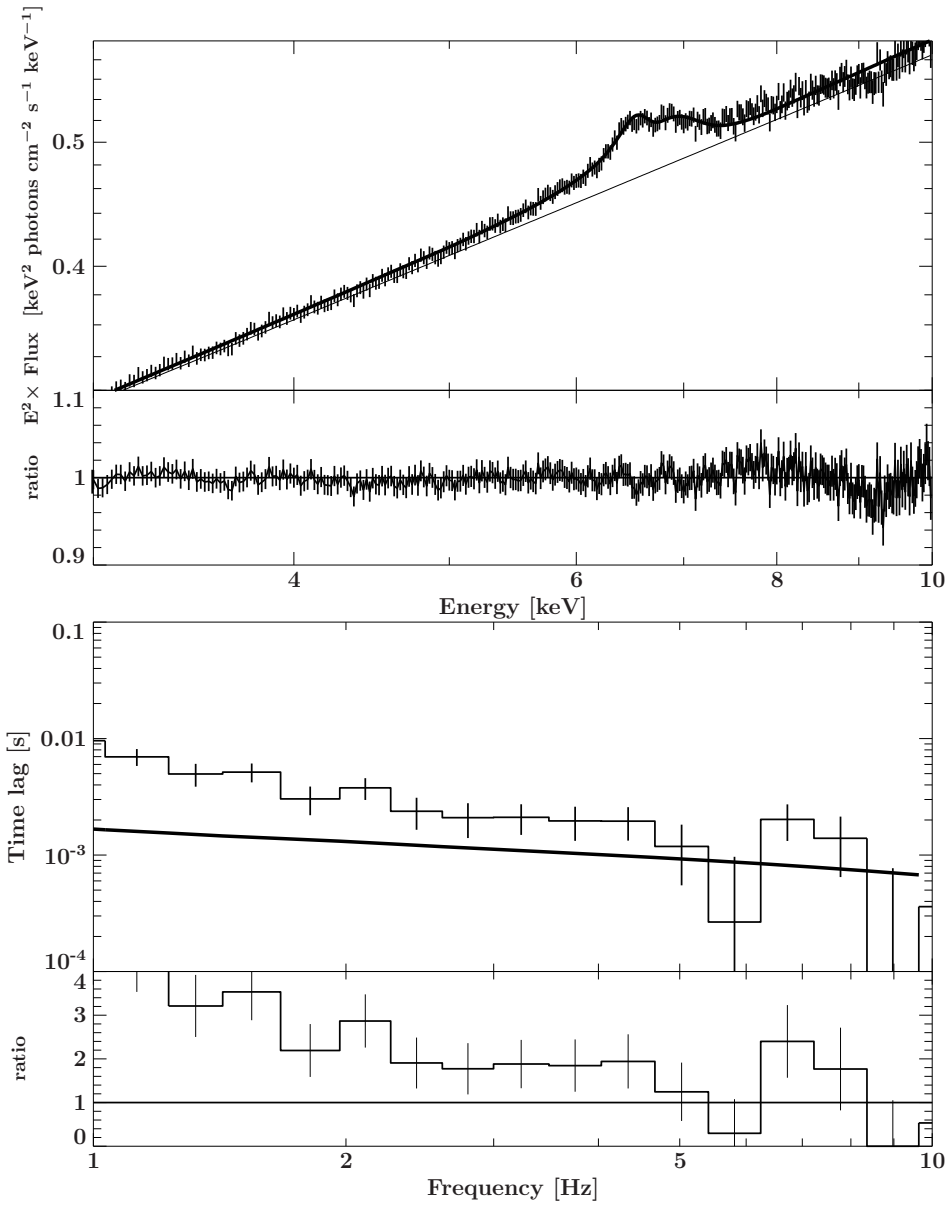

Figure 2.7: Obs. A (2004 spectrum (top) and lags vs frequency (bottom). The thick solid lines represent the model REFLAGS (including a Gaussian component in the case of the spectrum) for the best-fitting parameters (see Table 2.3, left column), in the frequency range 1–10 Hz.

In a scenario where the lags are caused by reflection only at higher frequencies (e.g. with a propagation model explaining the lower-frequency lags) the lags would correspond to reflecting regions close to the central source of direct emission. If this is the case, excluding the lags below a certain frequency would lead to a lower and more plausible $H_{\rm out}/R_{\rm out}$. We therefore exclude lags below 1 Hz from the fit (Table 2.3), since this frequency roughly corresponds to the threshold frequency below which the lags relative to softer energies are consistent with being due to propagation of disc fluctuations (Uttley et al., 2011). In addition, this range approximately corresponds to a range in frequency of the PSD (Fig. 2.9) where the PSD has a similar shape in both observations. Figs 2.7 and 2.8 show that when lags at frequencies < 1 Hz are excluded, the model cannot successfully reproduce either the lag shape or amplitude.

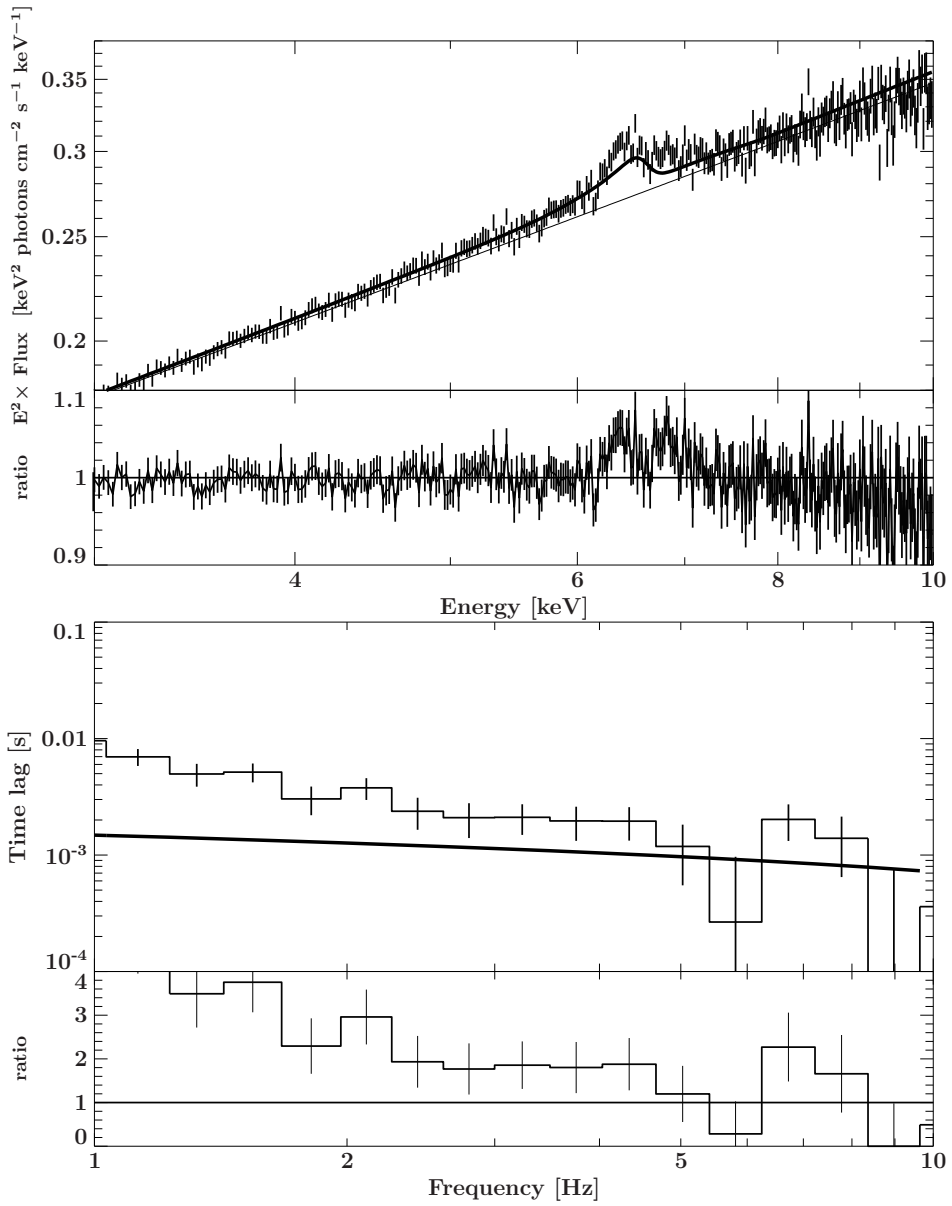

Figure 2.8: Obs. B' (2009 spectrum, 2004 lags) spectrum (top) and lags vs frequency (bottom). The thick solid lines represent the model REFLAGS for the best-fitting parameters (see Table 2.3, right column), in the frequency range 1–10 Hz.

2.4.1 Consistency of the reflection lags model with optical/UV data

The parameters inferred for the best-fitting reflection model, which are quite extreme and still do not provide a good fit to the lags, can also be checked for the implied effect on optical and UV emission from GX 339–4. X-ray heating of the outer disc could in principle produce a large optical/UV flux if there is a large solid angle illuminated by the continuum, as is the case for the geometry inferred from our lag model fits. We can assume that each illuminated cell in the disc absorbs a fraction of the illuminating continuum equal to $1-a(E)$, so that the

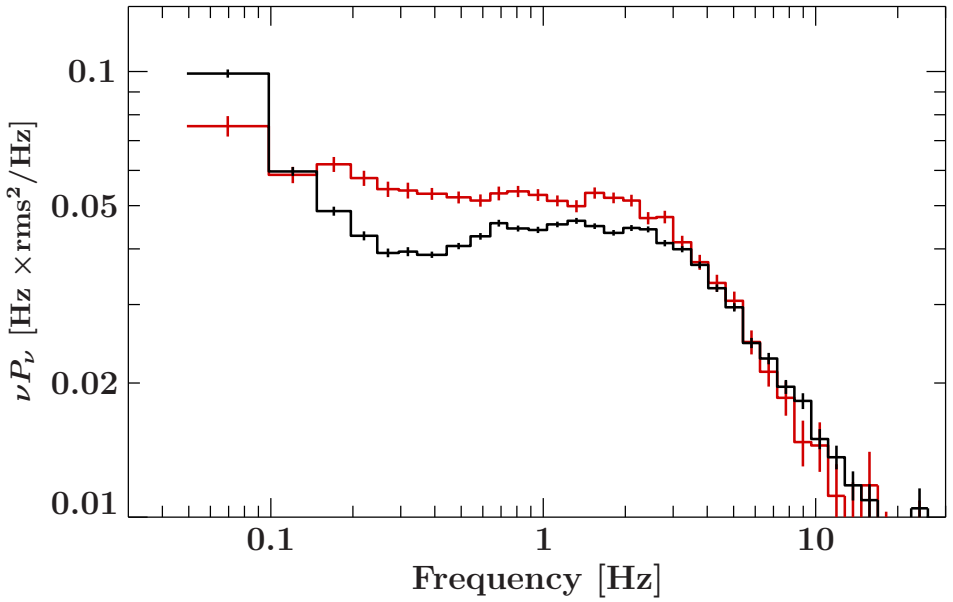

Figure 2.9: Power spectral densities for the 2004 (A, black) and 2009 (B, red) observations, in the energy range 2.0 – 10.0 keV

incident luminosity absorbed by each cell can be calculated for the best-fitting given continuum shape and model ionisation parameter. If we make the simplifying assumption that the absorbed luminosity dominates over any intrinsic blackbody emission, we can equate the luminosity that is re-emitted by the cell to the absorbed luminosity and so determine the temperature of blackbody radiation emitted by each cell, and hence determine the total reprocessed contribution to the Spectral Energy Distribution (SED).

To compare the predicted contribution to the optical/UV SED from the geometry required by the lags model, we have extracted *Swift*/UVOT (bands ubb, um2, uuu, uvv, uw1, uw2, uwh) spectra from a 1760 s observation of GX 339–4 made on 29 March 2009, two days after *XMM-Newton* observed the source. `uvot2pha` was used to extract spectra for source and background using a 6 arcsec radius, as well as extract response files. No additional aspect correction was required.

From the best-fitting spectral parameters found in fit B' and assuming a high-energy cut-off at 100 keV (Motta et al., 2009), we derive an X-ray luminosity of 3.5×10^{37} erg s^{-1} (assuming $d = 8$ kpc; Zdziarski et al. 2004). This value can be used to predict the reprocessed fluxes that are consistent with the geometries inferred from the fits, accounting for interstellar extinction using the XSPEC model `redden`.

Fig. 2.10 shows the expected reprocessed spectra for the two fits to the 2009 X-ray data (orange for B, red for B', black is the 2009 X-ray spectrum). The photon index Γ of the power-law that characterises the reprocessed spectra at the energies covered by the UVOT data is -0.97.

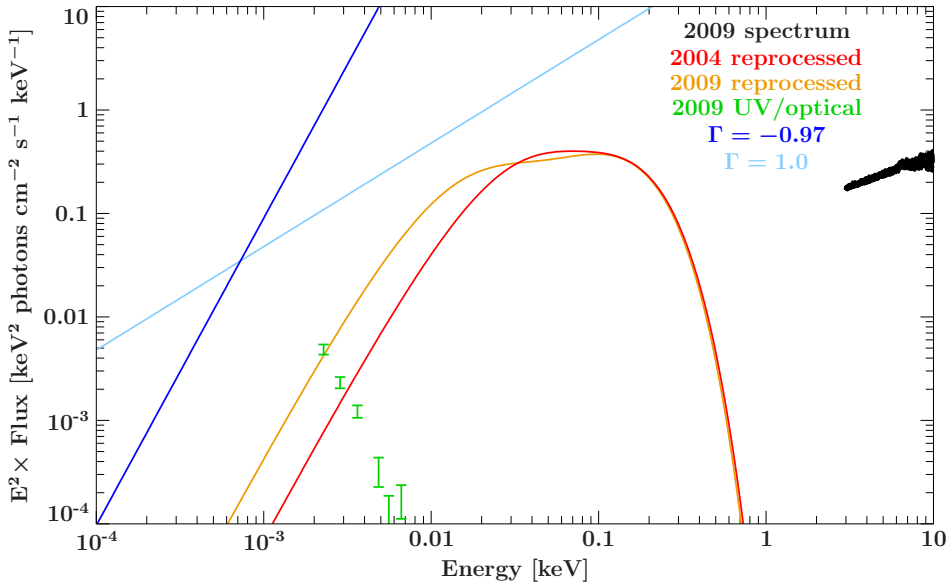

Figure 2.10: Comparison between expected reprocessed emission using 2009 spectrum with 2004 lags (red) and 2009 lags (orange) and power-law fits to the data with photon indices −0.97 (same as low energy tail of reprocessed spectra, in blue) and 1.0 (as expected from synchrotron emission coming from a jet, in light blue) after removing extinction. The data points from UVOT are shown in green.

In these energy ranges, dust extinction needs to be taken into account using the multiplicative model `redden` (Cardelli et al., 1989). E(B-V) = 0.933 is the value for the extinction calculated from IR dust maps along the line of sight towards our source[2]. By fitting a power-law with a photon index of -0.97 to the UVOT data, one finds the unabsorbed intrinsic power-law depicted in blue, which has a normalisation in the UVOT energy range that is several decades larger than that expected from reprocessing, and requires E(B-V) = 1.587 for χ^2/dof=0.8/4. Therefore the model severely underpredicts the observed flux. A more likely explanation for the optical/UV emission is flat-spectrum synchrotron emission from a compact jet (Maitra et al., 2009), or magnetised hot accretion flow (Veledina et al., 2011). Assuming a power-law photon index of 1.0 (spectral energy index of 0), one also obtains a good fit (χ^2/dof=0.64/4) and a lower extinction than in the previous case, E(B-V) = 1.089 (light blue line). This is consistent with the results found by Maitra et al. (2009) for the same object, fitting broadband data using only synchrotron and inverse-Compton models. Therefore, the UVOT data cannot be explained solely by reprocessing in the flared disc envisaged by the lags model and is more likely to be produced by a synchrotron process. However, a small contribution from reprocessing cannot be ruled out.

[2] `http://irsa.ipac.caltech.edu/`

2.5 Discussion

2.5.1 Physical implications

The analysis shown in Section 2.4 shows that extreme disc geometries are preferred for the model to reproduce the observed lags as closely as possible within the constraints given by the overall spectral shape, including the strength of the iron feature at 6.4 keV. There are several effects of the reprocessing geometry on the spectral and lag data that need to be highlighted to understand why the spectrum can be fitted well, whereas the lags cannot.

Firstly, while the amount of flux in the iron line is proportional to the solid angle subtended by the disc, the ability to determine how much of it is produced in the outer radii of the disc (where Doppler effects are weak) is limited by the resolution of the *XMM-Newton* EPIC-pn detector. Therefore, the description of the geometry that could be inferred by the spectral modelling alone is degenerate, since line emission from the largest radii (e.g. $\sim 10^5\ R_{\rm g}$) cannot be resolved from that at more modest (but still large) radii (e.g. $\sim 10^4\ R_{\rm g}$). The result of this effect is that the spectral fits are not sensitive to variations in the outer radius of the reflector at large radii.

On the other hand, the lags are very sensitive to the geometry at large radii. Firstly, the lags at low frequencies increase with both the solid angle and light-travel time to the reflector at large radii. A larger outer radius and more flared disc therefore corresponds to larger lags. However, the size-scale of the largest radius also corresponds to a characteristic low-frequency flattening in the lag versus frequency dependence. This is because the frequency-dependent drop in lags seen at higher frequencies is caused by smearing of the reflection variability on time-scales shorter than the light-travel size-scale of the reflector. The reflection variability amplitude is not smeared out for variability time-scales significantly longer than the light-travel time to the largest disc radii, and the lag at low frequencies quickly approaches the average light-travel delay from the reflector (diluted by the direct continuum emission which has zero intrinsic lag). The frequency of this characteristic flattening in the lag-frequency dependence is therefore a sensitive indicator of the size-scale of the reflector. However, it is not possible to reconcile the position of this flattening at frequencies ~ 0.2 Hz with the large amplitude of the lags at low frequencies, which imply an even larger reflector subtending an even greater solid angle to the continuum at large radii.

The result of these effects is that despite the already extreme inferred geometries the model is clearly underpredicting the lags. At this point, the maximum value of the lags is now constrained by the spectral modelling, which cannot place a tight constraint on the disc outer radius but does limit the solid angle of the reflector. It is instructive to consider the effects on the predicted spectral shape when the model is fitted to the 2004 lags alone (fixing $\Gamma = 1.5$) and the same parameters are used to estimate the resulting spectrum. This yields

χ^2/dof=66.4/16=4.15 for the fit to the lags and yields an apparent solid angle $\Omega^{\rm eff}/2\pi = 1.51$ subtended by the disc. The resulting spectral shape is compared to the data in Fig. 2.11. Therefore, even fitting the lag data alone with the model cannot produce a good fit, and the fit that is obtained shows that much larger solid angles of large-scale reflection are required than are permitted by the spectrum.

The lags cannot be explained solely by reflection, it is therefore necessary to invoke an additional mechanism to explain them. This result is perhaps not surprising, since we have previously found evidence that in hard state BHXRBs, fluctuations in the accretion disc blackbody emission are correlated with and precede the variations in power-law emission (Uttley et al., 2011). Although the disc variations seem to drive the power-law variations, this does not in itself explain the lags within the power-law band, which we consider in this Chapter, since the disc emission only extends up to ~ 2 keV. However, as we noted in Uttley et al. (2011), at frequencies < 1 Hz, the lags of the power-law emission relative to the disc-dominated 0.5–0.9 keV band show a similar frequency-dependence to the lags seen within the power-law band (i.e. $\tau \propto \nu^{-0.7}$). This strongly suggests that the lags intrinsic to the power-law are somehow connected to the mechanism which causes the power-law to lag the disc, most probably due to the propagation of accretion fluctuations through the disc before reaching the power-law emitting hot flow.

One possibility is that the disc is sandwiched by the hot-flow/corona which produces power-law emission which becomes harder towards smaller radii, leading to hard lags as fluctuations propagate inwards (e.g. Kotov et al. 2001; Arévalo and Uttley 2006). This model can explain the hard lags in terms of propagation times in the flow, which are much larger than light-crossing times and so can produce relatively large lags which the reflection model struggles to produce without leading to solid angles of large-scale reflection. Reflection may also contribute to the lags at some level, but is not the dominant mechanism, at least at frequencies < 1 Hz.

A more detailed analysis of the contribution of reflection to the observed lags could be performed using datasets with higher signal-to-noise, by e.g. searching for reflection signatures around the iron line investigated by Kotov et al. (2001). Lags vs. energy spectra for GX 339–4 are shown in Uttley et al. (2011) for the 2004 *XMM-Newton* observation, and demonstrate that the current quality of the data is not sufficient to detect these features.

It is also possible that multiple distinct components contribute to the lag, e.g. associated with the different Lorentzian features that contribute to the PSD. This possibility could explain the apparent 'stepping' of the lag vs. frequency that appears to be linked to the frequencies where the dominant contribution to the PSD changes from one Lorentzian component to another (Nowak, 2000).

2.5.2 Wider implications of combined spectral-timing models

In this work we have considered (and ruled out) a relatively simple model for the lags in terms of the light-travel times from a large scale reflector. However, it is important to stress the generalisability of our approach to other models. In particular, we have shown how it is possible to combine timing and spectral information to fit models for the geometry and spatial scale of the emitting regions of compact objects. Previous approaches to use the information from time-lags to fit models for the emitting region have focussed on fitting the lag data (e.g. Kotov et al. 2001; Poutanen 2002). However, since these models also make predictions for the spectral behaviour, it is possible to achieve stronger constraints on the models by fitting the lags together with the time-averaged energy spectrum, as done here, or with spectral-variability products such as the frequency-resolved rms and covariance spectra (Revnivtsev et al., 1999; Wilkinson and Uttley, 2009; Uttley et al., 2011).

In order to use these techniques more generally, one needs to calculate the energy-dependent response function for the emitting region, i.e. determine the emission as a function of time delay and energy. This approach can be used to test reverberation models for the small soft lags seen at high frequencies in AGN (Fabian et al., 2009; Zoghbi et al., 2010; de Marco et al., 2011; Emmanoulopoulos et al., 2011) and BHXRBs (Uttley et al., 2011), which offers the potential to map the emitting region on scales within a few gravitational radii of the black hole. Other models can also be considered, e.g. to test the propagation models for the low-frequency lags with the time-delay expressed in terms of propagation time through the accretion flow. Future, large-area X-ray detectors with high time and energy-resolution, such as the proposed *ATHENA* and *LOFT* missions, will allow much more precise measurements of the lags in combination with good spectral measurements, so that fitting of combined models for spectral and timing data could become a default approach for studying the innermost regions of compact objects. Future research in this direction is strongly encouraged.

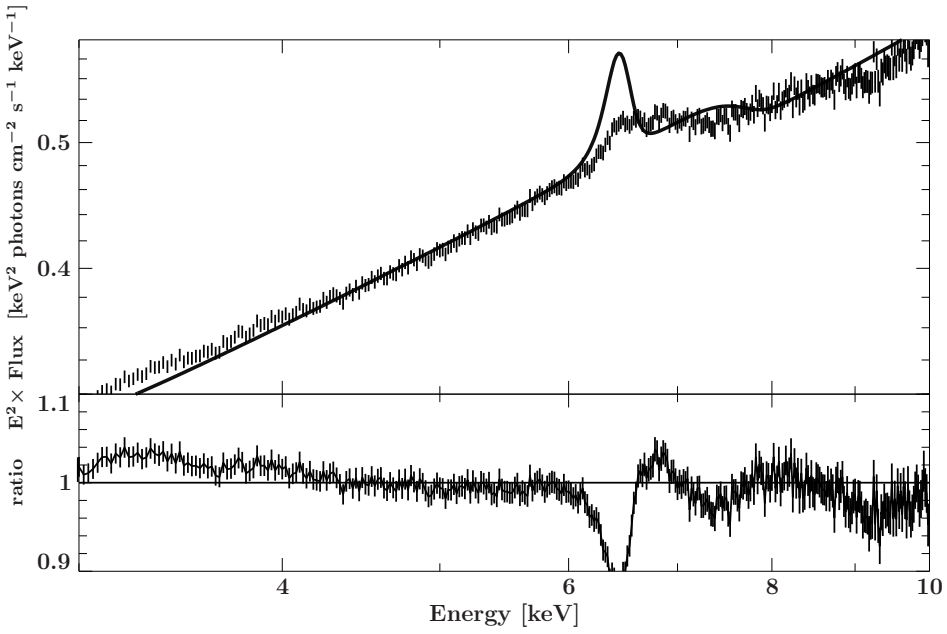

Figure 2.11: Spectral model (REFLAGS) compared to 2004 spectrum using the parameters that are derived from fitting the 2004 lags alone. The lower panel shows the ratio data/model.

CHAPTER 3

Accretion flow diagnostics with X-ray spectral-timing: the hard state of SWIFT J1753.5–0127

Recent *XMM-Newton* studies of X-ray variability in the hard states of black hole X-ray binaries (BHXRBs) imply that the variability is generated in the 'standard' optically-thick accretion disc that is responsible for the multicolour black-body emission. The variability originates in the disc as mass-accretion fluctuations and propagates through the disc to 'light up' inner disc regions, eventually modulating the power-law emission that is produced relatively centrally. Both the covariance spectra and time lags that cover the soft band strongly support this scenario.

Here we present a comparative spectral-timing study of *XMM-Newton* data from the BHXRB SWIFT J1753.5–0127 in a bright 2009 hard state with that from the significantly fainter 2006 hard state, to show for the first time the change in disc spectral-timing properties associated with a global increase in both the accretion rate and the relative contribution of the disc emission to the bolometric luminosity.

We show that, although there is strong evidence for intrinsic disc variability in the more luminous hard state, the disc variability amplitude is suppressed relative to that of the power-law emission, which contrasts with the behaviour at lower luminosities where the disc variability is slightly enhanced when compared with the power-law variations. Furthermore, in the higher-luminosity data the disc variability below 0.6 keV becomes incoherent with the power-law and higher-energy disc emission at frequencies below 0.5 Hz, in contrast with the coherent variations seen in the 2006 data. We explain these differences and the associated complex lags in the 2009 data in terms of the fluctuating disc model, where the increase in accretion rate seen in 2009 leads to more pronounced and extended disc emission. If the

variable signals are generated at small radii in the disc, the variability of disc emission can be naturally suppressed by the fraction of unmodulated disc emission arising from larger radii. Furthermore the drop in coherence can be produced by disc accretion fluctuations arising at larger radii which are viscously damped and hence unable to propagate to the inner, power-law emitting region.

3.1 Introduction

Variability on a broad range of time-scales, ranging from milliseconds to hours, is a well-known characteristic of black hole X-ray binaries (BHXRBs). The short-term variability (time-scales up to a few hundred seconds) is strongly dependent on the spectral state of the object, being particularly enhanced in the so-called hard state, with rms amplitudes up to ~ 40 per cent of the average flux (Belloni et al., 2005; Remillard and McClintock, 2006; Muñoz-Darias et al., 2011).

In the hard state, the X-ray emission is dominated by a hard power-law component (typically, $\Gamma \sim 1.4 - 2.1$) accompanied by weaker multi-colour black-body emission associated with the accretion disc, with $kT_{\rm in} < 0.5$ keV (Miller et al., 2006; Reis et al., 2010). While the source of black-body photons is strongly suspected to be the accretion disc, there is some controversy as to the origin and physical location of the power-law component, which could be produced by Compton scattering in a hot corona (Malzac and Belmont, 2009), in a hot inner flow (Zdziarski et al., 1998) resulting from the evaporation of the inner radii of the accretion disc (thus truncating the disc, see e.g. Narayan and Yi 1994), or at the base of the radio-emitting jet that is observed during this state (Markoff et al., 2005). Notwithstanding the actual physical source of hard photons (i.e. $E \gtrsim 2$ keV), there is general agreement that the hard photon-emitting region must be relatively central, within tens of gravitational radii from the central object in the bright hard state.

Previously, most X-ray variability studies of the hard state were conducted using instruments with hard X-ray sensitivity, e.g. the Proportional Counter Array (PCA) on the *Rossi X-ray Timing Explorer* (*RXTE*), which limited the study of the variability to the hard, power-law component. Due to the anti-correlation of rms variability amplitude with spectral hardness as a source transitions to the disc-dominated soft state (e.g. Belloni et al. 2005), a picture arose where the disc itself was intrinsically stable and variability is generated in the power-law emitting region, usually envisioned as an unstable hot inner flow (Churazov et al., 2001; Done et al., 2007).

However, recent work (Wilkinson and Uttley, 2009; Uttley et al., 2011) has used the soft X-ray response and timing capability of the EPIC-pn instrument on board *XMM-Newton* to carry

out the first spectral-timing studies of the disc in the hard state. This work has shown strong evidence, in the form of variability ('covariance') spectra and time-lags, that the disc is intrinsically variable, and that disc variations in fact precede (and likely drive) the power-law variations, at least on time-scales of seconds and longer. The role of the disc as the driver of variability can be explained by invoking the presence of mass-accretion rate fluctuations that originate in the disc at particular radii and propagate inwards. Such fluctuations vary at the local viscous time-scale corresponding to the radius where they are produced (Lyubarskii, 1997; Kotov et al., 2001; Arévalo and Uttley, 2006). Because the $\dot{m}$ variations at inner radii are modulated by the outer $\dot{m}$ via a multiplicative process in the context of this model, variability is seen at all frequencies that correspond to the viscous time-scales of the radii generating the propagating signals. Once the perturbations reach the inner regions of the system, these modulate the power-law emission that accounts for most of the hard X-ray flux. A direct implication of this model is that the variations in power-law X-ray emission must be observed after the disc black-body variations with a time-delay scaling with the viscous travel-time between the radii where the fluctuations originate and the power-law emitting region.

So far, evidence for disc-driven variability has been seen in *XMM-Newton* EPIC-pn observations of three hard state sources: GX 339-4, Cyg X-1 and SWIFT J1753.5–0127 (Uttley et al., 2011). These data were obtained during relatively long-lived hard states at moderate luminosities, around 1 per cent of the Eddington limit (with some uncertainty given the uncertain distances to GX 339-4 and SWIFT J1753.5–0127 as well as some uncertainty on their masses). More luminous hard states, corresponding to outburst rises or 'failed' state transitions, have previously not been studied with detailed spectral-timing into the soft X-ray band, but could contain important information about changes in the disc variability as the accretion rate increases and the disc emission strengthens. In this chapter, we carry out the first detailed spectral-timing study of a luminous hard state of SWIFT J1753.5–0127, a transient black hole candidate which was discovered on 2005 May 30 (Palmer et al., 2005). Its $\sim$ 3.2 hr period derived from optical lightcurves (Zurita et al., 2008; Durant et al., 2009) makes it the BHXRB with the second-shortest period known to date (Kuulkers et al., 2012). Since the start of the outburst, it has undergone a transition to and from the hard-intermediate state. It has never completed the outburst cycle toward softer states nor has it gone into quiescence. The long-term *Swift/BAT* lightcurve of SWIFT J1753.5–0127 is shown in Figure 3.1 and shows the epochs of the two *XMM-Newton* EPIC-pn observations. The first, obtained early in 2006, corresponds to a relatively faint hard state and shows clear evidence for disc-driven variability (Wilkinson and Uttley, 2009; Uttley et al., 2011). The second observation was triggered by us in response to a brightening of the hard state in September 2009, and enables a comparison of a more luminous hard state with the lower-luminosity states studied to date.

In Section 3.2 we describe our data reduction, and in Section 3.3 we carry out a detailed spectral-timing study of the 2009 data and compare our results with the 2006 observation. In Section 3.4 we summarise our key results and interpret them in terms of the disc fluctuation model which can explain the data obtained at lower luminosities.

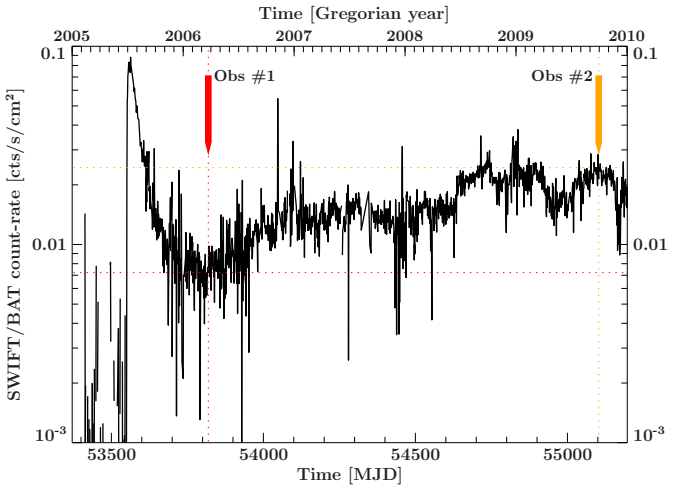

Figure 3.1: *Swift/BAT* lightcurve of SWIFT J1753.5–0127. The 2006 (#1) and 2009 (#2) *XMM-Newton* observations of the source are marked in red and orange, respectively.

3.2 Data reduction

SWIFT J1753.5–0127 was observed with *XMM-Newton* on 2009 September 29 (ObsID 0605610301) and is shown as Obs #2 on Fig. 3.1, together with the 2006 observation (Obs #1, ObsID 0311590901) that we use later for comparison. We have extracted the data from the EPIC-pn camera in Timing Mode (Kuster et al., 2002) using the SAS 10.0.0 tool `epproc` and the most recent Current Calibration Files (CCFs). We have used `evselect` for filtering and extracted only single and double events (`PATTERN <= 4`) during intervals of low, quiet background. The column filters `RAWX in [27:50]` was chosen for the 2006 observation, and the selection `RAWX in [27:35] || RAWX in [38:50]` was adopted for the 2009, thus excising the central two columns to mitigate pile-up effects. Although an excision of further columns results in a further decrease of pile-up, this effect does not change any of the results shown in this chapter. Due to telemetry drop-outs, the total useful exposure are ~ 35.5 ks (2006) and ~ 13.3 ks (2009).

The tool `evselect` was used to extract the spectrum and the response files were extracted with `rmfgen` and `arfgen`. Power spectra, cross-coherences, phase- and time-lags were extracted using data within the GTI intervals following Nowak et al. (1999) and Vaughan and Nowak (1997). We use a time bin of 0.005965 s and segment size of 16384 and 4096 bins/segment (for the 2006 and 2009 observations, respectively).

We omit any segments with gaps between successive events longer than 0.1 s. These gaps are likely to be due to telemetry drop-outs that are not always included in the GTI files. Covariance (Wilkinson and Uttley, 2009) and rms spectra (Revnivtsev et al., 1999) as well as the energy-dependent phase- and time-lags were also extracted for different frequency ranges as discussed in Section 3.3.3.

Given the Fourier transform of the signal in the i-th energy channel (or energy band) $S_i(\nu_l)$, the cross-spectrum between two bands 1 and 2 is defined as:

$$\langle C_{1,2}(\nu_l)\rangle = \langle S_1^*(\nu_l)S_2(\nu_l)\rangle \quad , \tag{3.1}$$

where the averages are performed over a number of independent segments and frequency bins m. Energy-dependent products such as phase- and time-lags as well as covariance spectra can be obtained from the above quantity by defining a reference band for each channel/band $S_{\bar{i}}(\nu_l) = \sum_{i'\neq i} S_{i'}(\nu_l)$. In reality, the reference band can be fine-tuned to incorporate a given set of channels (e.g. a soft or a hard band), as we will show in the Results section. With this definition, the cross-spectrum for the i-th channel becomes:

$$\langle C_i(\nu_l)\rangle = \langle S_{\bar{i}}^*(\nu_l)S_i(\nu_l)\rangle = \left\langle \sum_{i'\neq i} S_{i'}^*(\nu_l)S_i(\nu_l)\right\rangle \quad . \tag{3.2}$$

The energy-dependent phase-lag is then given by $\phi_i(\nu_l) = \arg\left[\langle C_i(\nu_l)\rangle\right]$ and its time-lag is $\tau_i(\nu_l) = \phi_i(\nu_l)/2\pi\nu_l$.

The covariance spectrum in the frequency range ν_{l_1} to ν_{l_2} can be defined as (Wilkinson, 2011):

$$\mathrm{N}_{i,l_1\to l_2} = \frac{\sqrt{\sum_{l=l_1}^{l_2}\left(\langle|C_i(\nu_l)|^2\rangle - B_i(\nu_l)\right)\Delta\nu_l}}{\sqrt{\sum_{l=l_1}^{l_2}\langle|S_{\bar{i}}(\nu_l)|^2\rangle\Delta\nu_l}} \quad , \tag{3.3}$$

where $\Delta\nu_l$ is the width of the l-th frequency bin and $B_i(\nu_l)$ denotes the bias caused by the contribution of Poisson power N_i, which is subtracted from the modulus-squared value of the observed cross-spectrum (see e.g. Vaughan and Nowak 1997):

$$\begin{aligned} B_i(\nu_l) =& (\langle|S_{\bar{i}}(\nu_l)|^2\rangle\langle|S_i(\nu_l)|^2\rangle + \langle|S_{\bar{i}}(\nu_l)|^2\rangle\langle N_i(\nu_l)\rangle \\ &+\langle N_{\bar{i}}(\nu_l)\rangle\langle N_i(\nu_l)\rangle)/m_l \quad , \end{aligned} \tag{3.4}$$

where m_l equals the product of the number of segments used for the averages, times the number of frequency bins that contribute to the sums over frequency bins in Eqn. 3.3. With the definition above, the amplitude of the covariance spectrum equals the value of the rms spectrum (Revnivtsev et al., 1999) when the coherence between each energy channel and their respective reference bands equals unity, and its value decreases according to the value of the coherence otherwise. The units are the same as those of the rms and mean spectra, therefore the covariance spectrum can be divided by the mean spectrum to show the amplitude of fractional correlated variability as a function of variability frequency (see e.g. Fig. 3.6).

Covariance spectra can be used to study the properties of the spectral components as they vary on different time-scales, as we shall see.

The spectral analysis was performed in ISIS 1.6.2-1 (Houck and Denicola, 2000), using a spectrum where a minimum of 3 channels per bin are grouped together, after adding 1% systematic error to account for instrumental uncertainties in the calibration of the instrument (Wilkinson and Uttley, 2009).

3.3 Spectral-timing analysis and results

3.3.1 Spectral analysis

We first apply a simple phenomenological model to the energy spectrum in order to understand qualitatively the contributions of the key continuum components, which will inform our choice of energy-bands for spectral-timing analysis. A cursory examination of the unfolded spectrum reveals a clear soft X-ray excess above the usual power-law emission, consistent with a stronger disc black-body component than is seen at lower luminosities in the hard state. Therefore we consider a model consisting of a multi-colour disc black-body emission component plus a thermal-Comptonisation component (`nthcomp`, hereinafter the 'power-law' component given the similarity in shape at the hard energies), both absorbed by neutral Galactic absorption with the model `tbnew` (Wilms et al., in prep.).[1] In ISIS, the model corresponds to `tbnew(1)*(diskbb(1)+nthcomp(1))`. In addition, we fix the neutral hydrogen column density N_H value to 0.194×10^{22} cm^{-2} (Wilkinson and Uttley, 2009) and the thermal-Comptonisation electron temperature kT_e to 53 keV (Reynolds et al., 2010). The seed photon temperature for Comptonisation is tied to the disc black-body inner disc temperature. The model is fitted only to the 0.7–10 keV energy range, following the recommendations of the EPIC Consortium not to consider the data below 0.6 keV (Guainazzi, M., 2011), however we plot the entire useful energy range covered by the data for completeness and to show clearly the predicted disc contribution.

The best-fitting parameters for the model are shown in the first column of Table 3.1 and the spectrum together with the best-fitting model is shown in Figure 3.2. Although the fitting procedure yields a high χ^2/d.o.f. ~ 2.6 (possibly because of underestimated instrumental systematic errors), our aim is not to provide an accurate spectral fit but rather to understand what emission components are the fundamental spectral sources of photons. Therefore, we do not include the non-primary spectral components in the fit, such as the disc reflection component.

[1] http://pulsar.sternwarte.uni-erlangen.de/wilms/research/tbabs/

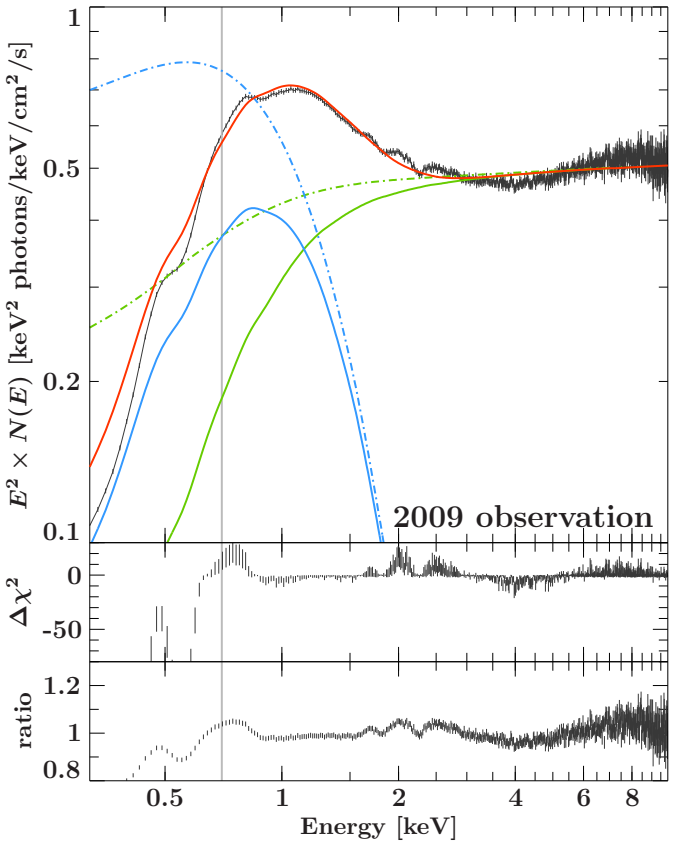

Figure 3.2: 2009 spectrum (black) and best-fitting model (red) using `tbnew(1)*(diskbb(1)+nthcomp(1))`. The χ^2 and ratio residuals are shown in the lower panels. The absorbed (solid line) and unabsorbed (dashed dotted line) contributions of the disc black-body and Comptonisation components are depicted in blue and green, respectively. The spectrum below 0.7 keV (left side of the vertical line) is shown but not used for the fit.

We use the same model to re-fit the EPIC-pn spectrum of the 2006 hard-state observation presented in Miller et al. (2006) and Wilkinson and Uttley (2009), which we will later use for a comparison with the 2009 data. The fitted parameters are shown in Table 3.1 and the model fit is plotted in Fig. 3.3. In this case we find a harder power-law ($\Gamma = 1.620 \pm 0.005$) and a much weaker disc black-body component with a lower temperature ($kT_{\rm in} = 0.212^{+0.020}_{-0.018}$) and normalisation only ~ 3 percent of the value from 2009.

3.3.2 2009 data: frequency-dependent spectral-timing properties

The presence of both a disc black-body component and a Comptonisation component that dominate different energy bands covered by EPIC-pn, as seen in Figure 3.2, suggests that the variability properties of the source may differ where the photon flux is dominated by either the disc black-body or the power-law components. Hence it is convenient to extract light curves for the energy ranges 0.3 – 0.6 keV, 0.6 – 1.0 keV, 2.5 – 4.0 keV and 4.0 – 10.0 keV,

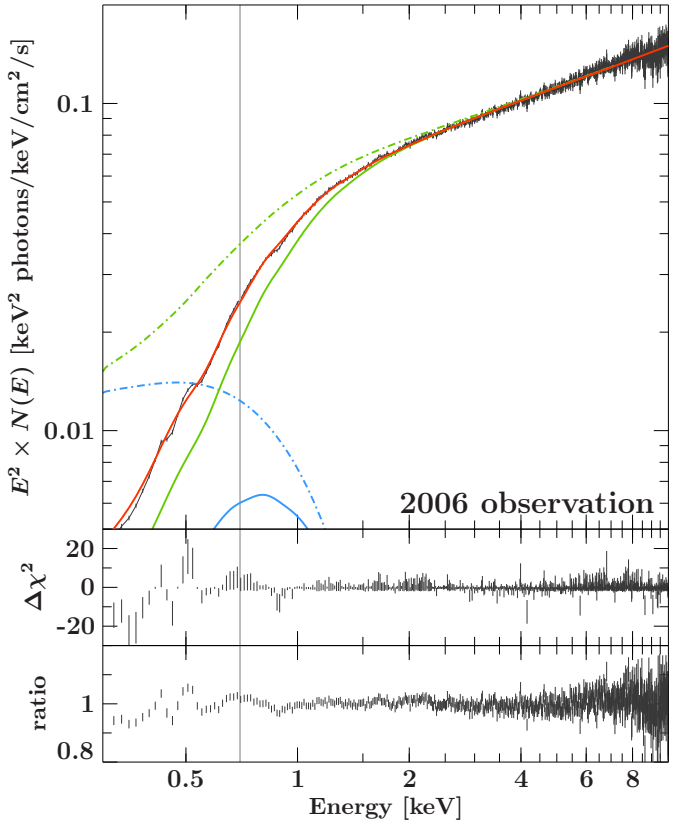

Figure 3.3: 2006 spectrum (black) and best-fitting model (red) using `tbnew(1)*(diskbb(1) +nthcomp(1))`. The χ^2 and ratio residuals are shown in the lower panels. The absorbed (solid line) and unabsorbed (dashed dotted line) contributions of the disc black-body and Comptonisation components are depicted in blue and green, respectively. The spectrum below 0.7 keV (left side of the vertical line) is shown but not used for the fit.

Table 3.1: Best-fitting values for the model `tbnew(1)*(diskbb(1)+nthcomp(1))` fitted over the range 0.7 – 10.0 keV. The disc black-body normalisation $N_{\rm dbb}$ is in units of $[(r_{\rm in}/{\rm km})/(D/10{\rm kpc})]^2\cos\theta$ (where $r_{\rm in}$ is the 'apparent' disc inner radius that is also dependent on the inclination θ of the disc, see Kubota et al. (1998)), whereas the Comptonisation normalisation is in units of photons cm^{-2}s^{-1}keV^{-1} at 1 keV.

Parameter	Value (2009)	Value (2006)
N_H [10^{22} cm^{-2}]	0.194 (fixed)	0.194 (fixed)
$kT_{\rm in}$ [keV]	0.248 ± 0.001	$0.212^{+0.020}_{-0.018}$
$kT_{\rm e}$ [keV]	53 (fixed)	53 (fixed)
$N_{\rm dbb}$	31490^{+810}_{-940}	1050^{+270}_{-50}
Γ	$1.961^{+0.004}_{-0.005}$	1.620 ± 0.005
$N_{\rm nth}$	0.429 ± 0.003	0.052 ± 0.003
χ^2/d.o.f.	1638/617	600/623

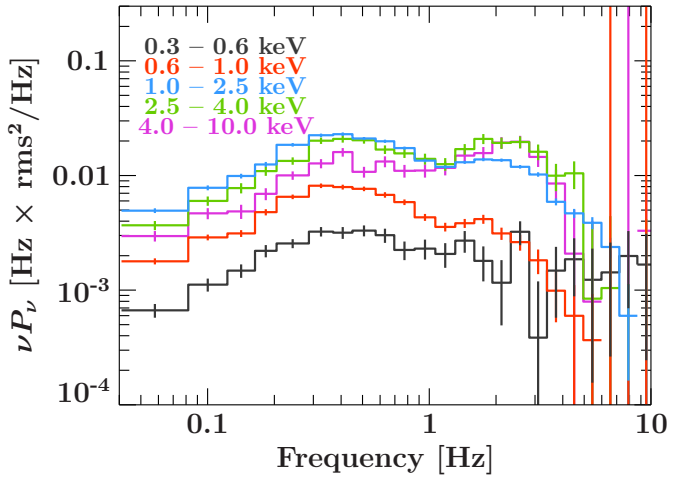

Figure 3.4: Power spectral densities (PSDs) for the 2009 observation in the energy bands 0.3 – 0.6 keV (black), 0.6 – 1.0 keV (red), 1.0 – 2.5 keV (blue), 2.5 – 4.0 keV (green) and 4.0 – 10.0 keV (magenta) showing two broad peaks whose relative amplitude varies as a function of the energy band chosen.

hereinafter ultrasoft (black), soft (red), intermediate (green) and hard (magenta) bands, respectively. For completeness we also extracted data from the the 1.0 – 2.5 keV band (blue), although since this band shows no unique properties compared to the other bands, we do not show it in our later spectral-timing results, but we do include it in showing the energy-dependent power-spectral densities (PSDs), which are shown in Fig. 3.4. As shown in Fig. 3.2, the ultrasoft and soft bands are dominated by the disc black-body component, whereas the intermediate and hard bands cover the power-law component. The range 1.0 – 2.5 keV represents the variability power for the photons emitted by a combination of the black-body and power-law components.

In addition, the ultrasoft, soft, intermediate and hard bands are also used to extract the coherence, phase- and time-lags, using the combinations ultrasoft-soft, ultrasoft-intermediate, ultrasoft-hard, soft-intermediate, soft-hard and intermediate-hard, shown in Fig. 3.5. The properties of coherence and lags with respect to the 1.0 – 2.5 keV band show an intermediate behaviour compared to other bands and are not shown here.

3.3.2.1 Amplitude of variability

The power spectral densities in Fig. 3.4 show evidence for two broad noise components, peaking at ~ 0.4 and ~ 2 Hz (hereinafter, the 'low-frequency' and 'high-frequency' peaks respectively). The relative amplitude between the two peaks changes, depending on the energy band considered. At the soft and ultrasoft energies where the disc black-body emission dominates, the low-frequency peak shows a larger peak-amplitude of variability than the high-frequency peak. Conversely, for the power-law-dominated intermediate and hard energies the high-frequency peak-amplitude becomes stronger with respect to the low-frequency

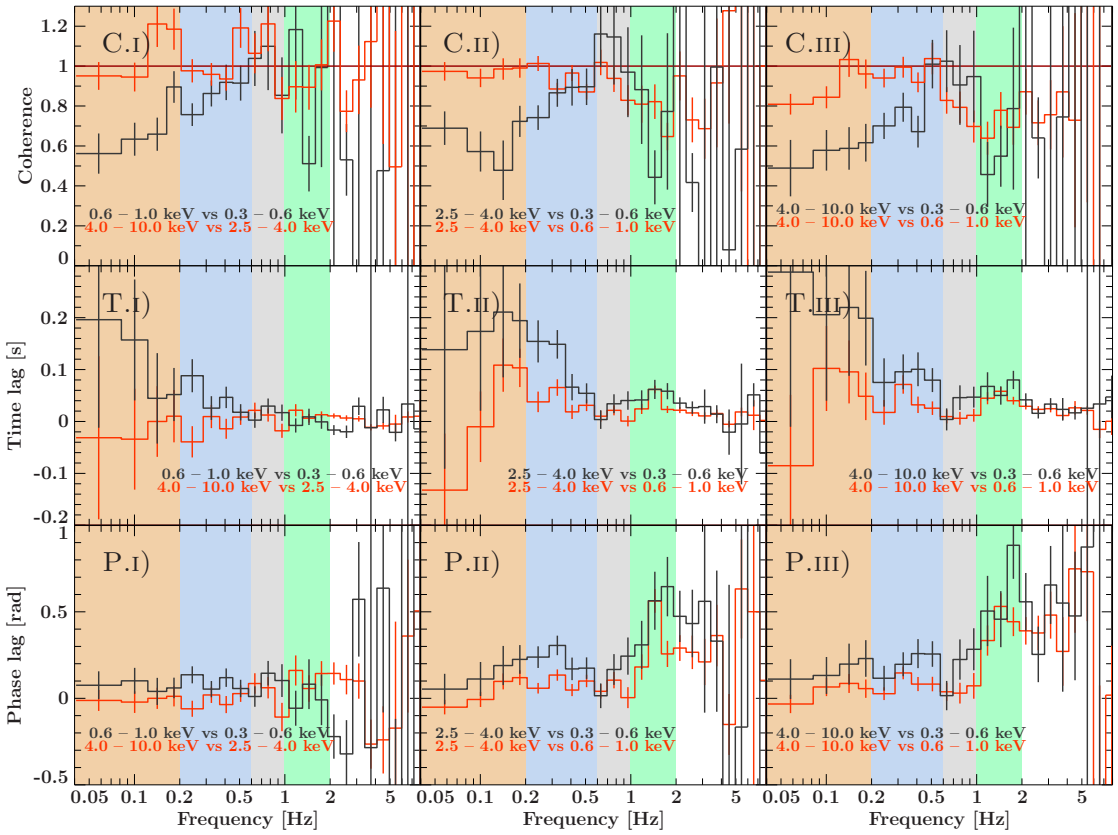

Figure 3.5: Coherence, time- and phase-lags for the 2009 observation, using different combinations of the bands 0.3 – 0.6 keV, 0.6 – 1.0 keV and 4.0 – 10.0 keV. A positive lag indicates that the emission in the harder band lags the emission in the softer band. The shaded regions indicate the frequency intervals used to average the energy-dependent products studied in Section 3.3.3.

peak and is consistent between the two bands, while the 1.0 – 2.5 keV band shows an intermediate behaviour between the softer and the higher-energy bands. Besides the energy-dependence of the PSD shape, there is also a clear change in normalisation, with the ultrasoft band showing the weakest variability at all frequencies, followed by the soft band and then the higher-energy bands. This behaviour contrasts with that seen in the 2006 data by Wilkinson and Uttley (2009), where the softer 0.5-1 keV band showed *larger-amplitude* variability than the 2-10 keV band at low frequencies.

3.3.2.2 Coherence

The coherence is an indicator of the degree of linear correlation between variations in the emission in two different bands (Vaughan and Nowak, 1997) at a particular frequency. Full coherence (when the coherence equals unity at the given frequency) is reached when the variations in emission in one band can be reconstructed linearly from the variations in emission in the other band. The upper panels in Fig. 3.5 show the coherence computed for all possible combinations of the above bands. For all plots in the figure (including time- and

phase-lags), the left panels show the frequency-dependent products computed between the pair of bands that cover the disc black-body region (black) or the pair of bands which cover the power-law region (red), i.e. the disc vs. disc behaviour and the power-law vs. power-law behaviour. The remaining panels show the same products made from intercombinations of the disc and power-law bands, combining the intermediate and hard bands (middle and right panels, respectively) together with the ultrasoft (black) and soft (red) bands that are dominated by the disc black-body emission.

In a simple case where all of the emission in one spectral component varies linearly together at all energies, the coherence between the observed emission in two bands where that component dominates should approach unity. This is similar to the case observed with the coherence between the intermediate and hard bands that sample the power-law (C.I), as well as between the intermediate (C.II) and hard bands (C.III) with respect to the soft band (red), up to ~ 0.6 Hz. This frequency roughly corresponds to the point where the low and high-frequency peaks in the PSD intersect (Fig. 3.4). The coherence deviates slightly from unity for frequencies below ~ 0.1 Hz between the hard and soft bands.

A very different behaviour is observed between the ultrasoft and the soft, intermediate and hard bands. In this case, unity coherence is reached at ~ 0.6 Hz, and decreases towards *both* lower and higher frequencies, reaching values of ~ 0.5 in some cases. Interestingly, the ultrasoft variations at low frequencies are just as incoherent with the soft variations as they are with the hard or intermediate variations. Thus, at low frequencies, while the variations in the two power-law bands are coherent with each other, a significant fraction of the variations in the two disc bands are incoherent with each other.

3.3.2.3 Time and phase lags

Previous observations of time-lags in the hard state show that they roughly follow a frequency-dependence $\tau(\nu) \propto \nu^{-0.7}$ (Miyamoto and Kitamoto, 1989; Miyamoto et al., 1992; Crary et al., 1998; Nowak et al., 1999), albeit with detectable 'steps' in the lag vs. frequency dependence where the data are good enough to discern them (Nowak et al., 1999). For the observation that we are studying, the frequency dependence does clearly vary as a function of energy bands chosen. For lags between the two disc-dominated bands (ultrasoft and soft) or between the intermediate and hard power-law bands, a relatively small amplitude is seen with a weak and smooth evolution of phase lag with Fourier-frequency. However, the lags between power-law and disc bands show a much more pronounced change with a significant 'step' (most clearly visible in the phase lags) above 1 Hz. Thus the strongest evolution of lags with frequency occurs when comparing the lags between the power-law and disc bands, with the power-law lagging the disc variations, as seen in other observations of hard state BHXRBs (Uttley et al., 2011).

3.3.3 2009 data: energy-dependent lags and covariance spectra

Given the complex behaviour in the coherence and lags shown in Fig. 3.5, we next pursue an energy-resolved analysis that can more clearly relate this behaviour to the different spectral components that contribute to the variability of the source.

We extract cross-spectra in the frequency ranges 0.04 – 0.20 Hz, 0.2 – 0.6 Hz, 0.6 – 1.0 Hz and 1.0 – 2.0 Hz (orange-, blue-, grey- and green-shaded regions in Fig. 3.5) that correspond to frequency ranges with comparable behaviour in terms of lags and coherence. The cross-spectra are computed between a broad, high signal-to-noise reference energy band, and each of 30 rebinned energy channels. Following Wilkinson and Uttley (2009) and Uttley et al. (2011), the reference band light curve is determined separately for each channel since it must have the signal in the channel of interest subtracted in order to avoid correlated Poisson noise effects. In these previous works, the data were largely coherent at all frequencies, so the results were not sensitive to the choice of reference band. In the observation presented here, our spectral-timing products could vary with the choice of reference band, because the lags and covariance spectra are effectively weighted by the coherence of each channel with the reference band. Therefore, to account for any effects due to the choice of reference band, we make cross-spectra for two reference bands: 0.5 – 2.5 keV and 3.1 – 10.0 keV. Following Uttley et al. (2011), we make energy-dependent lag and covariance spectra for each of these two bands, plotted as the left and right columns in Fig. 3.6.

3.3.3.1 Lag-energy spectra

The time and phase-lag versus energy are shown in the top two sets of panels in Fig. 3.6. Because the lag difference between two separate energy channels is relative to the choice of reference band, the offset on the y-axis is not meaningful in this case. The main consideration is the relative lag between channels, which is plotted here so that more positive values of lag are lagging smaller/more-negative values. We do not plot the time lags for the lowest-frequency range, since their large values would make it difficult to read the data for the other frequency ranges (these data can be seen in the phase lag plots, however).

The dependence of the lags on energy can be approximated as a 'broken' log-linear law, with a steeper slope below $\sim$ 1.0 keV and a flatter dependence above that energy. This turning point is close to the energy where the photon flux from the Comptonisation component overcomes the disc black-body photon flux (Fig. 3.2). In the frequency range 0.2 – 0.6 Hz, the lag between the power-law and disc components reaches values up to $\sim$ 0.2 s, reaching values close to zero as the frequency increases. This behaviour was also shown in Uttley et al. (2011) and is consistent with the disc variations leading (and probably driving) the power-law variations on time-scales of seconds. At frequencies of 1 – 2 Hz, the low-energy down-turn in lags is replaced by a small up-turn, also similar to what was observed for other observations of BHXRBs at similar frequencies (Uttley et al., 2011). However, it is interesting to note that in

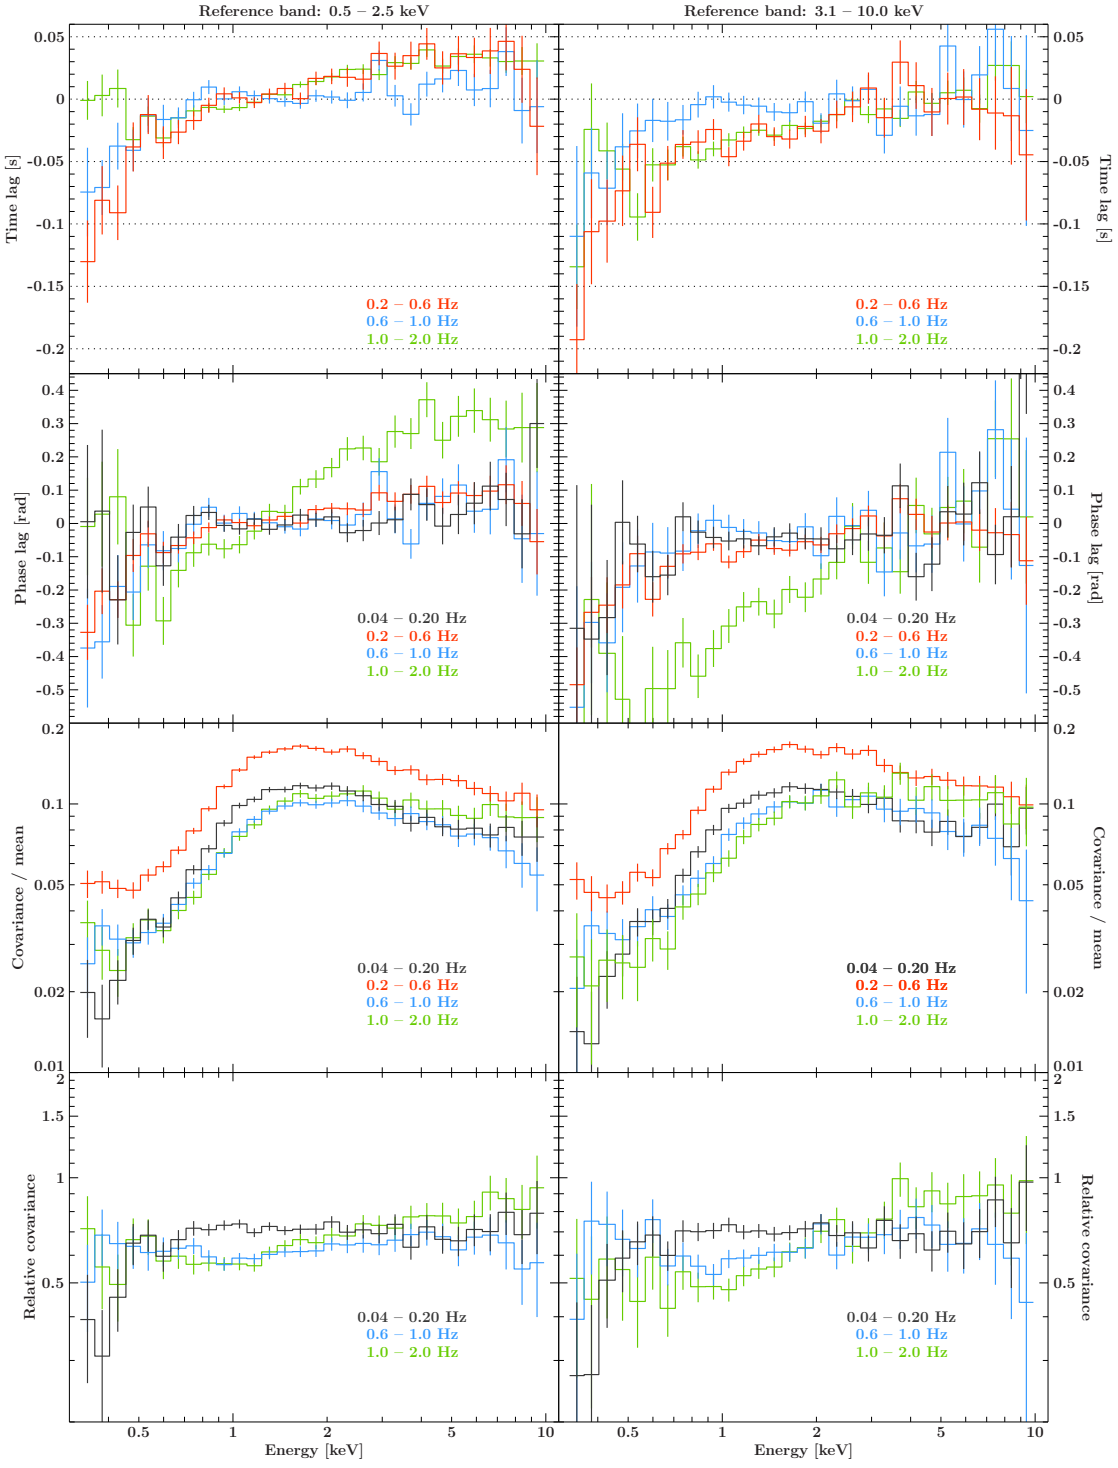

Figure 3.6: Time- and phase-lag spectra, covariance spectrum divided by average spectrum and covariance spectrum divided by covariance spectrum computed in the range 0.2 – 0.6 Hz for 2009 observation, using the soft (0.5 – 2.5 keV, left) and the hard (3.1 – 10.0 keV, right) reference bands. The frequency ranges used are 0.04 – 0.20 Hz (black), 0.2 – 0.6 Hz (red), 0.6 – 1.0 (blue) and 1.0 – 2.0 Hz (green).

this frequency range, the log-linear shape appears also to break downwards at around 2 keV (this behaviour is seen most clearly in the phase lag plot), suggesting that there is a composite shape, perhaps consisting of leading higher-energy disc emission together with lagging lower-energy disc emission. It should be noted that the general behaviour in the lags is replicated using both choices of reference band, albeit with small differences between them.

3.3.3.2 Covariance spectra

In addition to the causal information shown above, we extract covariance spectra from the cross-spectra (Wilkinson and Uttley, 2009; Uttley et al., 2011) in order to quantify, as a function of energy, the amplitude of variable emission that is linearly-correlated with the reference band of choice (see Section 3.2). The covariance spectrum depends on the coherence between the reference band of choice and each energy channel. In the limit where the coherence is unity, the covariance spectrum is equal to the rms spectrum (e.g. Revnivtsev et al. 1999), although with considerably smaller error-bars due to the effect of the broad reference band acting as a matched filter on the flux variations in individual channels.

The panels in the third row of Fig. 3.6 show the corresponding covariance spectra divided by the mean spectrum, which enables us to estimate the contribution of each component to the overall variability, analogous to the fractional rms spectrum. For energies above ~ 2 keV, the fractional covariance spectra show a roughly power-law slope that becomes harder for increasing frequencies. At the lowest frequencies, the power-law slope is softer than the mean, while it reaches the slope of the average spectrum at the highest frequencies, 1.0 – 2.0 Hz. This effect is mild if only the first three frequency ranges are considered, and becomes more obvious for the 1.0 – 2.0 Hz range.

Such an effect can also be seen in the lowest panels of Fig. 3.6, by computing the ratio between the covariance in the frequency bands 0.04 – 0.20 Hz, 0.6 – 1.0 Hz and 1.0 – 2.0 Hz and the covariance in the range 0.2 – 0.6 Hz. In this case, the variation in power-law slope is plotted independently of any constant component, which would not contribute to the covariance spectra used to normalise the data. Thus the effect cannot be due to any simple change in the ratio of a variable to constant spectral component which have different spectral slopes. Instead, the hardening of the power-law variations at higher frequencies is intrinsic to the variable power-law itself. Similar behaviour has been seen in frequency-dependent rms spectra of hard state BHXRBs obtained from *RXTE* data (e.g. Revnivtsev et al. 1999).

At the disc-dominated energies (i.e. below ~ 1 keV), the fractional covariance is up to a factor ~ 5 weaker than at the power-law-dominated energies, suggesting that, although the disc variability is apparently driving the power-law variability (as seen in the causality argument provided with the lags above, and in Uttley et al. 2011), the disc emission is less variable than the power-law emission.

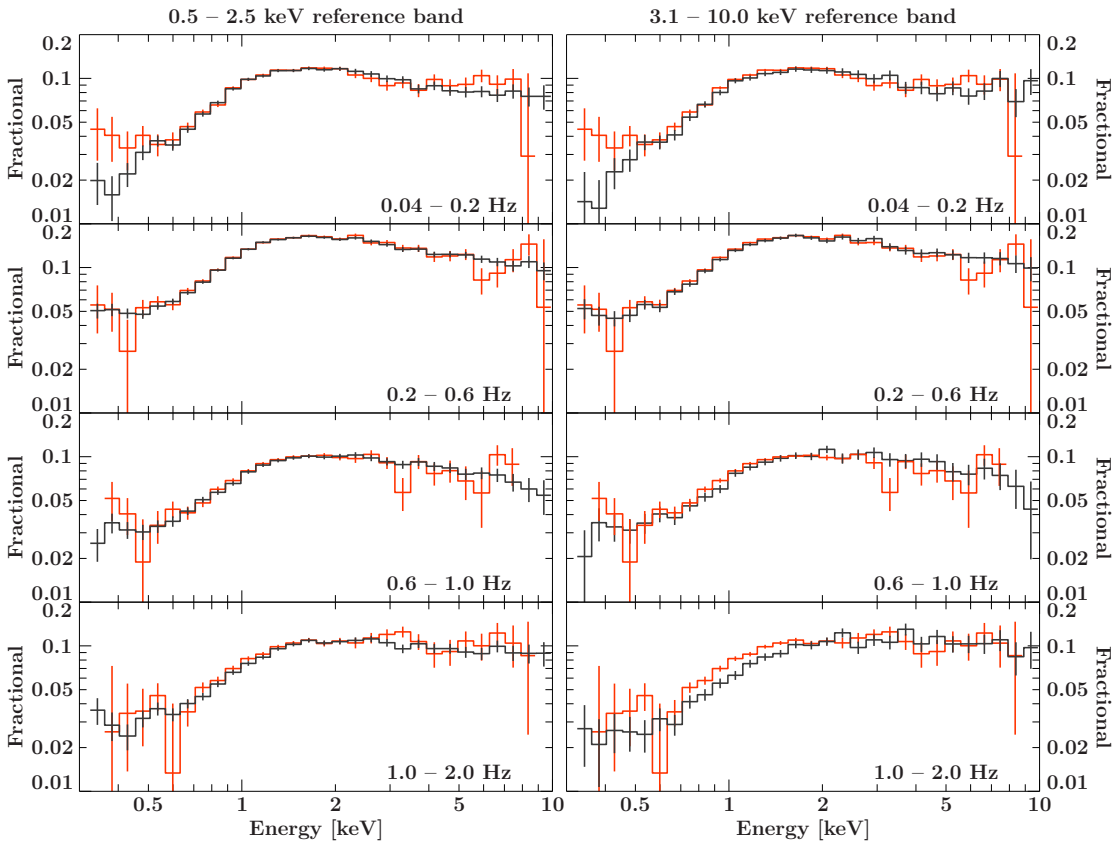

Figure 3.7: Fractional covariance (black) and fractional rms (red) in the 2009 observation for the frequency ranges 0.04 – 0.20 Hz, 0.2 – 0.6 Hz, 0.6 – 1.0 Hz and 1.0 – 2.0 Hz

In Fig. 3.5 we showed that the coherence clearly drops below unity at low frequencies in any of the combinations of energy bands studied where the ultrasoft band was considered. As in the previous cases shown above, this drop can be quantified in an energy-dependent manner by comparing the rms and covariance spectra. The fractional covariance spectra shown in Fig. 3.6 appear to flatten off at the lowest energies in the intermediate frequency bands, while in the lowest-frequency band (where coherence is low) it continues to drop. We can show that this effect is due to the low coherence of variations of the softest photons by comparing the fractional covariance spectra with the fractional rms spectra, shown in Fig. 3.7. Where the fractional covariance (black) and rms (red) are the same, the coherence is consistent with unity. However, at low frequencies, the fractional covariance is clearly lower than the fractional rms below ~ 0.5 keV, indicating sub-unity coherence at these energies, consistent with our findings from the frequency-dependent coherence.

On the other hand, above 0.5 keV the fractional rms and covariance are similar for the three lowest frequency ranges, indicating that the higher-energy part of the disc emission varies coherently with the power-law emission on time-scales longer than a second. However, as also expected from the frequency-dependent coherence measurements, the coherence of this soft-band emission with the higher energy emission also drops in the 1–2 Hz frequency bin. This drop manifests itself as a drop in fractional covariance compared to fractional rms,

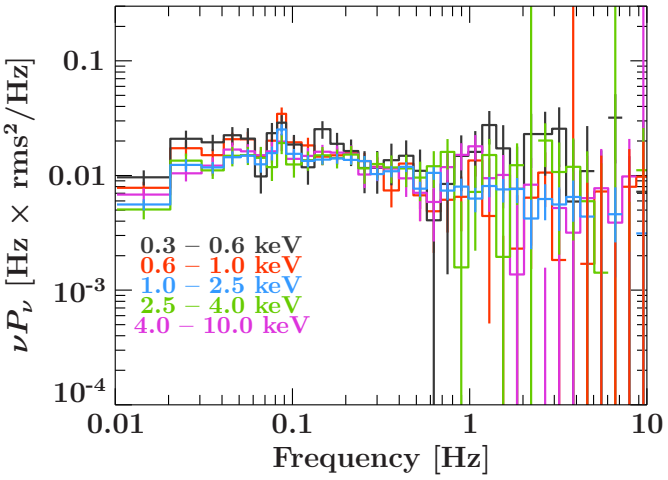

Figure 3.8: Power spectral densities for the 2006 observation

which is seen at low energies for the 3.1–10 keV reference band, and at high energies for the 0.5–2.0 keV reference band (since this band contains the disc emission).

3.3.4 Comparison with the spectral-timing properties in the 2006 hard state

The 2006 observation was obtained during a hard state that was a factor of ~ 3.4 fainter than the 2009 observation in the *Swift/BAT* lightcurve (red arrow in Fig. 3.1). The luminosity difference is even greater at the lower energies covered by *XMM-Newton*, since the 2006 data show a much harder power-law component ($\Gamma = 1.620 \pm 0.005$) which dominates the SED, since the disc component is much weaker (see Figs. 3.2 and 3.3 and Table 3.1 for a comparison). Due to the relatively low signal-to-noise in the 2006 data compared to 2009, we limit our comparison of the data to a few key aspects (such as the PSD and coherence) and the two lower frequency ranges for the energy-dependent comparison of lags and covariance.

The power spectral densities for the 2006 observation are shown in Fig. 3.8. In this observation, the data can be approximated by a power-law down to ~ 0.04 Hz with a similar slope and amplitude for all energy ranges considered, as opposed to the much clearer energy dependence of the power amplitude of the two peaks observed in 2009, whose relative amplitudes also changed with energy. The lower-frequency peak in the PSD seen in 2006 is consistent with the correlation of PSD changes with spectral evolution seen in moderate to high luminosity hard states of other BHXRBs, with lower-frequency features seen when the spectrum is harder (e.g. Belloni et al. 2005).

In the 2009 observation, an important property of the coherence at frequencies below ~ 0.6 Hz is a drop down to coherence values ~ 0.5 for each combination of bands that involved

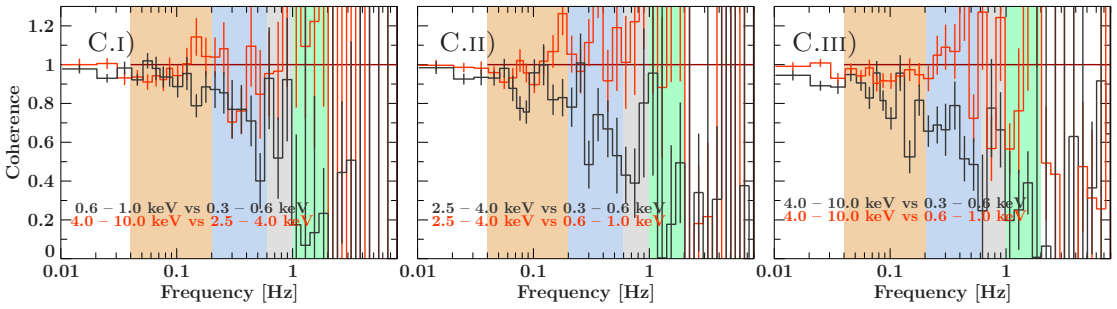

Figure 3.9: Coherence values for the 2006 observation, using the same energy bands as in Fig. 3.5

the ultrasoft band. The corresponding coherence plot for 2006 (Fig. 3.9) shows a value much closer to unity down to the lowest-frequencies probed.

Fig. 3.10 shows the energy-dependent lags and fractional covariance spectra for the 2006 data for the first two frequency-ranges (0.04 – 0.20 and 0.2 – 0.6 Hz), using the reference band 0.5 – 2.5 keV. The harder reference band gives consistent but noisier results (as expected, since the coherence is high at these frequencies). These plots show behaviour consistent with that previously reported for this observation by Wilkinson and Uttley (2009) and Uttley et al. (2011). The lowest-frequency lags show the characteristic drop below 1 keV that can be associated with the disc driving the variability at harder energies, consistent with the behaviour also seen in 2009 (and implying that there is intrinsic disc variability in both observations). However, the covariance spectra show a 'soft excess' at the lowest frequencies, implying that there is *extra* disc variability relative to the power-law variations at these frequencies. This behaviour contrasts with that seen in 2009 where the disc variations at all frequencies appear to be suppressed when compared with the power-law variations.

3.4 Discussion

To summarise, our main observational findings are:

1. The 2009 *XMM-Newton* observation of SWIFT J1753.5–0127 took place when the source was in a significantly more luminous hard state (~ 6 times more luminous in the 0.5–10.0 keV band) than in 2006. Correspondingly, the power-law spectral component is softer than in 2006, and importantly the soft excess associated with the disc black-body is significantly stronger. The inferred black-body temperature in 2009 is only marginally higher than in 2006, but the 2009 disc black-body normalisation is a factor ~ 30 larger than in 2006.

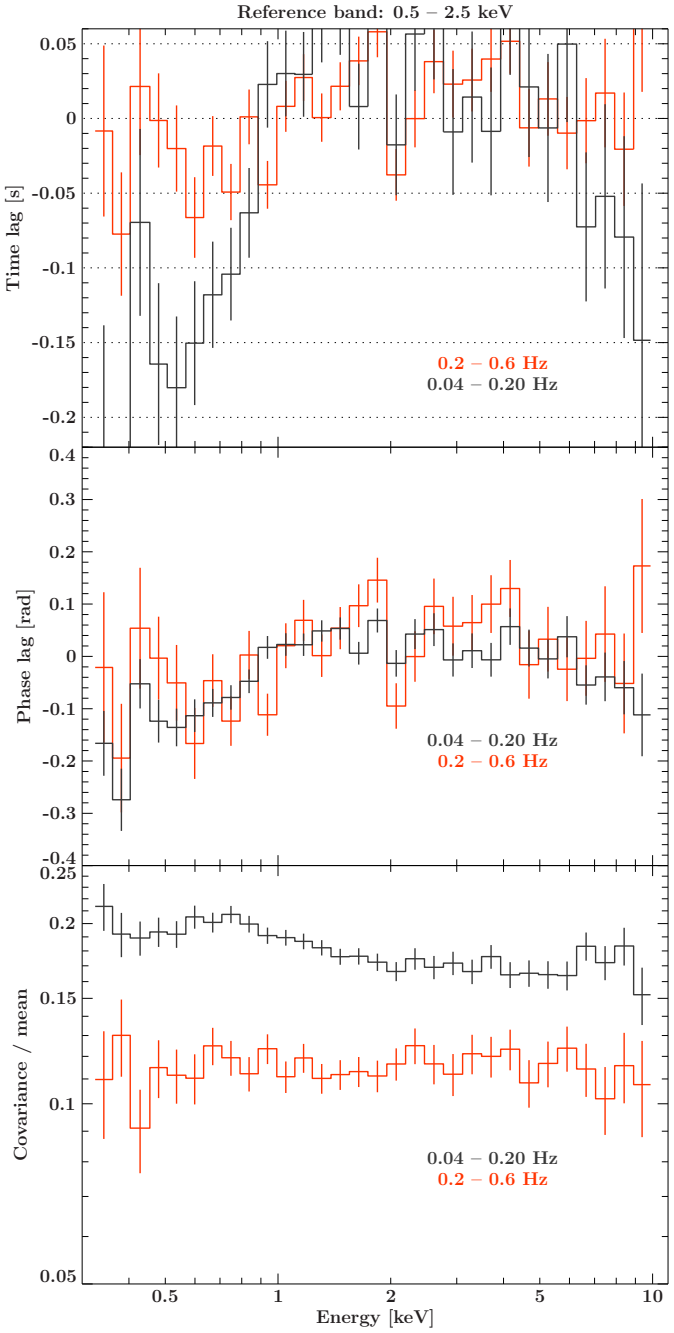

Figure 3.10: Time- and phase-lag spectra and covariance spectrum divided by average spectrum for 2006 observation. The frequency ranges used are 0.04 – 0.20 Hz (black) and 0.2 – 0.6 Hz (red).

2. The PSDs in 2009 are band-limited, and can be interpreted in terms of two broad frequency components which peak around 0.4 Hz and 2 Hz respectively.

3. The 2009 PSD shape and amplitude evolves with energy. The lowest-amplitude variability over all frequencies is seen in the softest 0.3–0.6 keV band, followed by the 0.6–1.0 keV band. The amplitude of the lowest-frequency component is largest in the 1.0–2.5 keV band before dropping slightly at higher energies, while the amplitude of the high frequency component increases with energy.

4. The covariance and rms spectra suggest that the energy-dependent PSD behaviour is linked to two effects. Firstly, there is a fall-off in fractional rms (seen in all frequency ranges) at low-energies where the disc black-body dominates the emission. Secondly, the variable power-law emission component becomes harder at higher temporal frequencies. Neither of these effects are seen in the 2006 data. In fact, in 2006 the opposite is seen at low energies, where the fractional rms increases in the energy range where the disc contributes to the spectrum.

5. Frequency-dependent lag measurements show that in general, the harder photons lag softer ones. Within the energy bands dominated by the power-law (2.5–4 keV and 4–10 keV) and the disc (0.3–0.6 keV and 0.6–1 keV) the lags are relatively short and the phase-lags evolve only smoothly and weakly (if at all) with frequency. However, the lags between power-law and disc bands are larger and show significant structure in the form of 'steps', with a significant increase in phase lag with frequency.

6. Lag-energy spectra show similar behaviour to that reported by us (Uttley et al., 2011) in other hard state observations of BHXRBs (including SWIFT J1753.5–0127 in 2006). The sharpest change in the lags is seen at low energies where the disc dominates the spectrum. The low-energy shape of the lag spectrum depends on the frequency probed: the highest frequencies show evidence for a shift in the break in the lag spectrum to higher energies, together with an upturn in the lowest energy bins (similar to that seen in GX 339-4, Uttley et al. 2011).

7. The coherence between the bands dominated by the power-law remains high over a broad frequency range. However, between the power-law and disc-dominated bands, and between the two disc-dominated bands, coherence drops significantly at frequencies below 0.5 Hz (where the softest disc band shows the least coherence with the other bands) as well as above 1 Hz (where both disc bands show low coherence with respect to the power-law bands). Comparison of the covariance and rms spectra confirm that low-frequency variations of the softest photons, below 0.5 keV, are only weakly correlated with variations at higher energies. In contrast, the 2006 data are relatively coherent at low frequencies.

The spectral-timing behaviour of SWIFT J1753.5–0127 in 2009 is complex but we will attempt to interpret it here in terms of the physical picture of mass-accretion rate fluctuations arising in and propagating through the accretion disc, suggested by Wilkinson and Uttley (2009) and Uttley et al. (2011) in order to explain the covariance and time-lag behaviour of hard state BHXRBs. Since we see similar lag behaviour in the SWIFT J1753.5–0127 2009 data, it is useful to understand whether this model can be extended to explore the more complex spectral-timing behaviour seen in this dataset. In doing so, we will focus primarily on explaining the differences between the spectral-timing properties of the 2009 data we present here and those from 2006 and observations of other BHXRBs reported in Wilkinson and Uttley (2009) and Uttley et al. (2011). Comprehensive testing of the propagating fluctuations model against these data is beyond the scope of this work, so here we simply sketch how the key observational results might be explained by the model.

The key feature of the propagating fluctuations model used to explain the soft X-ray behaviour of other hard state BHXRBs is that the disc plays an important role in generating and carrying the fluctuations (at least on fluctuation time-scales > 1 s). The 2009 data show much stronger disc emission than is seen in 2006. This difference is puzzling because the inferred disc temperatures are relatively similar, while the normalisations are very different. Taken at face value, the disc black-body normalisation in 2009 implies an inner disc radius 5–6 times larger than that in 2006, which is at odds with the standard interpretation of the other phenomenological changes between the two data sets, which assumes that the disc becomes less truncated as PSD characteristic frequencies increase and the disc emission becomes stronger. However, to interpret changes in disc black-body normalisation in terms of changes in inner disc radius, we must assume in that both data sets the disc emissivity profile is the same. Since the X-ray luminosity in 2006 is dominated by the power-law emission, it is likely that much of the disc emission is produced by X-ray heating of the disc by the power-law, which could produce a much more centrally peaked (and apparently 'smaller') disc black-body component (e.g. see Gierliński et al. 2008). In 2009, the disc emission is comparable to or even exceeds that from the power-law and most of the emission is likely to be intrinsic to the disc, leading to a more 'standard' emissivity profile. Since the peak energy of the disc black-body emission is comparable to the minimum energy included in the fit (0.7 keV), distinguishing among different emissivity profiles from the spectral fits would be an arduous task. Systematic errors deriving from spectral calibration do not, however, impact on the variability analysis performed here.

Assuming that the disc emissivity profile in 2009 has the standard R^{-3} form (in terms of bolometric emission) and hence the disc temperature scales with radius as $R^{-3/4}$, we can consider the dependence of cumulative disc emission on radius. For comparison and allowing for the likely possibility of substantial disc-heating by the power-law in 2006, we fitted the 2006 spectrum with a combined single black-body (representing a hot inner ring heated by the central power-law emission) and a standard disc black-body component. We forced the disc black-body to have the same normalisation (i.e. the same assumed inner radius) as the

disc in 2009 (note that due to degeneracy in the model the inner radius in 2006 may be even larger, as expected if the disc is more truncated). For this composite model the disc black-body temperature drops to 0.11 keV while the inner single black-body temperature is high at 0.31 keV, but corresponds to only a small emitting area at the inner edge of the disc. The best-fitting model parameters give χ^2/d.o.f.= 565/623, providing an even better fit to the data than the single disc black-body model. The resulting cumulative emission profiles are shown in Fig. 3.11 for the soft and ultrasoft energy bands (after allowing for the effects of the instrumental response on each band). Due to the lower disc temperature and the relatively hot inner ring, the disc emission in 2006 is much more centrally concentrated than in 2009. Furthermore, the ultrasoft emission is significantly more extended across the disc than the soft emission. We can use these differences to explain some of the interesting spectral-timing behaviour in the data. Note that the arguments that follow apply independently of the actual inner disc radius $r_{\rm in}$, since the emission profile in each band is set by the disc temperature at $r_{\rm in}$ (which is set by the spectral fit results) and the *relative* change in radius.

First, let us consider the 2009 data and suppose that a fluctuation in the disc arises at a relative small radius $R \sim 1.5\ r_{\rm in}$, and modulates the disc emission as it propagates inwards. Since only ∼40 per cent of the ultrasoft flux is contained within this radius, the observed variability amplitude will be suppressed by a factor ∼ 2.5 compared to the amplitude of the intrinsic fluctuation. In contrast, the soft band fluctuation will be suppressed by only a factor ∼ 1.5, while the power-law emission, if it originates within the disc truncation radius, will not be suppressed at all. This simple picture can explain the energy dependence of the PSD amplitude in 2009 and the corresponding drop in rms and covariance spectra at low energies. In this scenario, the same behaviour may not seen in other hard-state *XMM-Newton* observations of BHXRBs or in the 2006 data because of two effects. Firstly, since the PSDs show lower characteristic frequencies in these other observations, the fluctuations in accretion rate may arise at a larger radius, so that a significantly larger fraction of the disc emission can be modulated. Secondly, as can be seen in the 2006 emission profile, the combination of X-ray heating of an inner ring with a cooler disc at larger radii will lead to more concentrated emission in both disc bands, so that all or most of the flux will be modulated by a fluctuation arising at $R \sim 1.5\ r_{\rm in}$ or larger. Thus we do not see suppression of the PSD amplitude at soft X-ray energies.

We next consider the drop in coherence at low frequencies. The natural explanation for this behaviour is that there is some component of accretion fluctuations in the disc which are unable to propagate very far inwards and therefore appear as variations in the extended disc emission (in the ultrasoft band and to a lesser extent the soft band) but do not produce correlated variations in the power-law emission. One possibility is that these variations are generated at larger radii than the spectrally-coherent fluctuations, so that they are viscously damped before they are able to propagate to the innermost part of the disc. Again, we may see this effect in the 2009 data because of the strong intrinsic disc emission and extended emissivity profile. Disc emission variations driven by power-law heating of the disc will be much

more coherent with the power-law variations. Furthermore, the lower-frequencies seen in the PSD of the 2006 data could imply that signals can effectively propagate from larger radii in the disc and thus modulate a significant fraction of the intrinsic disc emission coherently.

The higher-frequency drop in coherence corresponds roughly to the overlap between the low and high-frequency PSD components, and could be produced by a similar effect to that described above. For example, if the higher-frequency signal corresponds to a second, even smaller radius in the disc which generates propagating accretion fluctuations, the signals generated at the same frequencies at larger radii may not propagate effectively through the disc, again producing fluctuations that are not seen in the power-law component. It is important to note however that the relatively small radii which we suggest as the origin of fluctuations (e.g. within 1.5 $r_{\rm in}$) also imply that the power-law emitting region must be compact. Otherwise, seed photons from the varying regions in the disc would upscatter to produce variations of the power-law that are coherent with the disc variations. Similarly, we expect some component of disc emission to be produced by X-ray heating by the power-law and hence be coherent with the power-law variations. These components are likely to exist at some level in the variable X-ray signal, but they must be relatively weak, at least on the timescales where the coherence is low.

Some level of interaction between the disc photon variations and the central power-law emitting region could potentially explain the observed hardening of the covariance and rms spectra at higher frequencies (and correspondingly, the peak in low-frequency power in the 1–2.5 keV energy range). The amplitude of soft-photon variations is largest in the low-frequency peak of the PSD. Compton cooling of the power-law emitting region by these soft photons could lead to a significant pivoting of the power-law in response to these variations, such that at low frequencies the fractional amplitude of power-law variations increases towards lower energies. This effect then has the converse implication that at high frequencies the variable power-law appears to harden.

Finally, we consider the lag behaviour. The relatively small lags and smooth dependence on frequency seen between the two power-law bands and also the two disc bands can be simply explained because here we are seeing the same spectral components. Fluctuations propagating through the disc will primarily modulate the ultrasoft emission first, but there is significant overlap between soft and ultrasoft emission across the disc, so we expect lags to be relatively small. The same is true for the power-law component – lags should be small between energies where the power-law component dominates. The largest lags are then expected between the disc and the power-law components, assuming that these are physically separated, i.e. the power-law emission primarily originates within the disc truncation radius.

We must still explain the clear structure in the lag vs. frequency dependence of the power-law relative to the disc-dominated bands, i.e. the apparent 'steps'. Similar steps can be seen in other hard state data from *RXTE*, but between power-law dominated energies (Nowak et al., 1999). In the SWIFT J1753.5–0127 2009 data, the effect is particularly strong: the time lags

actually increase with frequency above ~ 0.5 Hz before dropping above ~ 1.5 Hz. Nowak (2000) and Kotov et al. (2001) have suggested that steps in lag vs. frequency could be related to the transition between two variability components, since the steps seem to occur in between the 'bumps' or broad Lorentzians in the PSD. If each component shows a different lag, one would expect a stepped lag vs. frequency dependence consisting of discrete lags with relatively weak lag vs. frequency dependence where a single component dominates the variability, but with sharp steps occuring in the transition region between components. This picture may fit with the SWIFT J1753.5–0127 2009 data, but the signal-to-noise is such that it is difficult to find a precise match between the step in the lag and the transition between the low and high-frequency components in the PSD (which seems to occur around 1 Hz). It is important to note however that the rise in the time lag seems to coincide with the frequency range where the coherence peaks around unity. Incoherent signals which are caused by a mixture of uncorrelated variable components will show lags which are a mixture of the lags of each separate component. Thus, the rise in lag could be explained if the component causing the low-frequency drop in coherence contributes a small lag, reducing the total observed lag, but at higher frequencies of ~ 0.5 Hz this component disappears from the variable signal and the lag suddenly increases as a result. One possibility is that the incoherent fluctuations in the disc also modulate a small disc-corona component which is physically separate (and also spectrally distinct from) the inner power-law emitting region, which may be linked to a hot inner flow or the base of a jet. The interband lags of such a component would then be relatively small, and would give way to the larger lags (due to propagation effects) once the coherent signal becomes dominant. If the lags of the coherent signal themselves show a dependence on frequency, (e.g. a weak increase in phase lag with frequency as observed in other hard state BHXRBs) a complex pattern of rises and falls could be produced.

3.5 Conclusions

The soft X-ray coverage of the EPIC-pn camera onboard *XMM-Newton* has allowed us to carry out a comprehensive study of the spectral-timing properties of the BHXRB SWIFT J1753.5–0127 in the luminous 2009 hard state, and to perform a comparison with the fainter 2006 hard state.

In our previous study of a sample of hard-state BHXRBs (Uttley et al., 2011), we showed that (i) soft (disc) X-ray variations lead hard (power-law) variations up to frequencies of about 1 Hz, and (ii) at the same low frequencies, the fractional variability in the disc is larger than the variability observed in the power-law component. Mass-accretion rate fluctuations that propagate through the disc would produce variability on frequencies corresponding to the viscous time-scale where they are generated, and eventually reach a centrally-concentrated

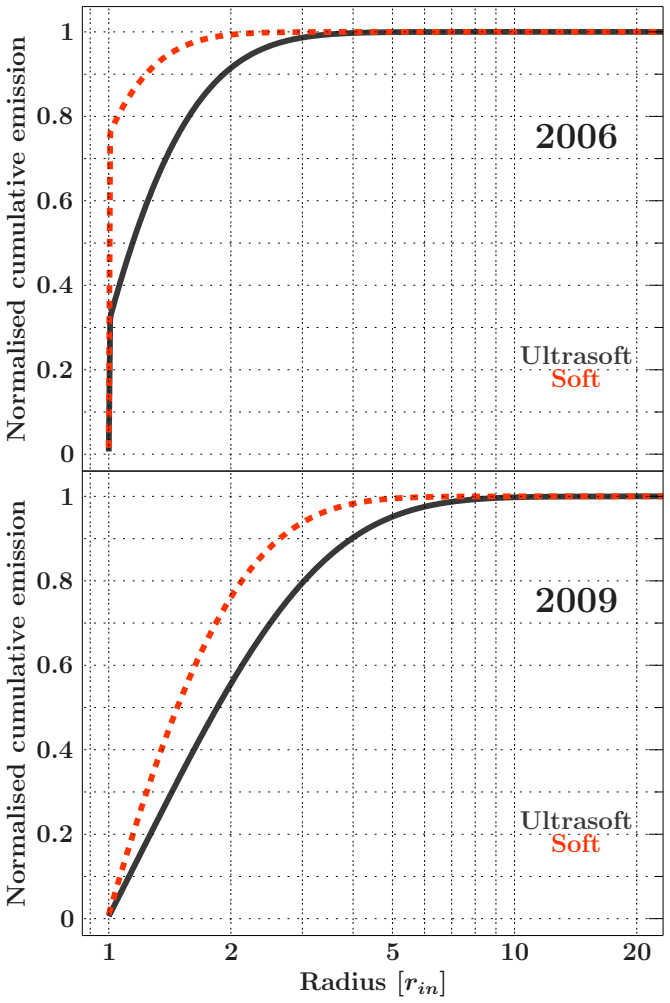

Figure 3.11: Normalised cumulative emission function in the ultrasoft (solid black) and soft bands (dashed red) as a function of disc radius. The top panel corresponds to the 2006 data, assuming a composite disc black-body and inner hotter black-body associated with X-ray heating by the power-law emission (see text for details). The bottom panel corresponds to the 2009 data, assuming a dominant intrinsic disc black-body profile.

source of power-law photons. This basic picture would both explain the hard-to-soft lags and give a reasonable explanation for the excess fractional variability seen in the disc as compared to the power-law.

The 2009 observation shows a breadth of previously unknown observational features, which we summarise as follows.

Firstly, the variability amplitudes in the PSD are strongly dependent on energy, with the harder band varying more than the softer bands. The explanation for this effect is seen in the fractional covariance spectrum, where the bright disc black-body component that is seen in the spectrum now appears suppressed in fractional-variability terms with respect to the power-law, at all frequencies considered. This is opposite to what found in less-luminous hard states by Uttley et al. (2011), where the disc was more variable than the power-law component. The higher luminosity (and thus, accretion rate) seen in this observation could imply

that the disc emissivity in the observed bands is less centrally peaked than in less-luminous hard states. In this case, a fluctuation that is produced at an inner disc radius (e.g. within 1.5 $r_{\rm in}$) will not modulate the emission from outer radii and this will effectively reduce the *observed* fractional variability of the disc as a whole. This is consistent with the shape of the energy-dependent lags.

Secondly, we have found that the low frequency ($\lesssim 0.5$ Hz) disc variability below < 0.6 keV appears partly uncorrelated with the rest of the emission produced by both the disc and power-law components. An explanation for this could exist if this incoherent accretion variability is produced at larger radii than the coherent variability, affecting the softest emission that is more extended while being unable to propagate to the innermost part of the disc and the power-law emitting region due to viscous damping.

Finally, the last feature to highlight is an up-turn in the frequency-resolved time lags above ~ 0.5 Hz. This could be produced if, e.g. a relatively central secondary extended corona that sandwiches the disc – and thus reduces the observed hard-to-soft lags caused by disc propagation variations – becomes swamped by a coherent source of variations associated with the high-frequency peak in the PSD (perhaps an even smaller unstable radius in the disc).

By identifying a number of spectral-timing signatures with the accretion disc, this work suggests that a wealth of information regarding the properties of accretion flows can be discovered by extending spectral-timing measurements to softer X-rays and using techniques that allow better spectral-resolution. By combining these powerful techniques with more detailed modelling, a comprehensive understanding of the variability may help to unravel many of the mysteries underlying accretion physics and associated phenomena such as the nature of state transitions and the production of the jets seen in the hard states of accreting black holes.

CHAPTER 4

Energy- and Frequency-Resolved Characterisation of Correlated Variability: The 2D Cross-Spectral Fitting Method

4.1 Introduction

The future understanding of the extreme physics that govern black hole X-ray binaries depends on both the development of more sensitive instruments and the development of novel techniques that are adequate for the capabilities of such detectors. Spectral-timing studies are at the forefront of scientific analysis of these sources. These studies require three important characteristics in detector capabilities, namely a large collecting area, high time resolution and CCD-quality energy resolution or better. A combination of these factors will yield high signal-to-noise observations where photons can be time-tagged accurately and their energy reconstructed to within a per cent or two.

The Timing mode for the EPIC-pn instrument onboard *XMM-Newton* is the only current instrument that provides all three characteristics, setting an example for upcoming detectors devoted to this kind of analysis, such as the Large Area Detector (LAD) onboard the proposed mission *LOFT*. Examples of the scientific work that can be done with the capabilities of EPIC-pn are shown in Chapters 2 and 3.

Standard spectral-timing methods include the study of power spectral densities, coherence, phase- and time-lags, rms (Revnivtsev et al., 1999) and covariance spectra (Wilkinson and

Uttley, 2009). These products are the result of collapsing Fourier-frequency- and energy-dependent variability information from multi-channel time-series onto either the energy or the Fourier-frequency plane. These products are either used for displaying results or for subsequent fitting, usually using Lorentzians in the case of power spectra, and spectral models for rms and covariance spectra. In the case of rms and covariance spectra (thus, in the energy domain), fitting will provide an estimate of the parameters that describe the variable part of the spectral components. Covariance and rms spectra will be equal for unit coherence, with covariance spectra having much smaller error-bars thanks to the reference band used to increase the signal-to-noise (Wilkinson and Uttley, 2009; Wilkinson, 2011).

The extension of the analysis of energy-dependent products to the Fourier-frequency domain is fundamental for our understanding of the physics around black holes. Where spectral fitting leads to model degeneracies and calibration is inaccurate, timing information can provide causal and correlation information, helping to physically locate the sources of X-ray emission as well as constrain the causality relations between them. On the other hand, timing information alone gives very little information about energetics, and has to be treated carefully. We aim to provide a self-consistent method for using energy and time information simultaneously.

Multi-channel phase information, also known as the lag spectrum, is of very difficult treatment in the context of 'traditional' fitting in the X-ray domain. Because phase models are in units of radians and not photons/s/cm^2, there is no straightforward way to fit energy-dependent phase data while considering the instrumental response that redistributes incoming photons among different energies and privileges the detection of photons where the effective area and quantum efficiency are larger. The instrumental response, therefore, needs to be taken into account to weight phase lags according to photon flux while following the properties of the redistribution matrix.

It is possible to imagine a situation where an emission line appears superimposed on a continuum in the spectrum, but its emission is delayed with respect to the continuum by a certain phase delay. Such an example can be seen in Fig. 4.1 for different line centroid energies. This figure shows that the instrumental response is, also in the case of phase lags, responsible for the appearance of wiggles and wings that appear around the centroid energy of the lines. These effects need to be considered when the analysis requires a high degree of accuracy, e.g. when studying the effects of reflection on the lags to carry out disc reverberation mapping studies, possibly to understand the effect of General Relativity on the regions around black holes. A second improvement to the current analysis techniques would be to fit the energy and Fourier-frequency (time-scale) dependence of variability simultaneously, overcoming the deficiency of covariance and rms spectra when it comes to fitting these products in different frequency ranges self-consistently.

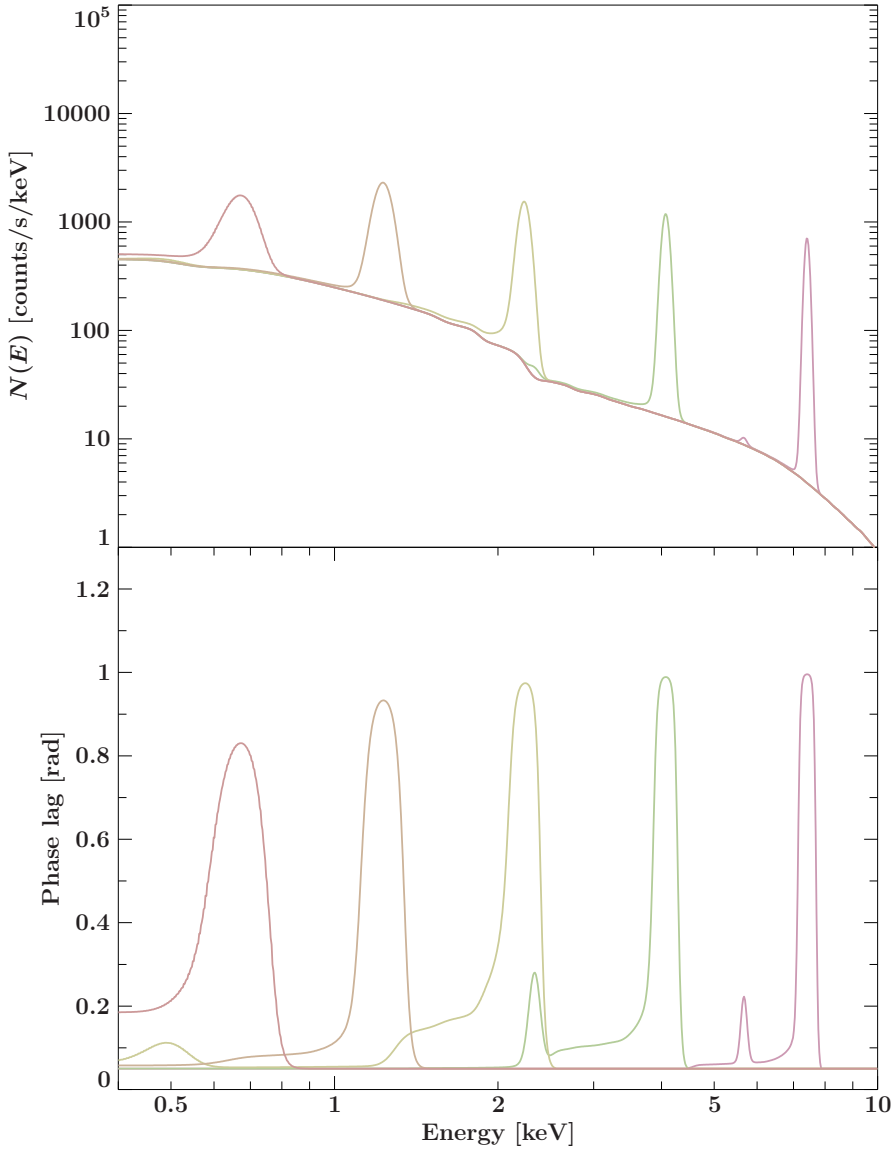

Figure 4.1: Example of the effect of the instrumental response of EPIC-pn onboard *XMM-Newton* on the shape of the lags. Top panel: Spectrum consisting of a power-law with Γ = 1.8 and a choice of five variable Gaussian lines centred at 0.67 keV, 1.23 keV, 2.24 keV, 4.07 keV and 7.4 keV, all of which have a full-width at half maximum set to 0.5 eV. The Gaussians are delayed by 1 radian with respect to the power-law emission. Bottom panel: resulting energy-dependence of the lags, showing excess phase lags corresponding to the wings of the emission lines produced by the instrumental response.

The technique proposed here overcomes the limitations mentioned above and is based on the extension of spectral fitting to correlated variability in two dimensions, energy and Fourier-frequency, simultaneously. This work expands the work conducted by Vaughan et al. (1994) to take into account variability information at each Fourier-frequency separately, as well as the instrumental response, the effect of which is important as shown earlier. The spectral-timing product to be fitted is the cross-spectrum, as we will see.

For this method to work, new 2D models can be constructed in a two-step process. Firstly, a time- and energy-dependent response function[1], $\psi(t)$, can be built from a desired process that produces variability (e.g. reflection, Comptonisation, jets...). Secondly, a Fourier transform of the response function is performed to produce a transfer function ($\hat{\psi}(\nu)$) that maps variability at each energy onto the Fourier-frequency space. This quantity is used to produce a cross-spectrum (Vaughan and Nowak, 1997), which is used in Chapter 2 to produce the frequency-dependent lag model.

In the time domain, one could imagine that a physical process $x(t)$ is a source of variability in a black hole X-ray binary, for example if it produces a mass-accretion rate fluctuation at a given radius in the accretion disc (e.g. Arévalo and Uttley 2006). One such process is 'latent', in the sense that its existence cannot be observed directly, but it *causes an effect* on a physical mechanism (characterised by its response function, $\psi(t)$) that is responsible for a particular spectral shape (e.g. power-law variability or disc reflection). The latter physical mechanism *alters* the variability properties of the incoming signal, and results in a third time series, $y(t)$, resulting from a convolution of $x(t)$ through $\psi(t)$. In the case of one single energy range, this corresponds to the one input – one output paradigm in signal theory (see e.g. Bendat and Piersol 2010). A real life equivalent of such paradigm can be found in the study of power spectral densities (PSDs) in one energy band, where the input cannot be distinguished from the output. In this case, modelling of PSDs cannot, by their own nature, constrain any of the spectral properties of the source.

As an example, we now assume that the signal $y(t)$ can be observed at two different energies, A and B. In the approximation of fully correlated variability, the time series $y_A(t)$ and $y_B(t)$ result from the convolution of one single input signal $x(t)$ through the response functions $\psi_A(t)$ and $\psi_B(t)$:

$$\begin{aligned} y_A(t) &= \int x(t')\psi_A(t-t')\mathrm{d}t' \\ y_B(t) &= \int x(t')\psi_B(t-t')\mathrm{d}t' \end{aligned} \tag{4.1}$$

[1] Not to be confused with the instrumental response.

In Fourier-frequency space, a convolution translates to a multiplication, hence the above reduces to:[2]

$$\begin{aligned} \hat{y}_A(\nu) &= \hat{x}(\nu)\hat{\psi}_A(\nu) \\ \hat{y}_B(\nu) &= \hat{x}(\nu)\hat{\psi}_B(\nu) \end{aligned} \tag{4.2}$$

where $\hat{x}(\nu)$ is the Fourier transform of $x(t)$ and $\hat{\psi}_A(\nu)$ and $\hat{\psi}_B(\nu)$ are the Fourier transforms of $\psi_A(t)$ and $\psi_B(t)$, respectively. The method that we will consider here here can be understood as an extension of the standard X-ray fitting method to the amplitudes of variability of the real and imaginary components of the observed signal in the Fourier domain. The energy- and frequency-dependent variability model is initially produced from a transfer function $\hat{\psi}(\nu)$ based on a set of parameters. For any given frequency, the complex-valued model is convolved through the response, and this is then used to form a cross-spectrum that is then used for a standard χ^2 minimisation from which the best-fitting values are obtained.

Via the minimisation technique, not only will the resulting best-fitting model try to approach the amplitudes of correlated variability, but it will also show the proper phase dependence on energy and Fourier-frequency. This is of fundamental importance for constraining causality in 'stationary'[3] systems such as BHXRBs, that cannot be done with standard spectral fitting.

4.2 Method

4.2.1 Source spectrum to detector counts

Given a source spectrum $S(E,t)$ (units: [photons/s/keV/cm^2]) and a constant detector response $R(E,h)$ during the integration time, the observed lightcurve at a time t_k in detector energy channel $h \in H$ equals:

$$x_h(t_k) = \int \mathrm{d}E\, R(E,h)\, A(E) S(E,t_k) + B(h,t_k) \quad , \tag{4.3}$$

where $A(E)$ is the detector effective area at energy E and $B(h,t_k)$ accounts for the contribution of the background counts.

[2] For N points in the time domain, the total computation time for convolution is $\propto N^2$. By performing a Discrete Fourier Transform $\sim O(N\log_2 N)$, although it has recently been reduced to $\sim O(\frac{34}{9}N\log_2 N)$, see e.g. Johnson and Frigo 2007), and a subsquent multiplication ($\sim O(N)$), the complexity decreases to $\sim N\log_2 N$.

[3] Stationarity over a time-scale similar to the length of an observation is a common approximation in spectral-timing analysis

The Fourier transform $\tilde{X}_h(\nu_l)$ is defined as

$$\tilde{X}_h(\nu_l) = \sum_{k=0}^{N-1} x_h(t_k) e^{2\pi t_k \nu_l i} \quad , \tag{4.4}$$

where $x_h(t_k)$ represents the detector counts in the energy channel h per time bin δt (units: [counts]). See Chapter 1 (pg. 11) for more details. Given that an integral in energy and the Fourier transform are interchangeable, the quantity in Eqn. 4.4 equals

$$\tilde{X}_h(\nu_l) = \int \mathrm{d}E\, R(h,E) A(E) \tilde{S}(E,\nu_l) + B(h,\nu_l) \quad , \tag{4.5}$$

where $\tilde{S}(E,\nu_l)$ is the Fourier transform of $S(E,t_k)$. In the following we will assume $B(h,\nu_l) = 0$.[4]

4.2.2 Cross-spectrum

From the quantity in Eqn. 4.5, the cross-spectrum (XSP) between channels h and h' can be defined as

$$\tilde{G}_{h'h}(\nu_l) = \tilde{X}^*_{h'}(\nu_l) \tilde{X}_h(\nu_l) \quad . \tag{4.6}$$

(see e.g. Vaughan and Nowak 1997), where $(\cdot)^*$ denotes complex conjugation ($Ae^{i\phi} \xrightarrow{*} Ae^{-i\phi}$). The cross-spectrum in Eqn. 4.6, if calculated for each pair of channels, yields a low signal-to-noise. A much better signal-to-noise ratio can be obtained by calculating the cross-spectrum between a channel h and its complement, $H^c_h = \{h' \in H | h' \neq h\}$. The subset H^c_h contains all the detector channels excluding the channel of interest h. Each channel will therefore have its own complement, also called 'reference band'. Thanks to the linearity of Fourier transforms and additivity of counts, the Fourier transform for the reference band of channel h can be constructed like so: $\tilde{X}^c_h(\nu_l) = \sum_{h' \neq h} \tilde{X}_{h'}(\nu'_l)$. The cross-spectrum of channel h with respect to its complement H^c_h (its reference band) is then:

$$\tilde{G}_h(\nu_l) = \tilde{X}^{c*}_h(\nu_l) \tilde{X}_h(\nu_l) = \left(\sum_{h' \neq h} \tilde{X}_{h'}(\nu_l) \right)^* \tilde{X}_h(\nu_l) = \sum_{h'} \tilde{X}^*_{h'}(\nu_l) \tilde{X}_h(\nu_l) - |\tilde{X}_h(\nu_l)|^2 \quad . \tag{4.7}$$

[4] Background counts, if not contaminated or caused indirectly by the source, are uncorrelated with the source counts, although they can produce unwanted energy-dependent correlated variability, which is not addressed here.

Aside from having a larger signal-to-noise, the cross-spectrum calculated as above has the advantage that it retains all the variability information contained in the reference band while avoiding the correlated Poisson noise that would appear if the reference band included the hth channel.[5] The amplitude of the cross-spectrum gives the amplitude of correlated variability of the model at a given frequency in a given energy channel with respect to its reference band. In the case where the variability process is fully coherent in each energy channel and for sufficiently narrow energy channels, the resulting cross-spectrum would be independent of the reference band chosen.

For real data, cross-spectra need to be averaged over a number of observation segments (e.g. following the Bartlett method as in Subsection 1.2.2) for the effect of Poisson-induced noise to vanish.

The relative amplitudes of the real and imaginary components of the averaged cross-spectrum give the phase-lag (Chapter 1):

$$\varphi_h(\nu_l) = \arctan \frac{\Im(\langle \tilde{G}_h(\nu_l) \rangle)}{\Re(\langle \tilde{G}_h(\nu_l) \rangle)} \quad . \tag{4.8}$$

The time-lag then equals:

$$\tau_h(\nu_l) = \frac{\varphi_h(\nu_l)}{2\pi\nu_l} \quad . \tag{4.9}$$

The definition of a spectral-timing model via a transfer function can thus be used to constrain both correlated variability amplitudes and phase-/time-lags via the transformations in Eqn. 4.8 and Eqn. 4.9.

4.2.3 Model optimisation

In general and for Normally-distributed data, the model parameters that best describe a particular set of experimental values can be obtained by minimising a χ^2 statistic:

$$\chi^2 = \sum_k (d_k - m_k)^2 w_k^2 \quad , \tag{4.10}$$

where d_k represents the data, m_k is the model, and w_k is the weight of the kth datum. Generaly, the weighting is $w_k^2 = \frac{1}{\sigma_k^2}$ where σ_k is the error on the kth datum d_k.

[5] Poisson noise between independent channels only manifests itself as a scatter of cross-spectral values that tend to vanish for a large number of segments, and as an obvious increase of the variance, but does not affect the cross-spectrum, on average.

The observed, thus segment-averaged, cross-spectrum defined in Eqn. 4.7 is redefined here as $\tilde{D}_h(\nu_l)$. A rebinning in frequency, e.g. logarithmic, will increase the signal-to-noise (Nowak et al., 1999).

The observed cross-spectrum $\tilde{D}_h(\nu_l)$ is a complex-valued matrix of L frequencies times H detector channels. In practical terms, the same reasoning used in the orthogonality of the sine and cosine functions in the definition of the Fourier transform (and in its discrete equivalent shown in Eqn. 1.19) ensures that the real and imaginary components of the cross-spectrum be independently-distributed. Thus, $\tilde{D}_h(\nu_l) = \tilde{D}_{R,h}(\nu_l) + i\tilde{D}_{I,h}(\nu_l)$. The χ^2 with $2 \times L \times H$ degrees of freedom in Eqn. 4.10 then equals:

$$\chi^2 = \chi^2_R + \chi^2_I \quad . \tag{4.11}$$

The weights for the real and imaginary parts of the cross-spectrum can be obtained independently by setting the weights to be equal to the inverse of the variance of the real and imaginary parts, $\sigma^2_{R,hl} = \langle \tilde{D}^2_{R,h}(\nu_l)\rangle - \langle \tilde{D}_{R,h}(\nu_l)\rangle^2$ and $\sigma^2_{I,hl} = \langle \tilde{D}^2_{I,h}(\nu_l)\rangle - \langle \tilde{D}_{I,h}(\nu_l)\rangle^2$, respectively (see also Vaughan et al. 1994). Because the data are extracted from a set of s independent segments (following Bartlett's method) and the frequencies are rebinned with m_l original frequencies per new frequency bin l, these variances should be divided by a factor $m_l \times s$.[6]

The model values m_k in Eqn. 4.10 correspond to the cross-spectral model of interest for the astronomer. These values are obtained by defining a model $\tilde{S}(E, \nu_l)$ that is convolved through the instrumental response as in Eqn. 4.5. For each energy channel, the cross-spectrum is then formed as in Eqn. 4.7, and rebinned in frequency to match the frequency grid of the observational data. The quantity resulting from the model, and useful for performing the minimisation technique in Eqn. 4.10, is $\tilde{G}_h(\nu_l) = \tilde{G}_{R,h}(\nu_l) + i\tilde{G}_{I,h}(\nu_l)$. The form of the model $\tilde{S}(E, \nu_l)$ above will be shown in the next Subsection and in Eqn. 4.13.

Taking into account the considerations above, Eqn. 4.11 becomes:

$$\chi^2 = \sum_h^H \sum_l^L \frac{(\tilde{D}_{R,h}(\nu_l) - \tilde{G}_{R,h}(\nu_l))^2}{\sigma^2_{R,hl}} + \sum_h^H \sum_l^L \frac{(\tilde{D}_{I,h}(\nu_l) - \tilde{G}_{I,h}(\nu_l))^2}{\sigma^2_{I,hl}} \quad . \tag{4.12}$$

Minimising the χ^2 is equivalent to maximising the likelihood that the model parameters produce a model that describes well the observed cross-spectrum $\tilde{D}_h(\nu_l)$.

[6] The cross-spectral matrices in a file format that is suitable for the cross-spectral analysis shown in this Chapter can be obtained from event lists using the nu Timing Suite that was written to extract all observational data in the present Thesis.

4.2.4 Implementation

The above cross-spectral fitting method has been implemented by John Houck (MIT) and Pablo Cassatella within the module `xspectra`[7] that is being tested for ISIS (Houck and Denicola, 2000) pre-2.0.0[8].

The combination of this ISIS module plus ISIS intrinsic capabilities allows the user to load cross-spectral matrices and instrumental response files from FITS files, define cross-spectral models, perform the fitting procedure to obtain parameter values that best describe the data, as well as obtain confidence intervals.

The model $\tilde{S}(E,\nu_l)$ that is convolved through the instrumental response to produce $\tilde{X}_h(\nu_l)$ should in general be formed by the product of two components, an input noise process $\hat{x}(E,\nu)$ (although generally independent of energy) and the transfer function of a variability process of interest $\hat{\psi}(E,\nu)$ that modulates the phase and amplitudes of the input noise process. In this case, the two quantities are multiplied together, $\tilde{S}(E,\nu_l) = \hat{x}(E,\nu)\hat{\psi}(E,\nu)$. The folding kernel function in `xspectra` ensures that the cross-spectra are formed after convolution of $\tilde{S}(E,\nu)$ through the instrumental response (Eqn. 4.7).[9]

Following the simple example in Eqn. 4.2, the cross-spectrum between two bands A and B (equivalently, between a channel h and its complement, or reference band) results in:

$$\begin{aligned}\hat{y}_A^*(\nu)\hat{y}_B(\nu) &= \hat{x}^*(\nu)\hat{x}(\nu)\psi_A^*(\nu)\psi_B(\nu) \\ &= |x(\nu)|^2\psi_A^*(\nu)\psi_B(\nu) \quad .\end{aligned} \tag{4.13}$$

The quantity $|\hat{x}(\nu)|^2$ represents the power of the noise process $\hat{x}(\nu)$. When performing fits of cross-spectra, the shape of $|\hat{x}(\nu)|^2$ cannot be observed directly because it is modified by the transfer functions. In e.g. the disc propagation scenario in black hole X-ray binaries, the power $|\hat{x}(\nu)|^2$ could be, however, very closely related to the power produced by mass-accretion rate ($\dot{m}$) perturbations in the accretion disc.

An example listing of commands in a typical cross-spectral fitting session within ISIS would include:

```
rsp = load_rmf("data/LOFT_Requirement_v4_rbnA.rsp");
xs = xs_load("data/j1753_sim.xsp");
id = xs_data(xs, rsp);
xs_model("xs_broken_pl(1)*refl(1)")
```

[7] Will be available at `http://space.mit.edu/cxc/isis/modules.html`

[8] Available via `git clone -b pre2 http://space.mit.edu/cxc/isis/isis.git` at the time of writing

[9] In real astrophysical systems, the same or different input processes may excite the variability produced by different, independent transfer functions $\hat{\psi}(\nu)$, leading to an observed variability that results in complicated phase and amplitude dependences on frequency and energy, as well as loss of coherence (see e.g. Vaughan and Nowak 1997). These complicated cases are beyond the scope of this Chapter.

```
compute_model();
fit();
conf_loop([1:5],1)
```

Listing 4.1: Sample cross-spectral fitting session in ISIS

In Listing 4.1, the response file and cross-spectral file are loaded in lines 1 and 2, respectively. The combination of cross-spectrum and response are loaded as a new dataset (line 3) before defining a model (line 4). The models `xs_broken_pl(1)` and `refl(1)` are examples of the functions $\hat{x}(\nu)$ and $\hat{\psi}(\nu)$.

As mentioned earlier, a variable signal that leaves a physical system is, in the Fourier domain, the result of the multiplication of an input (unknown) signal by a transfer function of the variability process in question, for each energy considered.

The commands that follow are used to evaluate the model, fit it, and get 90% confidence contours for the first 5 parameters of the model (`compute_model()`, `fit()` and `conf_loop()`, respectively).[10]

4.3 Simulations

The above technique can be tested by simulating cross-spectral values for a given set of input parameters and performing cross-spectral fits to recover the input parameter values, as described in Section 4.2.3.

The simulated cross-spectrum is obtained as follows. Firstly, flicker noise[11] obtained using the method in Timmer and Koenig (1995) is used to generate an input time series $x(E,t)$ with a fractional rms amplitude of 20% and a 0.3 – 10.0 keV count-rate of 10^6 counts per second (so as to emulate the count-rates that could be obtained using *LOFT*). The counts are distributed among energies following a black-body spectral distribution with a temperature parameter $kT = 0.4$ keV. An example lightcurve of such flicker noise process and its power-spectral distribution can be found in Fig. 4.2.

Secondly, a time-dependent response function $h(E,t)$ is defined. We choose a response function consisting of two sine functions of equal amplitude amplitude $A = 30$ and frequencies $\nu_1 = 2.3$ Hz (period of 0.434 s) and $\nu_2 = 14.2$ Hz (period of 0.07 s). These sinusoids are equal

[10] The command `conf_loop()` is executed in parallel, speeding up the parameter-search process by a factor roughly equal to the number of processors.

[11] A process whose power spectral density (PSD) has a $\propto \nu^{-1}$ dependence on frequency, and has the property that it carries the same amount of power in each octave considered.

in amplitude for every frequency, however their phases are set to be energy-dependent and are shifted with a log-linear relation, $\varphi(E) = \log_{10}(k_{\rm ll}, E)$ where $k_{\rm ll} = 0.3$.

At every energy, the model $y(E,t)$ that is used to calculate the cross-spectrum is obtained by convolving the input $x(E,t)$ with $h(E,t)$. The resulting $y(E,t)$ values are then convolved through the instrumental response for every time t and Poisson deviates are drawn to simulate the effects of counting noise that are present in real-life observations.[12] In practice, the above is done for 982 segments, which corresponds to 3 kiloseconds of exposure time using a time bin $\Delta t = 0.005965$ s and $N = 512$ time bins per segment. The spectral response used is that of EPIC-pn onboard *XMM-Newton*, rebinned to 30 channels that are logarithmically-spaced (the same used in Chapter 3).

For each segment, cross-spectra are constructed according to Eqn. 4.7. The resulting mean and variance of the distribution of the real and imaginary components of the cross-spectrum are rebinned geometrically in frequency into 7 bins, using a geometric rebinning factor of 1.15 (Nowak et al., 1999). These parameters yield cross-spectral matrices of size 30×7.

The fitting procedure using ISIS assumes the same model used for the simulations (defined with `xs_model("xs_broken_pl(1)*sinusoids(1)")`.

The model `xs_broken_pl(1)` is the square root of a doubly-broken power-law dependence on frequency, and we fix its power-law slope to -0.5 at all frequencies, so that its square will give -1.0 as desired in a power density spectrum of a flicker noise process.

The fitting procedure recovers the centroid frequencies of the sinusoids to within less than 3 per cent, while the black-body temperature and the log-linear dependence of phase on energy parameter match the input parameters perfectly (Table 4.1). Figures 4.3, 4.4 and 4.5 and 4.6 show the comparison between simulated and best-fit real and imaginary components, as well as resulting phase-lags, respectively.

The low reduced χ^2 value (~ 0.53) is likely caused by the error-bars (which are set here as the standard deviation of the real and imaginary components for each energy and frequency bin) being over-estimated. This effect can occur if these errors are correlated, which leads to a non-diagonal covariance matrix. The cross-spectral components derive from Eqn. 4.13, where both components of the cross-spectrum depend linearly on the modulus squared (or power) of the (energy-independent) input noise.[13] This component defines a relation whereby all the cross-spectral components at a given frequency will be linearly correlated, producing too large an error-bar, hence a χ^2 value lower than unity.

[12] In practice, it is straightforward to compute the product of the Fourier transform of $x(E,t)$ and $h(E,t)$, $\hat{x}(E,\nu)$ and $\hat{\psi}(E,\nu)$ respectively, and obtain the time-domain equivalent via an inverse Fourier transform afterwards (Eqn. 4.2), which is necessary to apply Poisson-induced noise.

[13] It is worth noting that this dependence affects only the norm of the cross-spectral vector in the complex plane, leaving its phase unaffected.

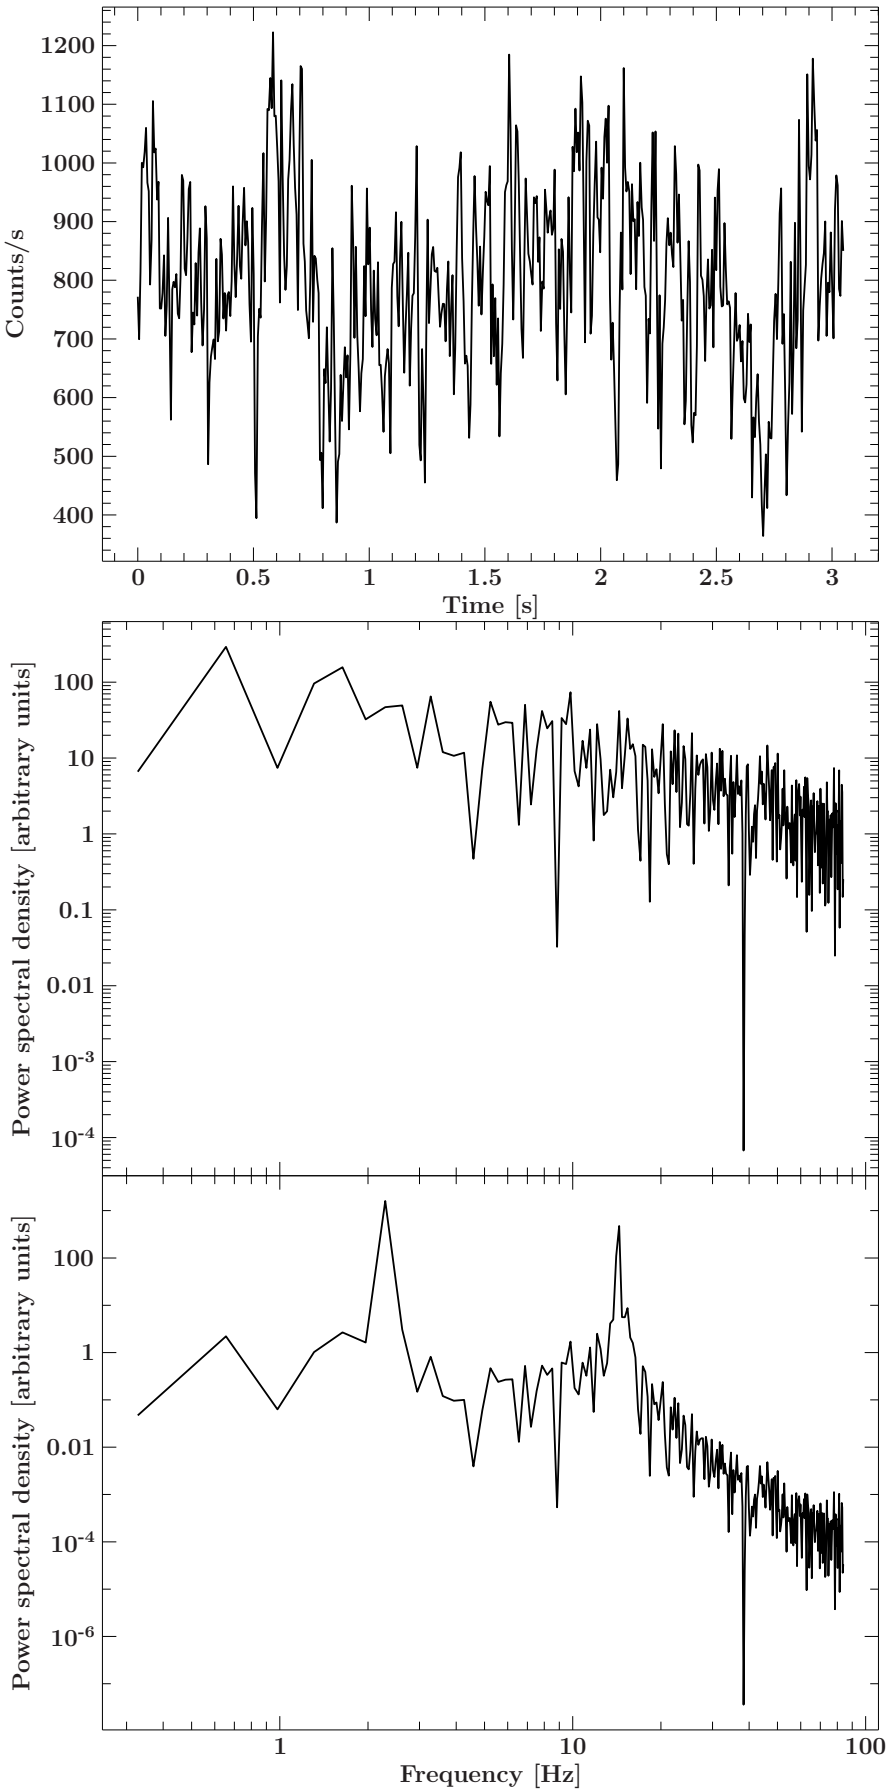

Figure 4.2: Lightcurve (top) and relative power spectral density of one possible realisation of the input flicker noise before (middle panel) and after (bottom panel) convolving with the sinusoid model

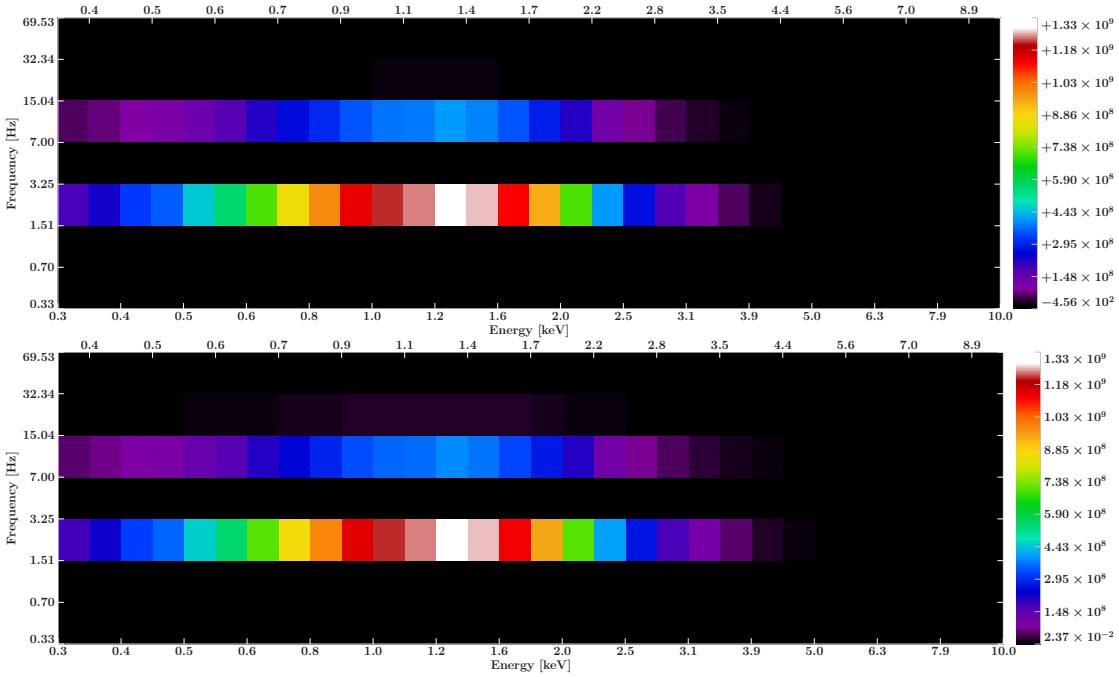

Figure 4.3: Real component of the simulated cross-spectrum (top) and real component of best-fit model (bottom)

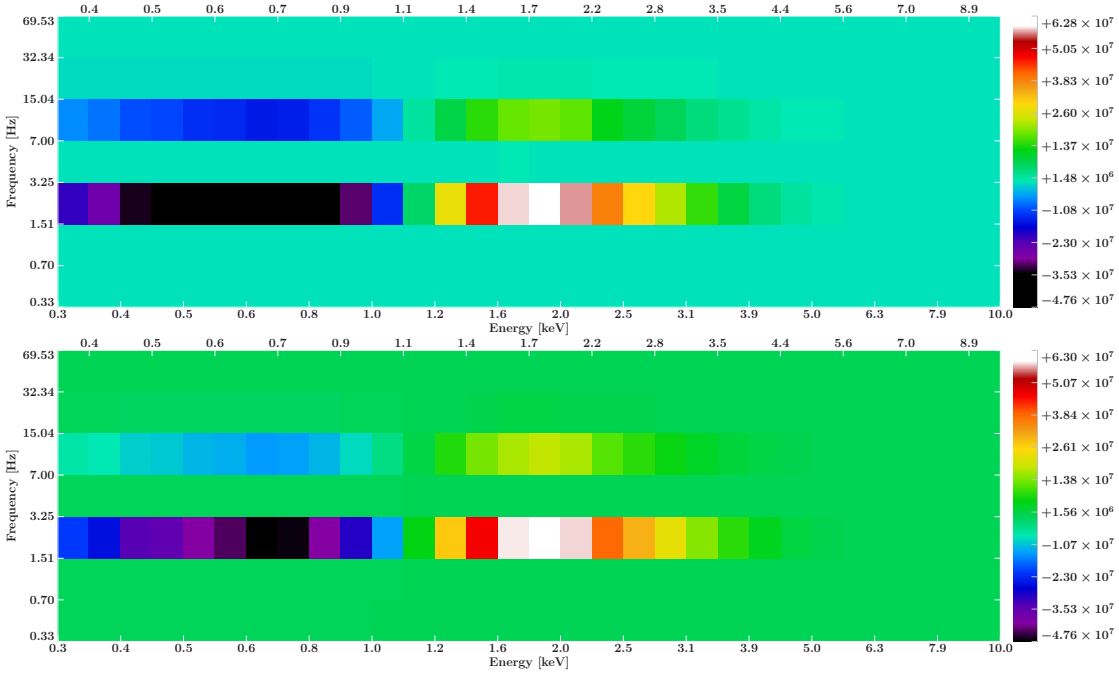

Figure 4.4: Imaginary component of the simulated cross-spectrum (top) and imaginary component of best-fit model (bottom)

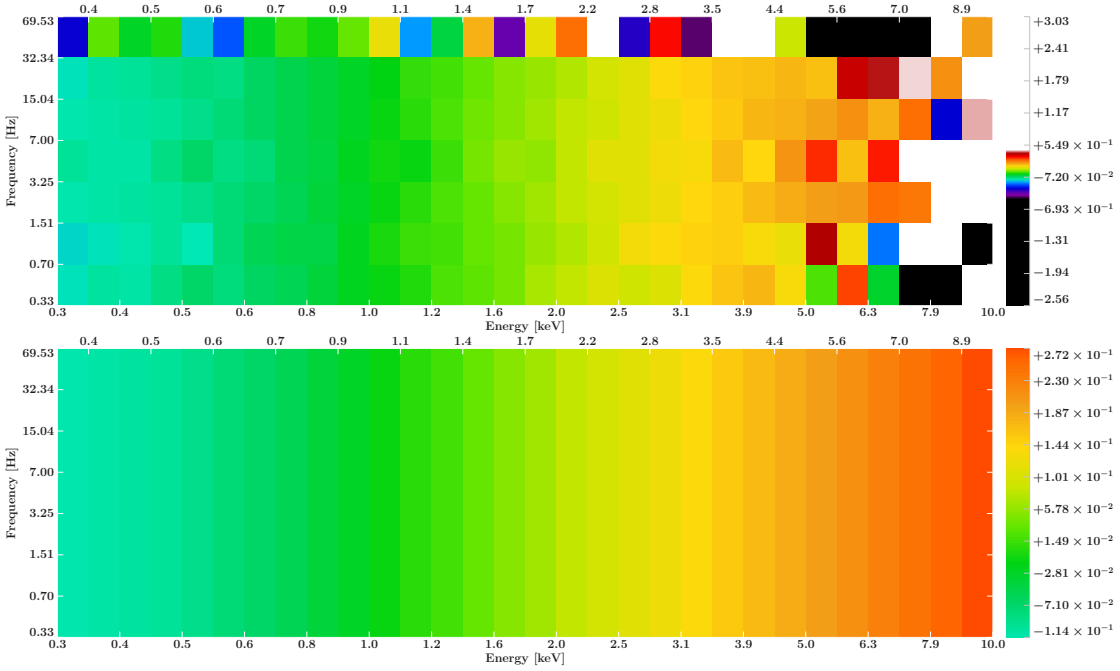

Figure 4.5: Phase-lag dependence of simulated cross-spectrum (top) and best-fit phase-lag dependence (bottom). The same colour scale is used to show the similarity.

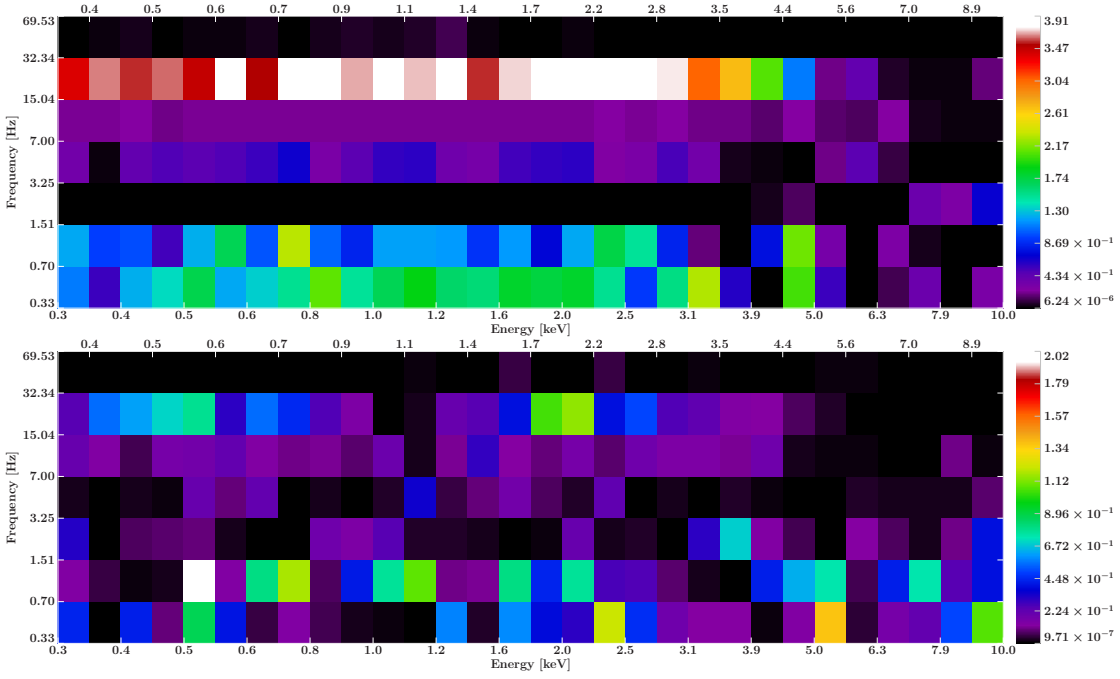

Figure 4.6: χ^2 value of best-fit model per real (top) and imaginary (bottom) cross-spectrum energy-frequency bin

Table 4.1: Input parameters and recovered parameters from the simulation.

	Input parameter	Recovered parameter
kT [keV]	0.40	0.400 ± 0.001
ν_1 [Hz]	2.30	2.264 ± 0.001
ν_2 [Hz]	14.20	14.500 ± 0.003
$k_{\rm ll}$	0.30	0.301 ± 0.004
χ^2 / d.o.f.	207 / 413	

4.4 Discussion

The method outlined here represents the first step towards constraining the X-ray properties of black hole X-ray binaries (as well as other objects such as neutron stars and Active Galactic Nuclei) from a concurrent combination of spectral and variability information. The simulations ran for this Chapter assume count-rates comparable to those expected from observations of bright black hole X-ray binaries with the proposed Large Area Detector onboard *LOFT* and show that the input parameters that link frequency- and energy-dependent information can be recovered with a relative error of less than 3 per cent. The extraction of the overall normalisation of the cross-spectrum as well as a more precise estimation of its errorbars is beyond the scope of this chapter, but should be addressed in the future (see e.g. Sivia and Skilling 2006 for a treatment of correlated errors).

This technique can be applied to several problems in the current understanding of X-ray binary physics where the link between spectral emission and the properties of variability (power spectral amplitudes, and time-lags) is necessary. In general though, these problems are extremely difficult to unravel with current techniques and instrumentation. Such problems include the understanding of disc truncation and its link to General Relativity.

Several methods have been used to understand whether the disc truncates in the low-hard state or whether it reaches the Innermost Stable Circular Orbit (ISCO). The understanding of the position of the ISCO provides a probe of General Relativity that cannot be equalled in laboratory experiments at such large scales. Tests with standard, time-average spectra include disc emission (McClintock et al., 2006; Kolehmainen and Done, 2010) as well as disc reflection (e.g. Reis et al. 2010) fitting methods. So far these methods have not produced reliable results, which are often instrument- and calibration-dependent. Reflection does however show imprints in time-scale-dependent spectra that can be observed in a model-independent way (Revnivtsev et al., 1999), as well as a correlation between the break frequency in power spectra and the amount of reflection (Gilfanov et al., 1999). Reflection is also known for producing time-lag signatures (Fabian et al., 2009) and can therefore be modelled to infer the geometry of the accretion disc (see Chapter 2 and e.g. Kotov et al. 2001) to complement the deficiencies of instrumental energy resolution and calibration. Figure 4.7 shows a smooth example of reflection imprints on the phase-lag dependence on frequency and energy (especially the

iron line at 6.4 keV and the Compton hump above 10–20 keV) of a $10M_{\odot}$ black hole using a model similar to that used in Chapter 2.

Another interesting challenge is the study of the timing signatures of the disc, to test both the propagation scenario (e.g. Arévalo and Uttley 2006) and the quasi-periodic signatures, in particular with the Lense-Thirring precession model (Ingram and Done, 2012).

The properties of quasi-periodic oscillations (QPOs) are, in fact, known to depend on energy in a non-trivial way (see e.g. Rodriguez et al. 2002). Figure 4.8 shows the $\sim 2-6$ keV power-spectral density of a ~ 3.4 ks observation of GRO 1655–40 taken with the Proportional Counter Array (PCA) onboard *RXTE* (observation 90704-04-01-00 on 2005 March 10). The count-rate is ~ 1700 counts/s. The PSD shows a strong QPO that is centred at ~ 2.3 Hz, with a harmonic component. Standard fitting of the QPO shape and underlying 'noise continuum' would be done with multiple Lorentzian components, whose properties may vary along the outburst evolution. Yet PSD fitting does not yield information about the energy-dependent properties of variability –i.e. what spectral components are mostly responsible for the QPO production and how the variability in the QPO is delayed with respect to the rest of the variable emission. In Fig. 4.8 we show the fractional rms squared amplitude of the cross-spectrum (absolute magnitude of the cross-spectrum divided by the product of the count-rate in each h channel with the count-rate of the hth reference band). The QPO clearly shows a dependence on energy that can be modelled to help understand the origin of this component. The phase-lags resulting from the model should also match the observed phase-lags (Fig. 4.10), so that the resulting best-fit model parameters produce a model that self-consistently reproduces frequency and energy dependence on variability as well as causality in the system.

Finally, the method could also be used with future models that can test the combination of lags produced by viscous disc propagation and disc reflection via a combination of transfer functions, as well as to understand incoherent variability in the disc (see Chapter 3).

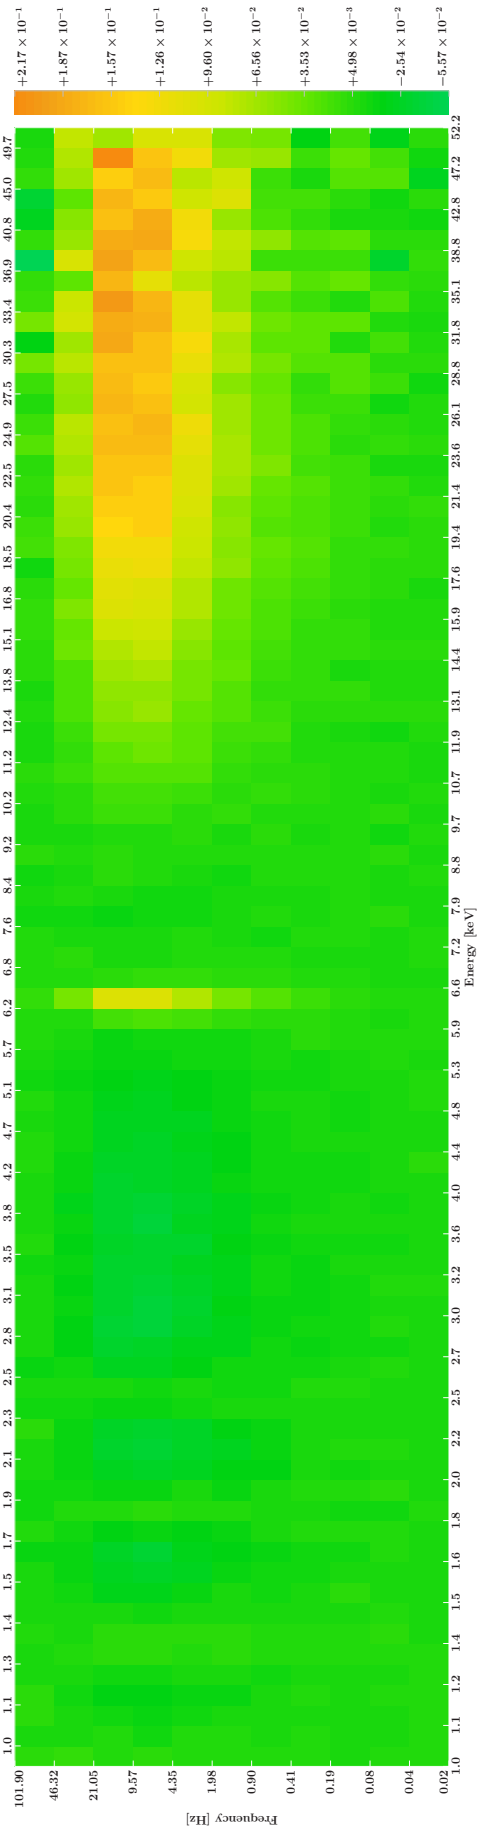

Figure 4.7: Simulated effect of reflection on the dependence of phase-lags on Fourier-frequency and energy of a $10M_{\odot}$ black hole observed during a 3 ks exposure with *LOFT*, using the preliminary response files (version 4). The count-rate is 10^6 counts/s. The parameters used for the disc geometry are comparable to those inferred in Chapter 2.

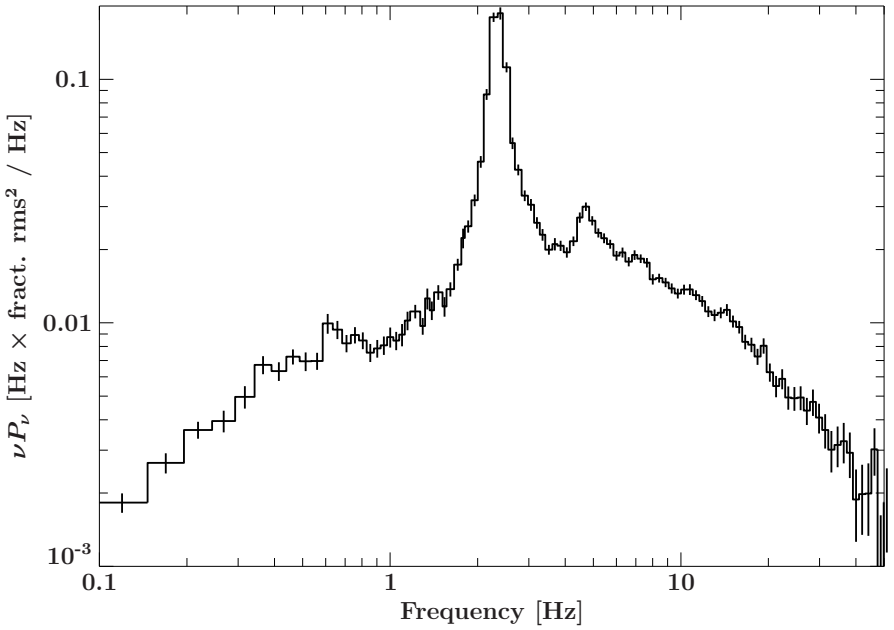

Figure 4.8: Power Spectral Density of GRO 1655–40 observed on 2005 March 10 in the band ~ 2 – 6 keV. It shows a clear, strong QPO component at ~ 2.3 Hz and a harmonic at twice the QPO frequency.

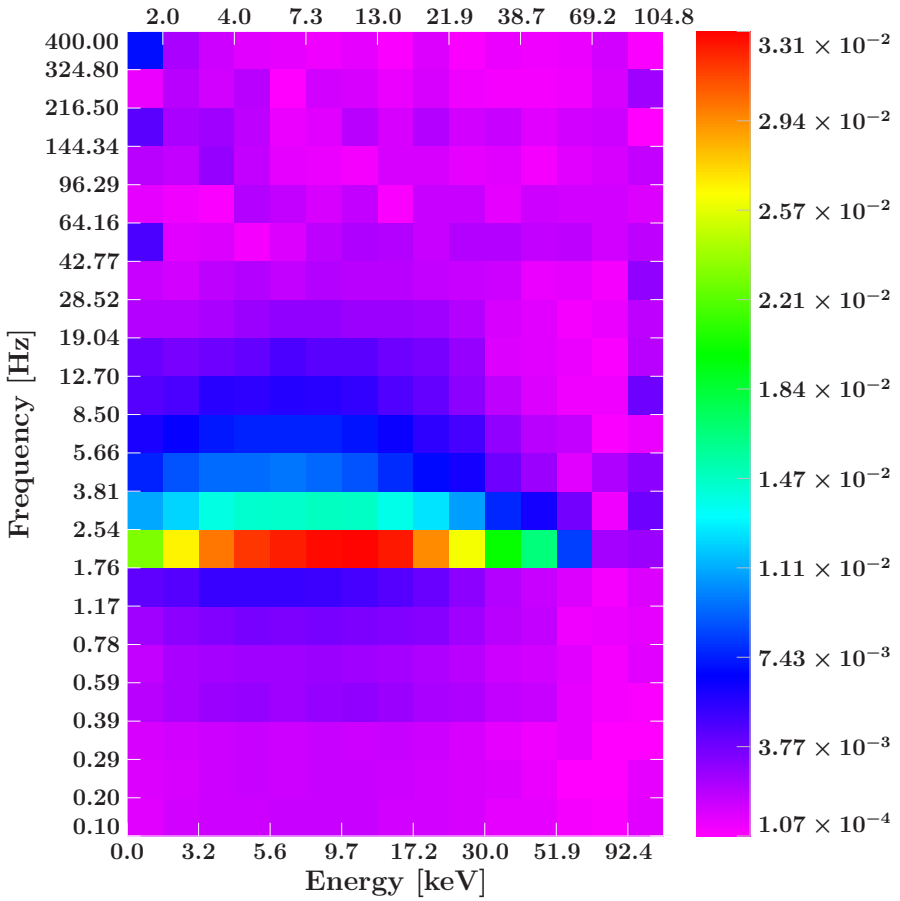

Figure 4.9: Fractional squared rms variability per frequency-energy bin that is correlated with the reference band (fractional cross-spectrum), using GRO 1655–40 data. The QPO component seen in Fig. 4.8 is clearly dominating the contribution to the total fractional variability. The Figure also shows that the QPO frequency width varies as a function of energy. This energy dependence can be tested with the method explained in this Chapter and an appropriate variability model.

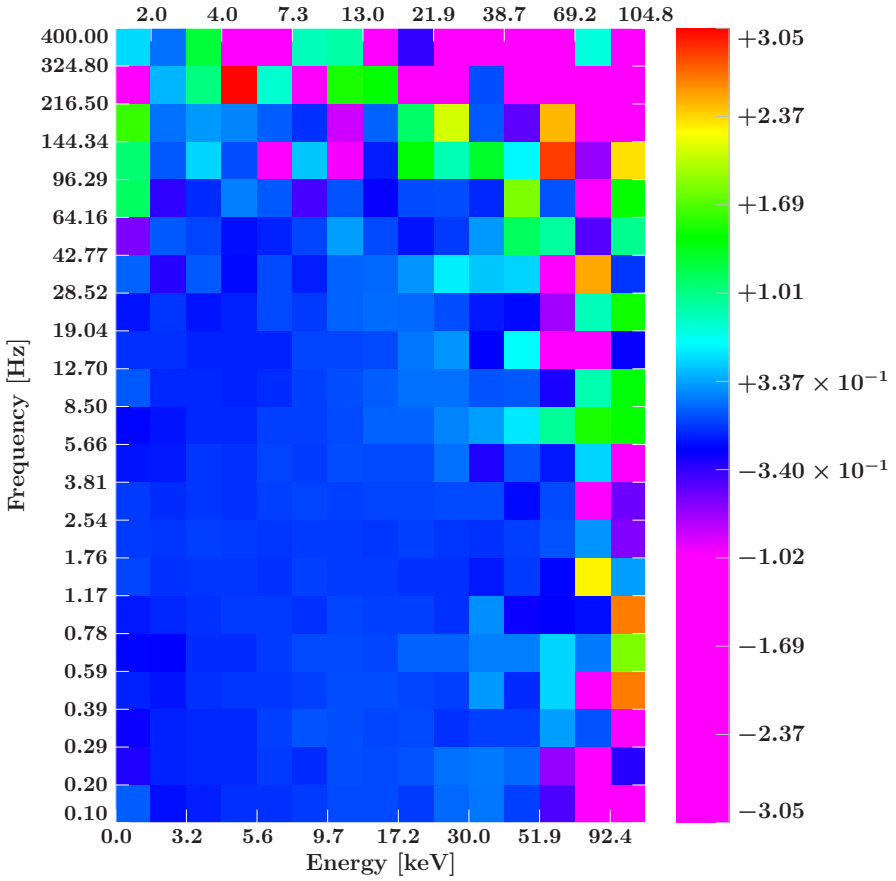

Figure 4.10: Phase-lag dependence on frequency and energy for the GRO 1655–40 observation

CHAPTER 5

Conclusions

5.1 Introduction

The focus of this Thesis has been on applying state-of-the art techniques in spectral-timing to the observed emission of black hole X-ray binaries. We have approached the study of the complex physics of black hole X-ray binary emission from the spectral-timing point of view, where the time variability seen in lightcurves in different energy ranges (alternatively, in multi-channel time series) can be studied in the Fourier space. This transform of the data is able to provide important information regarding some of the mysteries revolving around the properties of accreting systems.

Fifty years ago we had not even seen a black hole X-ray binary, and it was only thirty years ago that we saw the first case of a quasi-periodic oscillation (in Cen X–1, Giacconi et al. 1971), using nothing but naked-eye inspection of a lightcurve. Since then, X-ray instrumentation has achieved a much higher degree of complexity and is now pushing the frontiers of human knowledge in the field. Analysis tools and methods have also evolved and X-ray astronomy has borrowed many analysis techniques from fields such as acoustics, electronics and geological sciences.

Despite the improvement in imaging capabilities of focusing telescopes together with CCD detectors, they will never[1] be able to resolve the core of black hole X-ray binaries, e.g. the innermost regions of the accretion disc that should be in a non-Euclidean space-time such as the metrics predicted in General Relativity under Strong Gravity. Spectral analysis is next

[1] Never say never! Let us say highly unlikely given our current knowledge and predicted future technology.

in the pipeline for studying these regions, and much effort is being put on improving energy resolution and increasing the collecting area of future instruments, such as the Large Area Detector onboard *LOFT*. In X-ray binaries, spectral analysis has made it possible to discover the three main mechanisms that produce the spectrum seen in the X-rays, namely Comptonisation (or synchrotron emission, which produces a similar power-law in the X-ray band), disc-blackbody emission and disc reflection. In principle, understanding these three components in a variety of objects would give a clear picture with regards to their origin, e.g. in terms of energetics and solid angles.

5.2 Importance of this work

We have treated the spectral-timing study of black hole X-ray binaries from all possible fronts, namely spectral-timing modelling, data analysis and development of a novel technique for inferring physical parameters from multi-channel cross-spectral values. We have shown that a broad exploratory space is still open for study in the field, and most importantly, spectral-timing helps break the degeneracies that are inherent to spectral modelling when performed alone. Future X-ray timing missions will provide an unprecedented combination of energy and timing properties that are optimal for these studies. In the following, we highlight the most important findings that have been shown here.

- The reflection model discussed in Chapter 2 was developed to provide more precise measurements of the geometrical properties of the accretion disc by combining GX 339–4 hard-state energy spectra and frequency-dependent time-lags together. The modelling has shown that, in principle, spectral-timing modelling and fitting can be performed by extracting response functions from processes that transfer variability from one component to another. The best-fit results show that only extreme and unlikely disc geometries could reproduce the large lags observed while producing a picture that is consistent with the observed shape of the reflected iron line in the spectrum. Although these geometries are extreme, they imply only a small fraction of the emission observed in the optical/ultra-violet band. In addition, the fitting work done here confirms that there must exist a variability process that produces frequency-dependent lags that have to sum up to the observed lag amplitudes. The mass-accretion propagation model is a natural candidate, as it can explain already the intermediate-to-soft and hard-to-soft lags in the context of softer disc perturbations modulating the harder power-law emission on viscous time-scales.
- In Chapter 3, we have studied the differences in the spectral-timing properties of the source SWIFT J1753.5–0127 in the hard state at two different luminosities, the 2009

observation showing an approximately six times more luminous hard state than the other observation. This study has been possible thanks to a careful manipulation of time-tagged events to provide a self-consistent picture of frequency-dependent emission properties. We have shown that, despite both observations having been obtained in the same spectral state, clear differences appear beyond simple differences between the relative amount of emission in the disc and power-law components. A remarkable difference is that the disc becomes less-variable at higher luminosities and its variability appears 'suppressed', in clear contrast with the dimmer 2006 observation where the highest fractional rms variability was seen in the disc emission. We have shown that the difference in disc variability amplitudes can be reproduced if the same disc emissivity profile is assumend and assuming that ~ 75 per cent of the disc emission in 2006 was produced by reprocessing in a very tiny inner region of the disc. If the disc variability originates at the same disc radius, the difference in disc amplitudes can be reproduced. By comparing covariance and rms spectra (together with the coherence), we have shown that the outer, colder disc in 2009 may be producing slow accretion fluctuations that are likely damped before they can reach the region mostly responsible for the coherent fluctuations that are transferred to the power-law variability. These observed properties may be linked to the shape of power density spectra, where the noise extends to lower frequencies when only coherent variations are seen. This work opens up the exploratory space towards the understanding of the filtering properties of the disc as well as the link between the (non-)linearity of its modulations and properties such as the energy balance between the disc and power-law emission as well as $\dot{m}$.

- Chapter 4 shows a novel technique that has been developed to take full advantage of combined spectral and timing information contained in high-time resolution observations. The method allows one to fit complex-valued cross-spectral matrices as a function of energy and Fourier-frequency using 2D models of variability, and enables the astronomer to break some of the model degeneracies that appear in time-averaged spectral fitting. This technique is especially appropiate for high signal-to-noise observations with satellites such as *RXTE*, *LOFT* and *NuSTAR*.

5.3 Future prospects

Much progress towards the understanding of black hole X-ray binaries has been made in recent years, with an enormous contribution from *RXTE* in the spectral-timing world until the end of 2011 (see e.g. Nowak et al. 1999, Revnivtsev et al. 1999, Kotov et al. 2001).

The launch of *XMM-Newton* led to the extension and further development of the spectral-timing techniques used in earlier years to energies that cover the disc-dominated band much

better than *Ginga*, see e.g. Vaughan et al. 1994 for a study of phase-lags as a function of energy down to 1 keV using *Ginga* in the case of the neutron star GX 5-1. Wilkinson and Uttley (2009) and Uttley et al. (2011), together with the work presented in Chapter 3, are examples of spectral-timing studies of BHXRBs with soft X-ray coverage, which show that the study of the properties of accretion flows can benefit greatly from the spectral-timing approach.

There are a few topics that are perfectly in line with the studies presented in this Thesis and would deserve a careful analysis. These are the following:

1. Spectral-timing modelling of the accretion disc as a function of Fourier-frequency as well as energy to fit data in packages such as ISIS or XSPEC. This would be a perfect complement to the reflection modelling work shown earlier and could be used to confirm that the drop in the energy-dependent lags at low energies and at low frequencies is in fact caused by the disc leading the power-law variability. A more extended study of the behaviour of these lags could include a treatment of viscosity, density and pressure beyond the α-model, depending on computational resources and future technology.

2. Comprehensive study of the coherence function along the hardness-intensity diagram for different sources and energy bands, in order to understand what relation links drops in coherence to the emission and variability properties. The drop in coherence may be linked to instabilities in the disc (e.g. radiation pressure).

3. A spectral-timing analysis of the energy-dependent time-lags could expand the current knowledge of the behaviour of the boundary layer in neutron stars.

4. The log-linear law governing the time-lags above the disc energies. Future work could be done towards understanding whether the slope of the log-linear law correlates with parameters such as viscous time-scale, mass-accretion rate and α parameter in the accretion flow.

5. The cross-spectral method shown in this Thesis deserves further development and new computational tools to decrease the computation time of complex-valued multi-dimensional matrices. A possibility would be to exploit the capabilities of modern Graphics Processing Units (GPUs) to perform the Fourier transforms required for this method, and similar methods for fast instrumental response convolution.

6. The cross-spectral method should be used in conjunction with QPO models such as the Lense-Thirring precession model shown in Ingram and Done (2012) to constrain the origin of the QPO variability.

APPENDIX A

The causal connection between disc and power-law variability in hard state black hole X-ray binaries

The present Chapter has been published as Uttley et al. (2011) and has been added to this Thesis in its integrity. The contribution of the author of this Thesis to this Chapter is limited to the reduction of the data presented here.

We use the *XMM-Newton* EPIC-pn instrument in timing mode to extend spectral time-lag studies of hard state black hole X-ray binaries into the soft X-ray band. We show that variations of the disc blackbody emission substantially lead variations in the power-law emission, by tenths of a second on variability time-scales of seconds or longer. The large lags cannot be explained by Compton scattering but are consistent with time-delays due to viscous propagation of mass accretion fluctuations in the disc. However, on time-scales less than a second the disc lags the power-law variations by a few ms, consistent with the disc variations being dominated by X-ray heating by the power-law, with the short lag corresponding to the light-travel time between the power-law emitting region and the disc. Our results indicate that instabilities in the accretion disc are responsible for continuum variability on time-scales of seconds or longer and probably also on shorter time-scales.

A.1 Introduction

Rapid X-ray variability on time-scales of seconds to milliseconds is a key characteristic of X-ray binaries (XRBs) which has been studied almost since their discovery. It was suggested quite early on that the rapid variability observed in XRBs and on much longer time-scales in active galaxies, was produced by instabilities in the 'standard' accretion disc (Lightman and Eardley, 1974; Shakura and Sunyaev, 1976). Such a model is attractive because the disc can produce variations over a broad range of time-scales, corresponding to the range of unstable disc radii, to match well with the observed variability (Lyubarskii, 1997).

However, for many years the variability has been associated with the still-mysterious hard X-ray power-law emitting component (e.g. Galeev et al. 1979; Poutanen and Fabian 1999; Churazov et al. 2001; Done et al. 2007 and references therein). This is because, although strong variations in the luminosity of the disc blackbody component are seen on long time-scales corresponding to transitions between the soft and hard spectral states, the disc-dominated soft states show very little rapid variability (fractional rms $\sim$ 1 per cent, e.g. Belloni et al. 2005). In contrast, rms of several tens of per cent is seen in the hard spectral states which are energetically dominated by the power-law emission thought to be associated with a hot optically-thin flow, magnetised corona, or even the base of the persistent radio-emitting jets seen in these states (e.g. Zdziarski et al. 1998; Droulans et al. 2010; Markoff et al. 2005). In addition, the strong QPOs observed during the state-transitions of black hole XRBs (BHXRBs), when both disc and power-law emission are strong, show hard rms-spectra consistent with an origin predominantly in the power-law emitting component (Sobolewska and Życki, 2006).

In the hard states our understanding of the origin of the variability has been limited by a lack of spectral coverage of the disc, which emits at $kT < 0.5$ keV and so is not covered by the bandpass of detectors such as the Proportional Counter Array (PCA) on the *Rossi X-ray Timing Explorer*, which is the prime workhorse of X-ray timing studies to date. Recently we addressed this problem by studying the X-ray variability 'covariance spectra'[1] of the hard state of the black hole candidate GX 339−4 (Wilkinson and Uttley, 2009) using the soft X-ray response of the EPIC-pn CCD detector on board the *XMM-Newton* satellite. In that work, we found that the disc blackbody emission does vary with an rms amplitude of tens of per cent and that the blackbody variability amplitude is greater relative to the correlated power-law variations on longer time-scales (> 1 s). We used these results to argue that, although on time-scales < 1 s disc blackbody variability could be explained by X-ray heating of the disc by the varying power-law emission, on longer time-scales the disc is intrinsically variable. A natural physical interpretation is that instabilities in the standard disc are responsible for

[1] The covariance spectrum measures the spectral shape of variations which are correlated with a given reference band, and so is ideal to search for correlated disc and power-law variability, see Wilkinson and Uttley (2009) for details.

driving the observed power-law variability on time-scales at least down to a few seconds, perhaps through mass-accretion variations which propagate through the disc (Lyubarskii, 1997; Arévalo and Uttley, 2006), before reaching the power-law emitting region.

An even stronger argument can be made for disc-driven variability if we can establish the *causal* relationship between correlated variations in the disc and power-law components, through the detection of X-ray time-lags between variations in different energy bands. Previous work to measure time-lags in BHXRBs has used proportional counter instruments to measure lags between variations at harder energies, above 2 keV (e.g. Miyamoto and Kitamoto 1989; Nowak et al. 1999), however we can also do this analysis using the same *XMM-Newton* EPIC-pn timing mode data that we use to study the covariance spectrum. In this Letter we use these data to present the first time-lag study of the hard state which is extended to soft X-rays to cover the disc component. We confirm that there is indeed a clear signature in the lags at soft energies which shows that variations on time-scales of seconds or longer in the disc blackbody emission lead correlated variations in the power-law emission. We also show that the sign of the lag changes on time-scales < 1 s, so that the disc component lags behind the power-law variations by a few milliseconds. These lags are consistent with light-travel time-lags expected by disc thermal 'reverberation', caused by X-ray heating of the disc by the power-law which dominates over the intrinsic disc blackbody variability on these short time-scales.

A.2 Observations and Data Reduction

We analysed EPIC-pn (Strüder et al., 2001) timing mode data from the 2004 March 16-19 *XMM-Newton* observations of GX-339 in a stable hard state. The data were reduced in the standard manner using SAS 10.0.0, processing the raw data products using the SAS tools `epsplitproc` and `epfast`, taking account of background flaring and extracting only events with RAWX from columns 31 to 45 using the SAS tool `evselect`. Although the tool `epfast` was used on the events file, the Charge Transfer Inefficiency (CTI) corrections were insufficient to remove features at 1.8 and 2.2 keV (Si and Au edges) in the 2004 spectrum. Although the source is relatively bright, previous analyses have shown that the EPIC-pn data are not significantly piled up (Done and Diaz Trigo, 2010). The SAS tools `arfgen` and `rmfgen` were used to generate the ancillary response file (ARF) and redistribution matrix file (RMF) for the 15 columns used for source extraction. Background was not accounted for as it is very small compared to the source and in any case background variations are uncorrelated with those of the chosen reference energy bands and so do not contribute to the lags or covariance spectra. The total exposure time for the combined 2004 data was 155 ks.

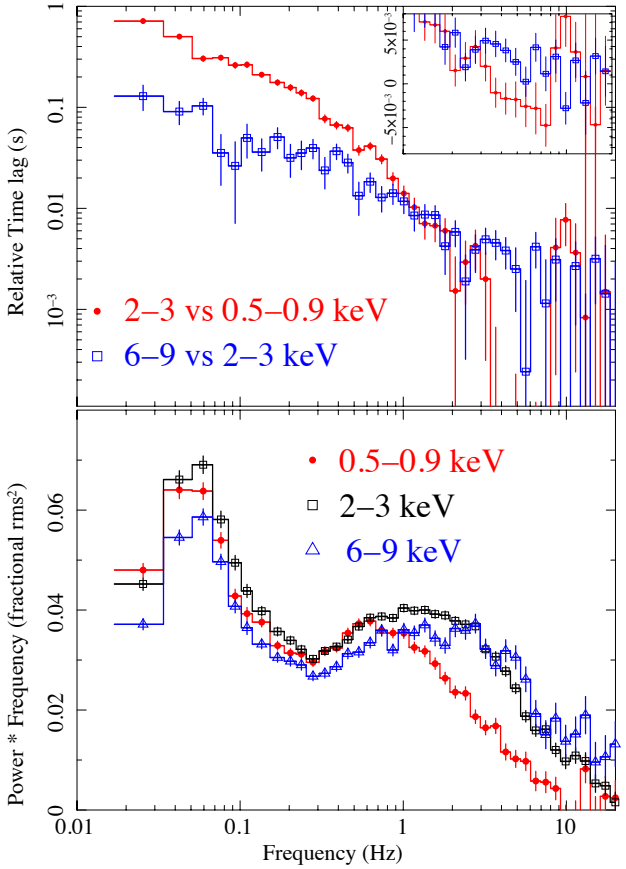

Figure A.1: Top panel: GX 339−4 2004 hard state lag versus frequency dependence for the medium-soft and hard-medium energy bands. Positive values correspond to harder photons lagging softer photons. The inset shows the lags on a linear y-axis in the frequency range where the medium-soft lags drop below the hard-medium lags (the x-scale is matched to the scale on the main figure). Bottom panel: Power spectra for the soft, medium and hard energy bands.

A.3 Analysis and Results

A.3.1 Power and time-lags versus frequency

We used the event lists, together with Good Time Intervals (GTIs) excluding telemetry dropouts, to extract light curves in three geometrically-spaced energy bands (soft 0.5–0.9 keV, medium 2–3 keV and hard 6–9 keV)[2] in order to measure the cross-spectra and determine lags between bands which are separated by equal steps in $\log(E)$. Cross-spectra for pairs of these energy bands were measured in the standard way (Nowak et al., 1999) for contiguous segments of data of length 59 s (the segment-length necessitated excluding data containing short

[2]There is a factor ~ 3 separation in average event energy between each band, with values of 750 eV, 2420 eV and 7140 eV for soft, medium and hard bands respectively.

gaps due to telemetry drop-outs). The cross-spectra were averaged over segments and in geometrically spaced frequency bins with frequency separation $\nu \rightarrow 1.15\nu$. We measured phase lags from the argument of the average cross-spectrum in each bin and divided by 2π times the average frequency of the bin to give the time-lag. The resulting lag-frequency spectra are shown in Fig. A.1, together with the power spectra of the measured energy bands.

From the lag-frequency spectra, we note two striking differences between the lag behaviour in the soft band and the behaviour at harder energies which has already been well-studied by *RXTE*. First, although we chose the energy bands to be separated by roughly-equal steps in $\log(E)$, at low frequencies the medium-soft lags are a factor ~ 6 times larger than the hard-medium lags. Therefore, by extending the measurement of lags into the soft band we have uncovered a significant increase in the lags compared with the values expected from extrapolation of the log-linear scaling with energy observed by *RXTE* above 3 keV (e.g. Nowak et al. 1999), which would give similar lags between pairs of bands separated by the same $\log(E)$. It is also interesting to note that although at low frequencies the medium-soft and hard-medium lag-frequency relations show similar power-law shapes (time-lag $\tau \propto \nu^{-0.7}$, as previously observed), at frequencies approaching 1 Hz the medium-soft lags start to drop rapidly, to the point (above 3 Hz) where they drop below the hard-medium lags. This drop in medium-soft lags seems to correspond roughly to a drop in the soft band power relative to other bands at frequencies approaching and above ~ 1 Hz. In Wilkinson and Uttley (2009) we used the covariance spectra measured on different time-scales to associate this drop in soft band power at high frequencies with a drop in the amplitude of variability intrinsic to the blackbody-emitting disc, so we might expect the disc to play a similar role in explaining the medium-soft lag behaviour.

A.3.2 Lag versus energy and covariance spectra

To better understand the origin of the lag behaviour associated with the soft band, we carried out a Fourier-frequency-resolved study of lag versus energy. To optimise signal-to-noise we measured the cross-spectra for many narrow energy bins with respect to a fixed broad reference band covering the range 0.54–10.08 keV (the reference energies are selected to match the start and end of energy bins). In order to remove the correlated Poisson noise part of the variability from the cross-spectrum we must remove the energy bin light curve from the reference band light curve before measuring the cross-spectrum. Strictly-speaking this means that the reference band is slightly different for each energy bin, but due to the narrow bin widths, the effects on the relative lags are negligible compared to the errors (see also the discussion by Zoghbi et al. 2010).

After measuring the cross-spectrum for each channel with respect to the reference band, we can determine the relative lag versus energy for a selected frequency-range by averaging the cross-spectrum over that range (e.g. see Nowak et al. 1999; Kotov et al. 2001). We can use the

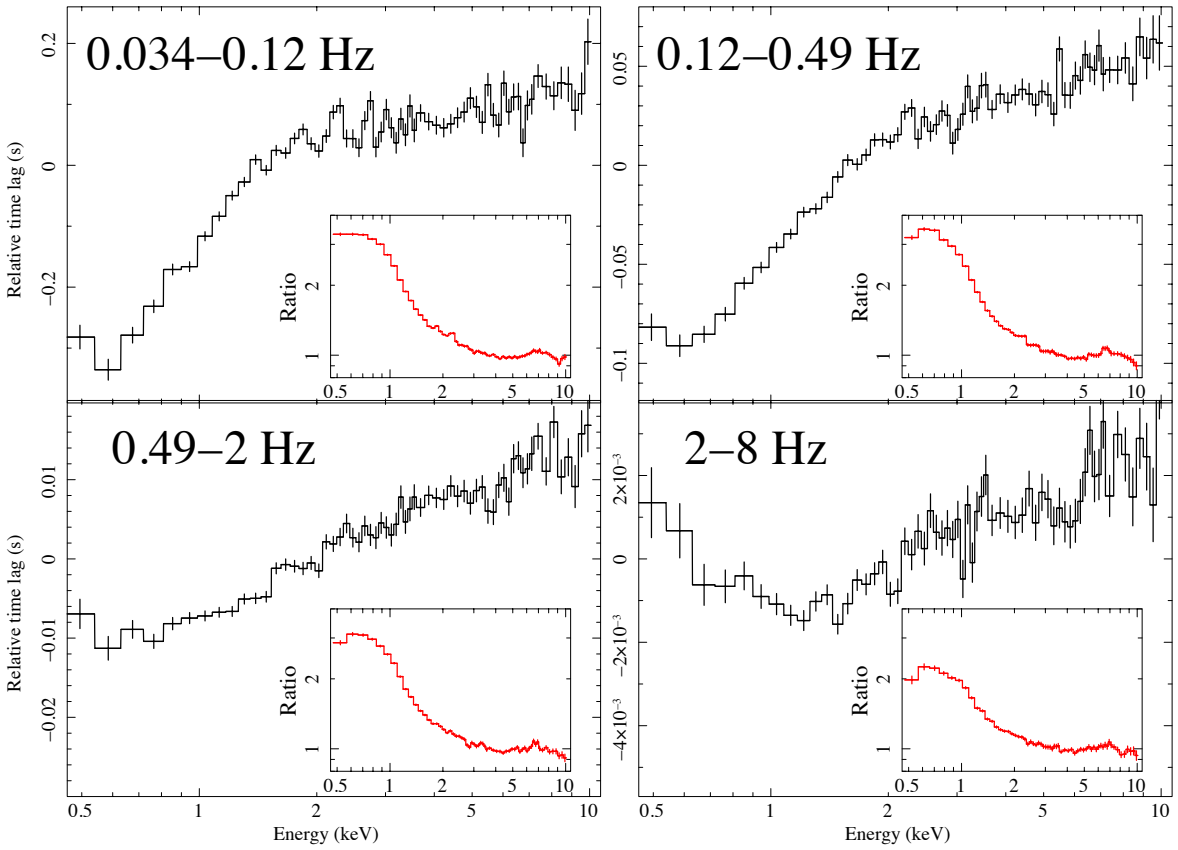

Figure A.2: Lag-energy spectra for four temporal frequency ranges. The lags are all measured relative to the same 0.54-10.08 keV reference band and plotted so that energies with more-positive values of lag are lagging behind energies with less-positive values of lag. The insets show the covariance spectra for the same energy range and frequencies, plotted as a ratio to a power-law with photon index fixed at $\Gamma = 1.55$, absorbed by a 6×10^{21} cm^{-2} neutral Galactic column and fitted to the 3–10 keV range. Note that the scales for the covariance plots are the same, showing that the switch in soft lag behaviour at high frequencies corresponds to a drop in amplitude of disc variability relative to the power-law variability.

frequency-averaged cross-spectra to determine phase lags and hence time-lags in the usual way. We can also use the averaged cross-spectrum to make frequency-resolved covariance spectra, which can be obtained by dividing the square-root of the modulus-squared of the cross-spectrum by the rms of the reference band, in a manner analogous to the way the time-domain covariance spectrum is measured using the cross-correlation function (Wilkinson and Uttley, 2009).

Fig. A.2 shows the frequency-resolved lag-energy spectra for four frequency ranges, each corresponding to an approximately factor 4 increase in frequency starting with 0.034 Hz and ending at 8 Hz. Covariance spectra for the same frequency ranges are shown as insets and are plotted as ratio plots with respect to a simple absorbed power-law model (see Figure caption for more details).

The frequency-resolved lag-energy spectra show clearly that the soft excess emission associated with the disc blackbody is responsible for the complex soft lag behaviour seen in Fig. A.1. The increased amplitude of the low-frequency medium-soft lags results from a significant downturn in the lag-energy relation, away from the log-linear law which applies above 2 keV. The covariance spectra show that this downturn is associated with energies where variable disc blackbody emission starts to become significant. Therefore, the disc blackbody variations are significantly *leading* the variations of the power-law emission which dominate at harder energies.

The covariance spectra show that at high frequencies the disc blackbody variability is still present but drops in amplitude. The same result was found in the time-domain covariance spectra (Wilkinson and Uttley, 2009), our interpretation being that at these frequencies, intrinsic disc variability is small compared to blackbody variations driven by X-ray heating of the disc by the power-law emission. This interpretation appears to be confirmed by the lag-energy spectra. The downturn in low-frequency lag-energy spectra becomes an upturn at high frequencies, so that now the soft photons lag the medium-energy photons (which explains the sharp change in medium-soft lags seen above 1 Hz in Fig. A.1). This behaviour can be simply explained if the mechanism for disc blackbody variability is switching from being dominated by variable viscous heating intrinsic to the disc on time-scales below 1 s, to being dominated by X-ray heating by the illuminating power-law on shorter time-scales. The high-frequency soft lags, which correspond to tens of $R_{\rm G}$ light-crossing time, can be associated with the light travel time from the power-law emitting regions to the blackbody-emitting regions of the disc.

A.3.3 Lags in other hard state systems

To show that other hard state systems show the same low-frequency lag behaviour as GX 339−4, we plot in Fig. A.3 lag versus energy for the 0.125-0.5 Hz range for EPIC-pn timing mode

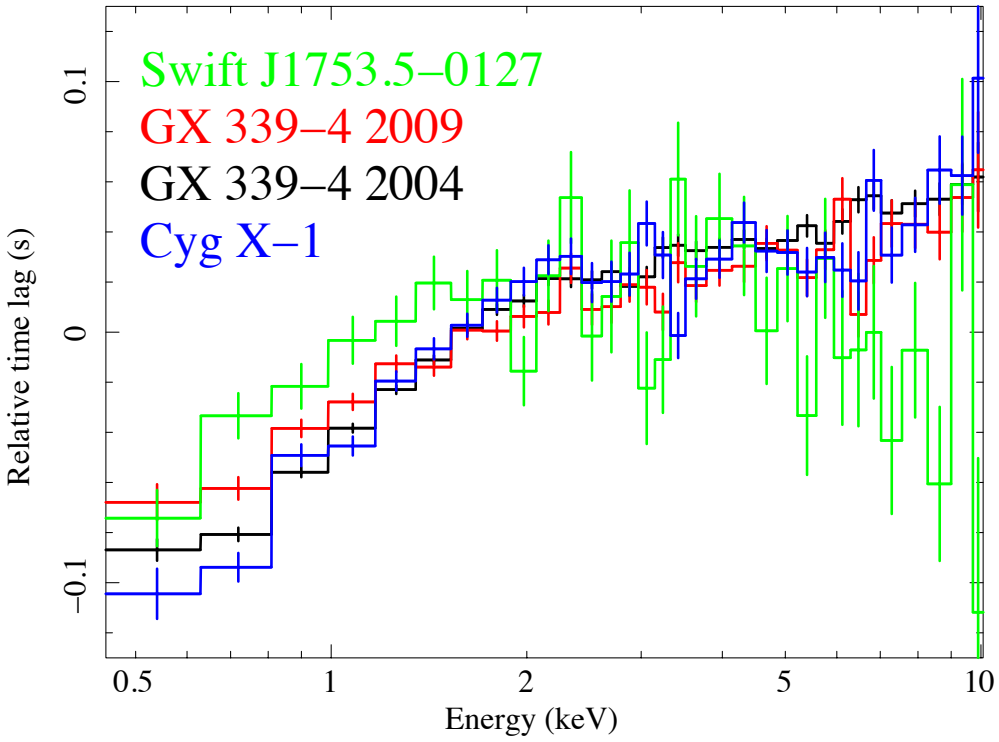

Figure A.3: Lag-energy spectra of 0.125–0.5 Hz variations for several hard state systems

data from *XMM-Newton* observations of hard states in GX 339−4, Swift J1753.5−0127 and Cyg X-1 as well as for GX 339−4 in a lower-flux state than observed in 2004[3]. The lags at higher frequencies cannot be compared with the GX 339−4 2004 data due to larger error bars caused by the shorter exposure times of these observations (40 ks, 16 ks and 31 ks for Swift J1753.5−0127, Cyg X-1 and GX 339−4 2009 respectively). However, in the 0.125-0.5 Hz range the other systems, as well as the GX 339−4 2009 data, show a downturn in lag versus energy in soft X-rays which is similar to that seen in the GX 339−4 2004 data.

A.4 Discussion

By extending the measurement of energy-dependent lags to soft energies, we have shown that on time-scales of seconds or longer, disc blackbody variations lead variations in the power-law component. The natural implication of this result is that the disc variations drive

[3]Swift J1753.5−0127 was observed on 2006 March 24, GX 339−4 and Cyg X-1 were observed on 2009 March 26 and 2009 December 2 respectively.

the power-law variations, so we now consider the physical origin of this causal connection. The medium-soft lags of 0.1 s or longer seen in variations on time-scales of seconds correspond to light-travel times of thousands of gravitational radii. Coronal Comptonisation models have been invoked to explain the log-linear form of the lags seen at harder energies (Payne, 1980; Kazanas and Hua, 1999), but they struggle to explain even the more moderate hard-medium lags seen at low frequencies without invoking an unfeasibly large Comptonising region (Nowak et al., 1999). A recent more energetically-feasible model invokes Comptonisation by the persistent hard state jet to produce the hard lags (Reig et al., 2003; Kylafis et al., 2008). In this model, seed photons are fed into the base of the jet (e.g. from the disc), and the anisotropy of Compton scattering along the jet axis means that upscattered photons can traverse large distances in the jet before escaping to the observer, to produce large lags relative to the photons undergoing fewer upscatterings. However, in such a model we would expect the medium-soft lags to be reduced relative to the hard-medium lags, not increased as observed here, since the direct seed photons from the disc at soft energies will dilute the log-linear lags introduced by the upscattered continuum. The medium-soft lag would only increase relative to the hard-medium lag if the medium-soft lag represents the actual travel time from the seed photon source to the scattering electron population, i.e. the seed photons must first interact with the jet not at its base but thousands of gravitational radii above the disc. This geometry raises significant problems for explaining the energetics and variability of the emission (since the jet at that height would subtend only a small solid angle as seen from the disc), not to mention the shape and strength of the iron $K\alpha$ line, which due to geometric and beaming effects would be much narrower and weaker than observed, if the disc were illuminated by the jet from such a great height (e.g. Markoff and Nowak 2004).

A likely explanation of the large medium-soft lags observed on time-scales of seconds is that the lags are associated with the generation of and propagation of accretion fluctuations within the inner regions of the disc before they reach the corona (Lyubarskii, 1997; Kotov et al., 2001; Arévalo and Uttley, 2006). In fact, the observed lag behaviour on these and sub-second time-scales is exactly that expected from our earlier disc-variability interpretation of the covariance spectrum (Wilkinson and Uttley, 2009). In the present work, we have used causal information to greatly strengthen that interpretation, so that accretion instabilities can be firmly identified as the source of X-ray variability, at least on time-scales of seconds. Furthermore, the unstable accretion flow is now shown to be the standard disc, with all the attendant physical implications (see discussion in Wilkinson and Uttley 2009), and not a hot optically thin flow.

The disc propagation model can explain the medium-soft lags, however the hard-medium lags - produced where the disc blackbody does not contribute to the X-ray spectrum - still require explanation. These lags may still be related to Compton upscattering of seed photons as the accretion fluctuations reach the base of the jet, e.g. in a hybrid model of disc propagation and the model of Reig et al. (2003). However, given the common $\nu^{-0.7}$ frequency-dependence of the low-frequency lags which seems to be independent of the energy-bands

chosen, it seems likely that the lags at hard energies are directly related to the same underlying propagation mechanism which produces the medium-soft lags. For example, lags at harder energies can be simply produced if the mass-accretion fluctuations in the disc preferentially generate softer and then harder power-law emission as they propagate inwards. Such effects could be produced if there is a soft coronal component above the disc while the hard emission is produced centrally (e.g. Kotov et al. 2001; Arévalo and Uttley 2006).

The switch in lag behaviour on short time-scales suggests that the disc lags the power-law variations by a few ms on these time-scales, consistent with the light-travel lags expected from a power-law component separated from the disc by only tens of gravitational radii at most. This lag signature represents the first evidence in BHXRBs for a disc thermal reverberation lag. This is the time-lag due to the light travel-time from the central power-law emitting region to the disc where the hard X-rays are reprocessed into blackbody emission. In order to use these lags to map the disc, e.g. to constrain the inner radius, we would need to make assumptions about the geometry of the power-law emitting region and the emissivity profile of the disc blackbody. This modelling effort is beyond the scope of this work.

Finally, it is important to note that although the lags observed for variations on time-scales of less than a second appear to be caused by X-ray heating of the disc, this does not necessarily imply that the variability on these shorter time-scales is not also generated in the disc. An accretion instability at small disc radii will modulate only disc emission inside that radius (e.g. Arévalo and Uttley 2006), so the intrinsic variability in disc emission will be small and likely to be dominated by X-ray heating effects when the mass fluctuations reach the central power-law emitting region.

Acknowledgments

We would like to thank the anonymous referee for valuable comments. PU is supported by an STFC Advanced Fellowship and TW is supported by an STFC postgraduate studentship grant. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number ITN 215212 "Black Hole Universe". This work was partly funded by the Bundesministerium für Wirtschaft and Technologie through Deutsches Zentrum für Luft- und Raumfahrt grants 50 OR 0701 and 50 OR 0808. This work is based on observations obtained with *XMM-Newton*, an ESA science mission with instruments and contributions directly funded by ESA Member States and NASA.

APPENDIX B

Reflection transfer functions: further considerations

In Chapter 2 it was shown that, under certain assumptions, the total luminosity that is observed in the presence of a source of direct continuum and a reflector is:

$$L_{\rm tot}(E,t|\alpha,i) = L_{\rm PL}(E) \times \left(\delta(t) + a(E)\frac{\Omega_{\rm eff}}{2\pi}\sum_l \kappa_{\alpha,l}(E,i)\delta(t-\tau_l)\right) \quad . \tag{B.1}$$

The above equation can be rewritten as:

$$L_{\rm tot}(E,t|\alpha,i) = L_{\rm PL}(E)\,T(E,t|\alpha,i) \quad , \tag{B.2}$$

where

$$\begin{aligned} \tilde{L}_{\rm tot}(E,\nu,i) &= \mathscr{F}\left\{L_{\rm PL}(E)\,T(E,t|\alpha,i)\right\} \\ &= L_{\rm PL}(E)\,\tilde{T}(E,\nu|\alpha,i) \quad . \end{aligned} \tag{B.3}$$

Equation B.3 is the sum of two terms, a frequency-independent term (because $\mathscr{F}[\delta(t)] = 1$) and a term that depends on both energy and Fourier-frequency:

$$\begin{aligned} \tilde{L}_{\rm tot}(E,\nu,i) &= L_{\rm PL}(E)\left(1 + a(E)\frac{\Omega_{\rm eff}}{2\pi}R(\nu)e^{i\phi_{\rm r}(\nu)}\right) \\ &= L_{\rm PL}(E) + a(E)L_{\rm PL}(E)R(\nu)e^{i\phi_{\rm r}(\nu)} \quad . \end{aligned} \tag{B.4}$$

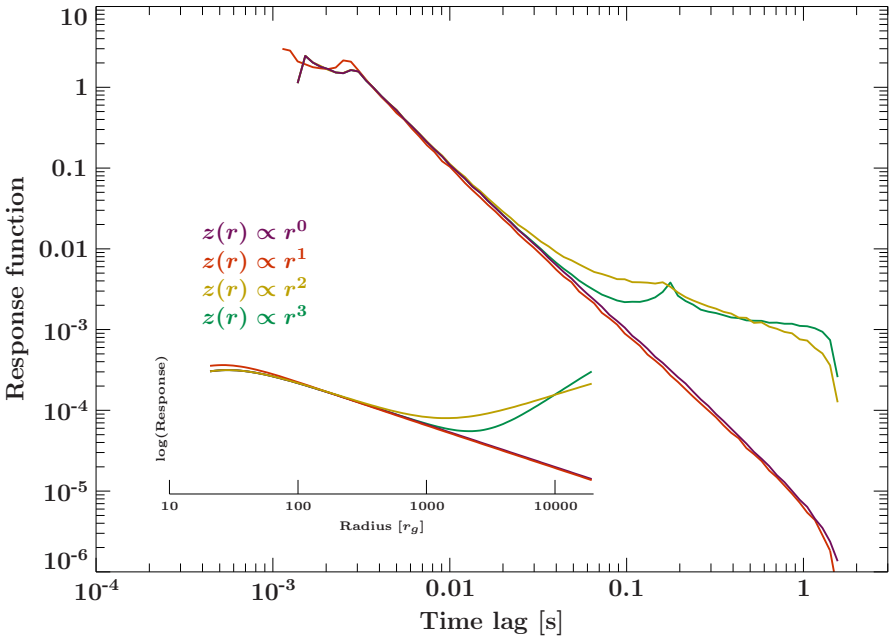

Figure B.1: Examples of time-dependent response of the disc to an incident pulse of emission produced at a certain height above the disc ($R_{\rm in} = 20r_{\rm g}$, $R_{\rm out} = 2\times10^4 r_{\rm g}$, $H_{\rm out}/R_{\rm out} = 0.2$, $h_{\rm src} = 10r_{\rm g}$, $i = 45$ degrees, for a 10 $M_\odot$ black hole). As the disc flaring increases, a secondary peak appears at ~ 0.2 s, corresponding to large-scale reflection produced where the combination of the solid angle of the disc as seen by the observer, as well as the disc as seen by the source of incident continuum emission reaches a maximum. The inset shows the radial distribution of the reflected emission, which is proportional to the solid angle subtended by each cell of the disc as seen by the source. This does not take into account corrections due to inclination angle, which modify the azimuthal structure of the reflected component.

The function $R(\nu)e^{i\phi_{\rm r}(\nu)}$ equals the Fourier transform of the reflection geometry-dependent term $\frac{\Omega_{\rm eff}}{2\pi}\sum_l \kappa_{\alpha,l}(E,i)\delta(t-\tau_l)$ in Eqn.B.1. Both terms are, by assumption, independent of energy, all the energy-dependent information being subsumed in $a(E)$ (the albedo) and the strength of the illuminating continuum $L_{\rm PL}(E)$. The shape of $R(\nu)e^{\phi_{\rm r}(\nu)}$ depends on the geometry of the reflecting accretion disc, as well as on the height of the source of illuminating power-law photons above the disc. Fig. B.1 shows examples of the disc time-dependent response for disc flaring indices ranging from 0 (flat disc) to 3.

The direct continuum term in Eqn. B.4, which is independent of Fourier-frequency, has a null projection on the imaginary axis, hence the frequent statement that this component has 0 intrinsic lag (or, more appropriately, 0 intrinsic phase). The effect of adding the direct continuum emission to the frequency-dependent reflected emission (that carries frequency-dependent phase information that depends on the geometry of the reflecting medium as well as on projection effects) is an effective reduction of the observed phase-lags that are obtained from the cross-spectrum, i.e. an increase of the real part of the cross-spectrum does not carry a commensurate increase of its imaginary part. This can be seen by taking two energies, 1 and 2 (where $E_2 > E_1$), as the cross-spectrum will equal: $C(\nu) = L_1^*(\nu)L_2(\nu) = L_{d,1}L_{d,2}\left(1 + a_1R(\nu)e^{-i\phi_{\rm r}(\nu)} + a_2R(\nu)e^{i\phi_{\rm r}} + a_1a_2R^2(\nu)\right)$.

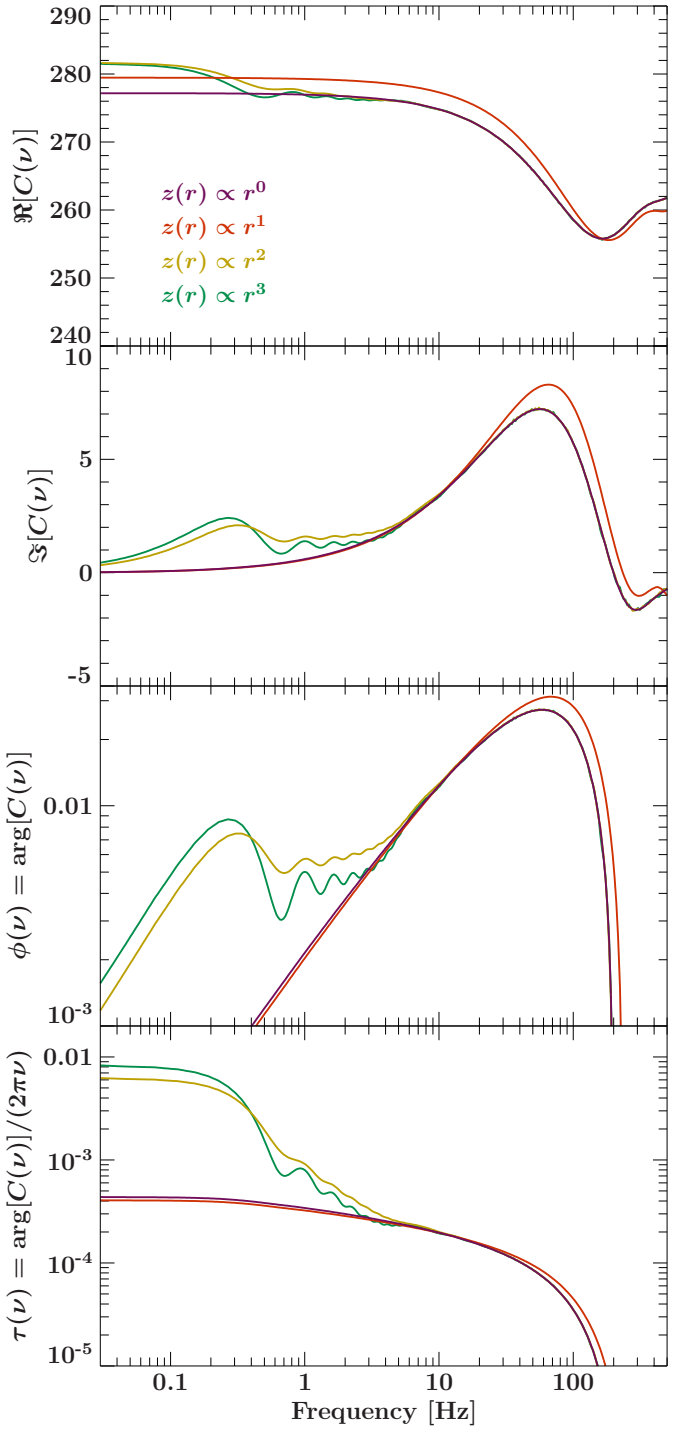

Figure B.2: Real and imaginary components of the cross-spectrum resulting from reflection using the parameters in Fig. B.1, as well as the phase- and time-lag resulting from it. The energy bands used here are 2.0 – 3.5 keV and 4.0 – 10.0 keV (as in Chapter 2). A positive lag indicates that hard photons lag soft photons.

The corresponding phase-lag $\phi(\nu)$ will be:

$$\tan[\phi(\nu)] = \frac{R(\nu)(a_2 - a_1)\sin\phi_{\rm r}}{1 + a_1 a_2 R^2(\nu) + (a_1 + a_2)R(\nu)\cos\phi_{\rm r}} \quad . \tag{B.5}$$

Note that the above phase-lags are independent of the direct input luminosity. If $a_1 \ll a_2$ (as is the case in Chapter 2),

$$\tan[\phi(\nu)] \approx \frac{\sin\phi_{\rm r}}{\frac{1}{a_2 R(\nu)} + \cos\phi_{\rm r}} \quad . \tag{B.6}$$

The equation above shows that observing the direct incident continuum produces a reduction in the lags, but also that the lags increase as the amount of reflected emission and the albedo in the second energy range increase.

It can be seen from Fig. B.1 that concave discs produce a secondary peak of emission at large lags. This is possible in the case where a range of cells have solid angles with respect to the source and the observer that produce reflected emission that is seen by the observer almost simultaneously. It is important to stress that the cells responsible for this emission need not be contiguous.

The Fourier transform of the disc response functions shown in Fig. B.1 that include the direct continuum term are shown in Fig. B.2.

Between 50 and 100 Hz, the phase-lags reach a maximum in all cases, meaning that it is on these time-scales that most of the reflected component is produced. A feature is however not present in the time-lags –these are of the order of the light-travel path difference between the observed reflected emission and the observed direct emission which, on average, increase slowly as a function of disc radius.

On time-scales much shorter than the the light travel-time from the source of power-law photons to the reflector, the disc reflection does not modulate the observed variability. For this reason, the observed phase-lags approach the phase of the direct continuum for increasing frequencies. However, the imaginary part of the Fourier transform of the incident emission is zero, therefore the phase-lags approach this value.

Below ~ 1 Hz, concave geometries produce a secondary peak that corresponds to the excess reflection at large lags (and radii) as seen in Fig. B.1. The value of the time-lag at these frequencies approaches the light travel-time time-lag to the outermost regions. The amount of reflection at these large radii has to be comparably significant to be seen, otherwise these lags are completely diluted by the direct continuum (see Eqns. B.5 and B.6).

Bibliography

Abramowitz, M. and Stegun, I. A. (1965). Handbook of mathematical functions. 1965. *Applied Mathematics Series*, 55.

Arévalo, P. and Uttley, P. (2006). Investigating a fluctuating-accretion model for the spectral-timing properties of accreting black hole systems. *MNRAS*, 367:801–814.

Axelsson, M., Borgonovo, L., and Larsson, S. (2005). Evolution of the 0.01-25 Hz power spectral components in Cygnus X-1. *A&A*, 438:999–1012.

Balbus, S. A. and Hawley, J. F. (1991). A powerful local shear instability in weakly magnetized disks. I - Linear analysis. II - Nonlinear evolution. *ApJ*, 376:214–233.

Ballantyne, D. R., Iwasawa, K., and Fabian, A. C. (2001). Evidence for ionized accretion discs in five narrow-line Seyfert 1 galaxies. *MNRAS*, 323:506–516.

Bartlett, M. S. (1948). Smoothing periodograms from Time-Series with continuous spectra. *Nature*, 161(4096):686–687.

Belloni, T., Homan, J., Casella, P., van der Klis, M., Nespoli, E., Lewin, W. H. G., Miller, J. M., and Méndez, M. (2005). The evolution of the timing properties of the black-hole transient GX 339-4 during its 2002/2003 outburst. *A&A*, 440:207–222.

Belloni, T., Psaltis, D., and van der Klis, M. (2002). A Unified Description of the Timing Features of Accreting X-Ray Binaries. *ApJ*, 572:392–406.

Bendat, J. and Piersol, A. (2010). *Random Data: Analysis and Measurement Procedures*. Wiley Series in Probability and Statistics. John Wiley & Sons.

Cardelli, J. A., Clayton, G. C., and Mathis, J. S. (1989). The relationship between IR, optical, and UV extinction. In L. J. Allamandola & A. G. G. M. Tielens, editor, *Interstellar Dust*, volume 135 of *IAU Symposium*, page 5P.

Chakrabarty, D. (1996). *Hard X-Ray Detection and Timing of Accretion-Powered Pulsars with BATSE*. PhD thesis, Division of Physics, Mathematics, and Astronomy, California Insititute of Technology, Pasadena, CA.

Chen, W., Shrader, C. R., and Livio, M. (1997). The Properties of X-Ray and Optical Light Curves of X-Ray Novae. *ApJ*, 491:312.

Chiang, C. Y., Done, C., Still, M., and Godet, O. (2010). An additional soft X-ray component in the dim low/hard state of black hole binaries. *MNRAS*, 403:1102–1112.

Churazov, E., Gilfanov, M., and Revnivtsev, M. (2001). Soft state of Cygnus X-1: stable disc and unstable corona. *MNRAS*, 321:759–766.

Crary, D. J., Finger, M. H., Kouveliotou, C., van der Hooft, F., van der Klis, M., Lewin, W. H. G., and van Paradijs, J. (1998). Hard X-Ray Lags in Cygnus X-1. *ApJ*, 493:L71.

Cui, W., Zhang, S. N., Focke, W., and Swank, J. H. (1997). Temporal Properties of Cygnus X-1 during the Spectral Transitions. *ApJ*, 484:383.

de Marco, B., Ponti, G., Uttley, P., Cappi, M., Dadina, M., Fabian, A. C., and Miniutti, G. (2011). PG 1211+143: probing high-frequency lags in a high-mass active galactic nucleus. *MNRAS*, 417:L98–L102.

Doi, K. (1978). Physical reality of shot noise model for the short-term variability of CYG X-1. *Nature*, 275:197.

Done, C. (2010). Observational characteristics of accretion onto black holes. *ArXiv e-prints*.

Done, C. and Diaz Trigo, M. (2010). A re-analysis of the iron line in the XMM-Newton data from the low/hard state in GX339-4. *MNRAS*, 407:2287–2296.

Done, C., Gierliński, M., and Kubota, A. (2007). Modelling the behaviour of accretion flows in X-ray binaries. Everything you always wanted to know about accretion but were afraid to ask. *A&A Rev.*, 15:1–66.

Droulans, R., Belmont, R., Malzac, J., and Jourdain, E. (2010). Variability and Spectral Modeling of the Hard X-ray Emission of GX 339-4 in a Bright Low/Hard State. *ApJ*, 717:1022–1036.

Dubus, G., Hameury, J.-M., and Lasota, J.-P. (2001). The disc instability model for X-ray transients: Evidence for truncation and irradiation. *A&A*, 373:251–271.

Durant, M., Gandhi, P., Shahbaz, T., Peralta, H. H., and Dhillon, V. S. (2009). Multiwavelength spectral and high time resolution observations of SWIFTJ1753.5-0127: new activity? *MNRAS*, 392:309–324.

Emmanoulopoulos, D., McHardy, I. M., and Papadakis, I. E. (2011). Negative X-ray reverberation time delays from MCG-6-30-15 and Mrk 766. *MNRAS*, 416:L94–L98.

ESA: XMM-Newton SOC (2012). XMM-Newton Users Handbook, Issue 2.10.

Esin, A. A., McClintock, J. E., and Narayan, R. (1997). Advection-dominated Accretion and the Spectral States of Black Hole X-Ray Binaries: Application to Nova MUSCAE 1991. *ApJ*, 489:865.

Fabian, A. C., Rees, M. J., Stella, L., and White, N. E. (1989). X-ray fluorescence from the inner disc in Cygnus X-1. *MNRAS*, 238:729–736.

Fabian, A. C. and Ross, R. R. (2010). X-ray Reflection. *Space Sci. Rev.*, 157:167–176.

Fabian, A. C., Zoghbi, A., Ross, R. R., Uttley, P., Gallo, L. C., Brandt, W. N., Blustin, A. J., Boller, T., Caballero-Garcia, M. D., Larsson, J., Miller, J. M., Miniutti, G., Ponti, G., Reis, R. C., Reynolds, C. S., Tanaka, Y., and Young, A. J. (2009). Broad line emission from iron K- and L-shell transitions in the active galaxy 1H0707-495. *Nature*, 459:540–542.

Fender, R. P., Belloni, T. M., and Gallo, E. (2004). Towards a unified model for black hole X-ray binary jets. *MNRAS*, 355:1105–1118.

Frank, J., King, A., and Raine, D. (2002). *Accretion power in astrophysics*. Cambridge University Press.

Galeev, A. A., Rosner, R., and Vaiana, G. S. (1979). Structured coronae of accretion disks. *ApJ*, 229:318–326.

George, I. M. and Fabian, A. C. (1991). X-ray reflection from cold matter in active galactic nuclei and X-ray binaries. *MNRAS*, 249:352–367.

Giacconi, R., Gursky, H., Kellogg, E., Schreier, E., and Tananbaum, H. (1971). Discovery of Periodic X-Ray Pulsations in Centaurus X-3 from UHURU. *ApJ*, 167:L67.

Giannios, D. (2005). Spectra of black-hole binaries in the low/hard state: From radio to X-rays. *A&A*, 437:1007–1015.

Gierliński, M., Done, C., and Page, K. (2008). X-ray irradiation in XTE J1817-330 and the inner radius of the truncated disc in the hard state. *MNRAS*, 388:753–760.

Gilfanov, M., Churazov, E., and Revnivtsev, M. (1999). Reflection and noise in Cygnus X-1. *A&A*, 352:182–188.

Goodman, J. W. (1985). *Statistical optics*. New York, Wiley-Interscience, 1985, 567 p.

Groth, E. J. (1975). Probability distributions related to power spectra. *ApJS*, 29:285–302.

Guainazzi, M. (2011). Xmm-newton calibration technical note 18. Technical report, XMM-Newton Science Operations Centre.

Hanke, M. (2011). *Probing the Environment of Accreting Compact Objects*. PhD thesis, Dr. Karl Remeis-Sternwarte, Astronomisches Institut der Universität Erlangen-Nürnberg.

Homan, J., Wijnands, R., van der Klis, M., Belloni, T., van Paradijs, J., Klein-Wolt, M., Fender, R., and Méndez, M. (2001). Correlated X-Ray Spectral and Timing Behavior of the Black Hole Candidate XTE J1550-564: A New Interpretation of Black Hole States. *ApJS*, 132:377–402.

Houck, J. C. and Denicola, L. A. (2000). ISIS: An Interactive Spectral Interpretation System for High Resolution X-Ray Spectroscopy. In Manset, N., Veillet, C., and Crabtree, D., editors, *Astronomical Data Analysis Software and Systems IX*, volume 216 of *Astronomical Society of the Pacific Conference Series*, page 591.

Ingram, A. and Done, C. (2012). Modelling variability in black hole binaries: linking simulations to observations. *MNRAS*, 419:2369–2378.

Johnson, S. and Frigo, M. (2007). A modified split-radix fft with fewer arithmetic operations. *Signal Processing, IEEE Transactions on*, 55(1):111–119.

Kazanas, D. and Hua, X.-M. (1999). Modeling the Time Variability of Accreting Compact Sources. *ApJ*, 519:750–761.

Kazanas, D., Hua, X.-M., and Titarchuk, L. (1997). Temporal and Spectral Properties of Comptonized Radiation and Its Applications. *ApJ*, 480:735.

King, A. (2012). Accretion Disc Theory since Shakura and Sunyaev. *ArXiv e-prints*.

Kolb, U. (2010). *Extreme environment astrophysics*. Cambridge University Press.

Kolehmainen, M. and Done, C. (2010). Limits on spin determination from disc spectral fitting in GX 339-4. *MNRAS*, 406:2206–2212.

Kotov, O., Churazov, E., and Gilfanov, M. (2001). On the X-ray time-lags in the black hole candidates. *MNRAS*, 327:799–807.

Kubota, A., Tanaka, Y., Makishima, K., Ueda, Y., Dotani, T., Inoue, H., and Yamaoka, K. (1998). Evidence for a Black Hole in the X-Ray Transient GRS 1009-45. *PASJ*, 50:667–673.

Kuster, M., Kendziorra, E., Benlloch, S., Becker, W., Lammers, U., Vacanti, G., and Serpell, E. (2002). Timing with the EPIC pn Camera of XMM-Newton. *ArXiv Astrophysics e-prints*.

Kuulkers, E., Kouveliotou, C., Belloni, T., Cadolle Bel, M., Diaz Trigo, M., Homan, J., Ibarra, A., Kennea, J. A., Munoz-Darias, T., Ness, J.-U., Parmar, A. N., Pollock, A. M. T., van den Heuvel, E. P. J., and van der Horst, A. J. (2012). MAXI J1659-152: The shortest orbital period black-hole transient in outburst. *ArXiv e-prints*.

Kylafis, N. D., Papadakis, I. E., Reig, P., Giannios, D., and Pooley, G. G. (2008). A jet model for Galactic black-hole X-ray sources: some constraining correlations. *A&A*, 489:481–487.

Lasota, J.-P. (2001). The disc instability model of dwarf novae and low-mass X-ray binary transients. *New A Rev.*, 45:449–508.

Lightman, A. P. and Eardley, D. M. (1974). Black Holes in Binary Systems: Instability of Disk Accretion. *ApJ*, 187:L1.

Lightman, A. P. and White, T. R. (1988). Effects of cold matter in active galactic nuclei - A broad hump in the X-ray spectra. *ApJ*, 335:57–66.

Liu, B. F., Taam, R. E., Meyer-Hofmeister, E., and Meyer, F. (2007). The Existence of Inner Cool Disks in the Low/Hard State of Accreting Black Holes. *ApJ*, 671:695–705.

Lynden-Bell, D. (1969). Galactic Nuclei as Collapsed Old Quasars. *Nature*, 223:690–694.

Lyubarskii, Y. E. (1997). Flicker noise in accretion discs. *MNRAS*, 292:679.

Maccarone, T. J., Coppi, P. S., and Poutanen, J. (2000). Time Domain Analysis of Variability in Cygnus X-1: Constraints on the Emission Models. *ApJ*, 537:L107–L110.

Maitra, D., Markoff, S., Brocksopp, C., Noble, M., Nowak, M., and Wilms, J. (2009). Constraining jet/disc geometry and radiative processes in stellar black holes XTE J1118+480 and GX 339-4. *MNRAS*, 398:1638–1650.

Malzac, J. and Belmont, R. (2009). The synchrotron boiler and the spectral states of black hole binaries. *MNRAS*, 392:570–589.

Markoff, S. and Nowak, M. A. (2004). Constraining X-Ray Binary Jet Models via Reflection. *ApJ*, 609:972–976.

Markoff, S., Nowak, M. A., and Wilms, J. (2005). Going with the Flow: Can the Base of Jets Subsume the Role of Compact Accretion Disk Coronae? *ApJ*, 635:1203–1216.

Masai, K. (1997). Effect of iron ionization balance on X-ray spectral analysis. *A&A*, 324:410–416.

McClintock, J. E., Shafee, R., Narayan, R., Remillard, R. A., Davis, S. W., and Li, L.-X. (2006). The Spin of the Near-Extreme Kerr Black Hole GRS 1915+105. *ApJ*, 652:518–539.

Miller, J. M., Homan, J., and Miniutti, G. (2006). A Prominent Accretion Disk in the Low-Hard State of the Black Hole Candidate SWIFT J1753.5-0127. *ApJ*, 652:L113–L116.

Miyamoto, S. and Kitamoto, S. (1989). X-ray time variations from Cygnus X-1 and implications for the accretion process. *Nature*, 342:773.

Miyamoto, S., Kitamoto, S., Iga, S., Negoro, H., and Terada, K. (1992). Canonical time variations of X-rays from black hole candidates in the low-intensity state. *ApJ*, 391:L21–L24.

Miyamoto, S., Kitamoto, S., Mitsuda, K., and Dotani, T. (1988). Delayed hard X-rays from Cygnus X-1. *Nature*, 336:450–452.

Motta, S., Belloni, T., and Homan, J. (2009). The evolution of the high-energy cut-off in the X-ray spectrum of GX 339-4 across a hard-to-soft transition. *MNRAS*, 400:1603–1612.

Muñoz-Darias, T., Motta, S., and Belloni, T. M. (2011). Fast variability as a tracer of accretion regimes in black hole transients. *MNRAS*, 410:679–684.

Narayan, R. and Yi, I. (1994). Advection-dominated accretion: A self-similar solution. *ApJ*, 428:L13–L16.

Novikov, I. D. and Thorne, K. S. (1973). Astrophysics of black holes. In DeWitt, C. and DeWitt, B., editors, *Black Holes*, pages 243–450. Gordon and Breach.

Nowak, M. A. (2000). Are there three peaks in the power spectra of GX 339-4 and Cyg X-1? *MNRAS*, 318:361–367.

Nowak, M. A., Vaughan, B. A., Wilms, J., Dove, J. B., and Begelman, M. C. (1999). Rossi X-Ray Timing Explorer Observation of Cygnus X-1. II. Timing Analysis. *ApJ*, 510:874–891.

Palmer, D. M., Barthelmey, S. D., Cummings, J. R., Gehrels, N., Krimm, H. A., Markwardt, C. B., Sakamoto, T., and Tueller, J. (2005). BAT detection of a new gamma ray source SWIFT J1753.5-0127. *The Astronomer's Telegram*, 546:1.

Payne, D. G. (1980). Time-dependent Comptonization - X-ray reverberations. *ApJ*, 237:951–963.

Pottschmidt, K., Wilms, J., Nowak, M. A., Pooley, G. G., Gleissner, T., Heindl, W. A., Smith, D. M., Remillard, R., and Staubert, R. (2003). Long term variability of Cygnus X-1. I. X-ray spectral-temporal correlations in the hard state. *A&A*, 407:1039–1058.

Poutanen, J. (2001). Understanding spectral variability and time lags in accreting black holes. *Advances in Space Research*, 28:267–280.

Poutanen, J. (2002). Impact of reverberation in flared accretion discs on temporal characteristics of X-ray binaries. *MNRAS*, 332:257–270.

Poutanen, J. and Fabian, A. C. (1999). Spectral evolution of magnetic flares and time lags in accreting black hole sources. *MNRAS*, 306:L31–L37.

Press, W. (1992). *Numerical recipes in FORTRAN: the art of scientific computing*, volume 1. Cambridge University Press.

Psaltis, D., Belloni, T., and van der Klis, M. (1999). Correlations in Quasi-periodic Oscillation and Noise Frequencies among Neutron Star and Black Hole X-Ray Binaries. *ApJ*, 520:262–270.

Reig, P., Kylafis, N. D., and Giannios, D. (2003). Energy and time-lag spectra of galactic black-hole X-ray sources in the low/hard state. *A&A*, 403:L15–L18.

Reis, R. C., Fabian, A. C., and Miller, J. M. (2010). Black hole accretion discs in the canonical low-hard state. *MNRAS*, 402:836–854.

Remillard, R. A. and McClintock, J. E. (2006). X-Ray Properties of Black-Hole Binaries. *ARA&A*, 44:49–92.

Revnivtsev, M., Gilfanov, M., and Churazov, E. (1999). The frequency resolved spectroscopy of CYG X-1: fast variability of the Fe K_α line. *A&A*, 347:L23–L26.

Reynolds, M. T., Miller, J. M., Homan, J., and Miniutti, G. (2010). Suzaku Broadband Spectroscopy of Swift J1753.5-0127 in the Low-Hard State. *ApJ*, 709:358–368.

Rodriguez, J., Durouchoux, P., Mirabel, I. F., Ueda, Y., Tagger, M., and Yamaoka, K. (2002). Energy dependence of a low frequency QPO in GRS 1915+105. *A&A*, 386:271–279.

Rybicki, G. B. and Lightman, A. P. (1979). *Radiative processes in astrophysics*. New York, Wiley-Interscience, 1979. 393 p.

Shakura, N. I. and Sunyaev, R. A. (1973). Black holes in binary systems. Observational appearance. *A&A*, 24:337–355.

Shakura, N. I. and Sunyaev, R. A. (1976). A theory of the instability of disk accretion on to black holes and the variability of binary X-ray sources, galactic nuclei and quasars. *MNRAS*, 175:613–632.

Shiokawa, H., Dolence, J. C., Gammie, C. F., and Noble, S. C. (2012). Global General Relativistic Magnetohydrodynamic Simulations of Black Hole Accretion Flows: A Convergence Study. *ApJ*, 744:187.

Sivia, D. and Skilling, J. (2006). *Data analysis: a Bayesian tutorial*. Oxford University Press, USA.

Smak, J. (1984). Accretion in cataclysmic binaries. IV - Accretion disks in dwarf novae. *Acta Astron.*, 34:161–189.

Sobolewska, M. A. and Życki, P. T. (2006). Spectral and Fourier analyses of X-ray quasi-periodic oscillations in accreting black holes. *MNRAS*, 370:405–414.

Sorathia, K. A., Reynolds, C. S., Stone, J. M., and Beckwith, K. (2012). Global Simulations of Accretion Disks. I. Convergence and Comparisons with Local Models. *ApJ*, 749:189.

Strüder, L., Briel, U., Dennerl, K., Hartmann, R., Kendziorra, E., Meidinger, N., Pfeffermann, E., Reppin, C., Aschenbach, B., Bornemann, W., Bräuninger, H., Burkert, W., Elender, M., Freyberg, M., Haberl, F., Hartner, G., Heuschmann, F., Hippmann, H., Kastelic, E., Kemmer, S., Kettenring, G., Kink, W., Krause, N., Müller, S., Oppitz, A., Pietsch, W., Popp, M., Predehl, P., Read, A., Stephan, K. H., Stötter, D., Trümper, J., Holl, P., Kemmer, J., Soltau, H., Stötter, R., Weber, U., Weichert, U., von Zanthier, C., Carathanassis, D., Lutz, G., Richter, R. H., Solc, P., Böttcher, H., Kuster, M., Staubert, R., Abbey, A., Holland, A., Turner, M., Balasini, M., Bignami, G. F., La Palombara, N., Villa, G., Buttler, W., Gianini, F., Lainé, R., Lumb, D., and Dhez, P. (2001). The European Photon Imaging Camera on XMM-Newton: The pn-CCD camera. *A&A*, 365:L18–L26.

Terrell, Jr., N. J. (1972). Shot-Noise Character of Cygnus X-1 Pulsations. *ApJ*, 174:L35.

Thomas, J. (1969). *An introduction to statistical communication theory*. Wiley New York.

Thorne, K. S. and Price, R. H. (1975). Cygnus X-1 - an interpretation of the spectrum and its variability. *ApJ*, 195:L101–L105.

Timmer, J. and Koenig, M. (1995). On generating power law noise. *A&A*, 300:707.

Tomsick, J. A., Yamaoka, K., Corbel, S., Kaaret, P., Kalemci, E., and Migliari, S. (2009). Truncation of the Inner Accretion Disk Around a Black Hole at Low Luminosity. *ApJ*, 707:L87–L91.

Uttley, P. and McHardy, I. M. (2001). The flux-dependent amplitude of broadband noise variability in X-ray binaries and active galaxies. *MNRAS*, 323:L26–L30.

Uttley, P., McHardy, I. M., and Vaughan, S. (2005). Non-linear X-ray variability in X-ray binaries and active galaxies. *MNRAS*, 359:345–362.

Uttley, P., Wilkinson, T., Cassatella, P., Wilms, J., Pottschmidt, K., Hanke, M., and Böck, M. (2011). The causal connection between disc and power-law variability in hard state black hole X-ray binaries. *MNRAS*, 414:L60–L64.

Vaughan, B., van der Klis, M., Lewin, W. H. G., Wijers, R. A. M. J., van Paradijs, J., Dotani, T., and Mitsuda, K. (1994). The time-delay spectrum of GX 5-1 in its horizontal branch. *ApJ*, 421:738–752.

Vaughan, B. A. and Nowak, M. A. (1997). X-Ray Variability Coherence: How to Compute It, What It Means, and How It Constrains Models of GX 339-4 and Cygnus X-1. *ApJ*, 474:L43.

Veledina, A., Poutanen, J., and Vurm, I. (2011). A Synchrotron Self-Compton-Disk Reprocessing Model for Optical/X-Ray Correlation in Black Hole X-Ray Binaries. *ApJ*, 737:L17.

Wijnands, R. and van der Klis, M. (1999). The Broadband Power Spectra of X-Ray Binaries. *ApJ*, 514:939–944.

Wilkinson, T. (2011). *Using novel spectral-timing techniques to probe the rapid variability of X-Ray Binaries*. PhD thesis, University of Southampton.

Wilkinson, T. and Uttley, P. (2009). Accretion disc variability in the hard state of black hole X-ray binaries. *MNRAS*, 397:666–676.

Wilms, J., Juett, A., Schulz, N. S., and A., N. M. (in preparation). `http://pulsar.sternwarte.uni-erlangen.de/wilms/research/tbabs/`.

Zdziarski, A. A., Gierliński, M., Mikołajewska, J., Wardziński, G., Smith, D. M., Harmon, B. A., and Kitamoto, S. (2004). GX 339-4: the distance, state transitions, hysteresis and spectral correlations. *MNRAS*, 351:791–807.

Zdziarski, A. A., Poutanen, J., Mikolajewska, J., Gierlinski, M., Ebisawa, K., and Johnson, W. N. (1998). Broad-band X-ray/gamma-ray spectra and binary parameters of GX 339-4 and their astrophysical implications. *MNRAS*, 301:435–450.

Zoghbi, A., Fabian, A. C., Uttley, P., Miniutti, G., Gallo, L. C., Reynolds, C. S., Miller, J. M., and Ponti, G. (2010). Broad iron L line and X-ray reverberation in 1H0707-495. *MNRAS*, 401:2419–2432.

Zurita, C., Durant, M., Torres, M. A. P., Shahbaz, T., Casares, J., and Steeghs, D. (2008). Swift J1753.5-0127: The Black Hole Candidate with the Shortest Orbital Period. *ApJ*, 681:1458–1463.