Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

Mon. Not. R. Astron. Soc. **419,** 2369–2378 (2012) doi:10.1111/j.1365-2966.2011.19885.x

Modelling variability in black hole binaries: linking simulations to observations

Adam Ingram★ and Chris Done

Department of Physics, University of Durham, South Road, Durham DH1 3LE

Accepted 2011 September 23. Received 2011 August 31; in original form 2011 July 29

ABSTRACT

Black hole accretion flows show rapid X-ray variability. The power spectral density (PSD) of this is typically fit by a phenomenological model of multiple Lorentzians for both the broad-band noise and quasi-periodic oscillations (QPOs). Our previous paper developed the first physical model for the PSD and fit this to observational data. This was based on the same truncated disc/hot inner flow geometry which can explain the correlated properties of the energy spectra. This assumes that the broad-band noise is from propagating fluctuations in mass accretion rate within the hot flow, while the QPO is produced by global Lense–Thirring precession of the same hot flow. Here we develop this model, making some significant improvements. First, we specify that the viscous frequency (equivalently, surface density) in the hot flow has the same form as that measured from numerical simulations of precessing, tilted accretion flows. Secondly, we refine the statistical techniques which we use to fit the model to the data. We re-analyse the PSD from the 1998 rise to outburst of XTE J1550−564 with our new model in order to assess the impact of these changes. We find that the derived outer radii of the hot flow (set by the inner radius of the truncated disc) are rather similar, changing from ∼68 to $13R_g$ throughout the outburst rise. However, the more physical assumptions of our new model also allow us to constrain the scaleheight of the flow. This decreases as the outer radius of the flow decreases, as expected from the spectral evolution. The spectrum steepens in response to the increased cooling as the truncation radius sweeps in, so gas pressure support for the flow decreases. The new model, PROPFLUC, is publicly available within the XSPEC spectral fitting package.

Key words: accretion, accretion discs – X-rays: binaries – X-rays: individual: XTE J1550−564.

1 INTRODUCTION

Black hole binaries (BHBs) have X-ray emission which is variable on a broad range of time-scales. On the longest time-scales (∼weeks), these sources are seen to transition between quiescence, when they are hardly visible above the X-ray background, and outburst, when they are amongst the brightest X-ray objects in the sky. During the rise from quiescence to outburst, the source displays an evolution in spectral state. At the lowest luminosities, the source is in the low/hard state, with a spectral energy distribution (SED) dominated by a hard (photon index $\Gamma < 2$) power-law tail but also including a weak disc component and reflection features. As the average mass accretion rate, $\dot{M}_0$, increases, the disc increases in luminosity, the power law softens ($\Gamma \sim 2$) and the reflection fraction increases (intermediate state). As the luminosity increases further, the SED either becomes almost completely disc dominated (high/soft or ultrasoft state) or displays both a strong disc component and a soft ($\Gamma > 2$) power-law tail (very high state; see e.g. Done, Gierliński & Kubota 2007, hereafter DGK07).

The SED of the low/hard state can be explained by a two-component model whereby the standard cool, optically thick, geometrically thin accretion disc (Shakura & Sunyaev 1973) is truncated at some radius r_o which is greater than the last stable orbit, r_{lso}. Interior to this is a hot, optically thin, geometrically thick accretion flow, perhaps similar to an advection-dominated accretion flow (Narayan & Yi 1995). We observe a fraction of the disc emission directly but some of the photons emitted by the disc are incident upon the flow where they are Compton up-scattered by the hot electrons, thus creating the power-law tail. As the source evolves, the truncation radius decreases which increases both the amount of direct disc emission and the number of seed photons incident on the flow, thus cooling the Comptonizing electrons and softening the power-law index of the flow emission. The spectral evolution throughout the

★E-mail: a.r.ingram@durham.ac.uk

rise to outburst can be described by assuming that r_o decreases as the average mass accretion rate increases until $r_o \approx r_{lso}$ in the high/soft state (Esin, McClintock & Narayan 1997; DGK07). While this truncated disc model has been challenged by direct observation of the disc inner edge at the last stable orbit in the low/hard state (Miller, Homan & Miniutti 2006; Rykoff et al. 2007), this interpretation of the data has itself been challenged (Gierliński, Done & Page 2009; Done & Diaz Trigo 2010). Thus, there is no unambiguous evidence as yet which rules out these models, and all alternative geometries (e.g. Markoff, Nowak & Wilms 2005; Miller et al. 2006) run into other difficulties (DGK07). Hence, we use the truncated disc model as our overall framework, and explore how this can also be used to interpret the rapid variability seen within a single observation of $\sim$2000 s.

A power spectral analysis of this rapid variability reveals quasi-periodic oscillations (QPOs) superimposed on a broad-band noise of variability. The broad-band noise can be described using broad Lorentzians centred at characteristic frequencies f_b and f_h, often referred to as the low- and high-frequency breaks. The QPO (fundamental and higher harmonics) can be described by narrow Lorentzians centred at f_{QPO}, $2f_{QPO}$, etc. (e.g. Belloni, Psaltis & van der Klis 2002). As the SED evolves on long time-scales, so does the corresponding power spectral density (PSD). All the PSD frequencies, f_b, f_{QPO} and (to a lesser extent) f_h, increase with luminosity and are correlated (Psaltis, Belloni & van der Klis 1999; Wijnands & van der Klis 1999; van der Klis 2004; Klein-Wolt & van der Klis 2008; Belloni 2010). However, the amount of power at high frequencies ($>$10 Hz) remains constant despite the increase in f_h (Gierliński, Nikołajuk & Czerny 2008). Although the variability properties have been studied for over 20 years, the underlying physical processes are still poorly understood. Most previous attempts to physically model the variability properties focus on explaining either the QPO or the broad-band noise. We review these briefly below.

1.1 QPO models

There are many proposed QPO mechanisms in the literature, some of which are based on a misalignment between the angular momentum of the black hole and that of the binary system (e.g. Stella & Vietri 1998; Fragile, Mathews & Wilson 2001; Schnittman 2005; Schnittman, Homan & Miller 2006; Ingram, Done & Fragile 2009, hereafter IDF09) and others are associated with wave modes in the accretion flow (Titarchuk & Osherovich 1999; Wagoner, Silbergleit & Ortega-Rodríguez 2001; Cabanac et al. 2010). In IDF09, we outlined a QPO model based on the original relativistic precession model of Stella & Vietri (1998). Here, the QPO frequency is given by the Lense–Thirring precession frequency of a test mass (say, a hotspot; e.g. Schnittman 2005) at the truncation radius. Lense–Thirring precession is a relativistic effect which occurs because of the asymmetric gravitational potential present around a spinning black hole. A test mass orbit will precess around the black hole if it is misaligned with the black hole spin plane as space–time is being dragged around the black hole. As the particle reaches the starting point of its orbit (i.e. $\phi = 0, 2\pi,$ etc.), that point in space–time has rotated some way around the black hole.

Our extension of this model was to replace the test mass with the entire hot accretion flow interior to the disc truncation radius. There is a differential warp across the flow due to the radial dependence of Lense–Thirring precession ($f_{LT} \propto \sim r^{-3}$). Warps are communicated by bending waves which travel at (close to) the sound speed. For a hot flow the sound speed is fast, so the warp can be communicated across the whole hot inner flow on a time-scale which is *shorter* than the precession period on the outer edge. The entire hot flow can then precess as a solid body (though it is still differentially rotating), with a precession frequency given by the surface-density-weighted average of $f_{LT}(r)$. This global precession is confirmed by General Relativistic Magnetohydrodynamic (GRMHD) simulations of a tilted flow (Fragile et al. 2007; Fragile & Meier 2009). A cool, thin disc responds very differently as here the sound speed across the disc is much longer than the precession period. Viscous diffusion then results in a steady state warp of the disc into the plane of the black hole spin at small radii (Bardeen & Petterson 1975; Kumar & Pringle 1985; Fragile et al. 2001), rather than global precession about the black hole spin axis.

In IDF09, we showed how global precession of the entire hot inner flow interior to a stationary truncated disc can give the QPO. The increasing frequency can be produced by the outer radius of the hot flow (set by the inner radius of the truncated disc) moving inwards in the range $50 > r_o > 10$, as also implied by the SED evolution. This model is also attractive as it ties the QPO to the hot flow, so trivially modulates the Comptonized spectrum rather than the disc component, as required by the data (Rodriguez et al. 2004; Sobolewska & Życki 2006).

1.2 Broad-band noise models

The underlying viscosity mechanism in the flow is most likely the magnetorotational instability (MRI; Hawley & Balbus 1991). It is looking increasingly likely that this is also the underlying source of broad-band variability in the flow as it generates large fluctuations in all quantities (e.g. Krolik & Hawley 2002; Beckwith, Hawley & Krolik 2008; Dexter & Fragile 2011). The temporal variability generated by the MRI extends to very high frequencies but the emission is inherently linked to the reservoir of available gravitational energy and therefore should depend on the mass accretion rate, $\dot{M}$. It is commonly assumed that the variability in $\dot{M}$ at a given radius is characterized by the local viscous frequency, $f_{visc}(r)$, where $f_{visc}(r) \propto r^{-3/2}$ in the simplest case. This interpretation implies that $f_b \approx f_{visc}(r_o)$ and $f_h \approx f_{visc}(r_i)$, where r_i is the inner radius of the flow (e.g. Ingram & Done 2010). Although this is a rather crude approximation, it has the very attractive property that the low-frequency noise is tied to the truncation radius but the high-frequency noise is not. This can explain the observed correlation of the low-frequency break in the PSD with the low-frequency QPO, as both are set by r_o, while the high-frequency break is more or less constant as r_i does not change (Gierliński et al. 2008).

Another fundamental property of the data which must be reproduced by any variability model is the σ–flux relation (Uttley & McHardy 2001). This can be measured by splitting the light curve into multiple short segments and finding the average and standard deviation of each segment. After binning, the standard deviation is always seen to be linearly related to the average flux. This shows that the variability is correlated across all time-scales and therefore, in a picture where different temporal frequencies come predominantly from different spatial regions, there must be a causal connection between those regions. This rules out simple shot noise models, where the variability is independent. Instead, the σ–flux relation *can* be reproduced if the fluctuations in $\dot{M}$ generated at a given radius *propagate* inwards towards the black hole as might be intuitively expected in an accretion flow (Lyubarskii 1997; Kotov, Churazov & Gilfanov 2001; Arévalo & Uttley 2006, hereafter AU06)

1.3 Combining the QPO and broad-band noise

In Ingram & Done (2011, hereafter ID11), we explored combining these ideas of both QPO and noise by developing a model with propagating mass accretion rate fluctuations in a precessing flow. The mass accretion rate fluctuations produce the observed band limited noise and the precession frequency modulates the fluctuating light curve to create the QPO. Crucially, these two processes are linked. We use the same geometry (inner and outer radius of the hot flow, surface density of the hot flow) to make both the QPO *and* broad-band noise properties. Precession of the fluctuating flow modulates its observed emission, imprinting the QPO on the broad-band noise, while fluctuations in the flow cause fluctuations in the precession frequency, making a quasi-periodic rather than periodic oscillation.

Here we develop a more advanced version of the model which is in better agreement with the results of GRMHD simulations, together with better statistical techniques to fit these models to the PSD data. We have made this package publicly available within XSPEC (Arnaud, Borkowski & Harrington 1996) as a local model, PROPFLUC.

2 THE MODEL

As in ID11, the model consists of fluctuations in mass accretion rate which propagate towards the black hole [following Lyubarskii (1997), Kotov et al. (2001) and AU06] within a flow that is precessing. Here we develop the model to include a number of improvements which allow us to gain more physical insight from the best-fitting parameters. Most significantly, we change our underlying assumption about the viscous frequency $f_{\rm visc}(r)$. In ID11, we assumed that this was a power law between $r_{\rm i}$ and $r_{\rm o}$, the inner and outer radius of the precessing hot flow. Here we have it to be a smoothly broken power law, with the radius of the break being the bending wave radius, $r_{\rm bw}$, expected from a misaligned flow. The viscous frequency is related to the surface density profile, Σ, via the radial infall velocity $v_r(r)$ as $f_{\rm visc}(r) = -v_r(r)/R$ and mass conservation sets $\dot{M} \propto \Sigma 2\pi r v_r$. Hence, we can use the surface density profiles from the GRMHD simulations to derive $f_{\rm visc}(r)$, which is especially important as the QPO frequency is dependent on $\Sigma(r)$.

We also change the assumed emissivity from ID11, where $\epsilon \propto r^{-\gamma} b(r)$ [where $b(r)$ was an unknown boundary condition], to $\epsilon \propto r^{-\gamma}\Sigma(r)$, i.e. we tie the emission to where the mass is in the flow. We describe the details of the model below, mainly focusing on these improvements made since ID11. Note that, throughout the paper, we use the convention $R = rR_{\rm g}$, where $R_{\rm g} = GM/c^2$ is a gravitational radius and we always assume a $10\,{\rm M}_{\odot}$ black hole.

2.1 Steady state properties

The surface density of the flow sets the QPO frequency by global precession as

$$f_{\rm prec} = \frac{\int_{r_{\rm i}}^{r_{\rm o}} f_{\rm LT} f_{\rm k} \Sigma r^3 {\rm d}r}{\int_{r_{\rm i}}^{r_{\rm o}} f_{\rm k} \Sigma r^3 {\rm d}r}, \quad (1)$$

where $f_{\rm k}$ is the Keplerian orbital frequency and $f_{\rm LT}$ is the point particle Lense–Thirring precession frequency (given by equation 2 in ID11).

We use the GRMHD simulations of tilted flows to guide our description of $\Sigma(r)$ (Fragile et al. 2007; Fragile 2009; Fragile & Meier 2009). These can be well fitted by a smoothly broken power-law function (IDF09)

$$\overline{\Sigma(r_n, t)} = \frac{\Sigma_0 \dot{M}_0}{cR_{\rm g}} \frac{x^{\lambda}}{(1+x^{\kappa})^{(\zeta+\lambda)/\kappa}}, \quad (2)$$

where $x = r/r_{\rm bw}$ is radius normalized to the bending wave radius $r_{\rm bw} = 3(h/r)^{-4/5} a_*^{2/5}$, Σ_0 is a dimensionless normalization constant and $\dot{M}_0$ is the average mass accretion rate which we will assume stays constant over the course of a single observation. This gives $\Sigma \propto r^{\lambda}$ for small r and $\Sigma \propto r^{-\zeta}$ for large r, where κ governs the sharpness of the break. The bending wave radius occurs at radii larger than the last stable orbit because there are additional torques created by the misaligned black hole spin which result in additional stresses, i.e. enhanced angular momentum transport. The material infalls faster, so its surface density drops.

Mass conservation then sets the viscous frequency as

$$f_{\rm visc}(r_n) = \frac{\dot{M}_0}{2\pi R^2 \overline{\Sigma(r,t)}} = \frac{1}{2\pi r_{\rm bw}^2 \Sigma_0} \frac{(1+x^{\kappa})^{(\zeta+\lambda)/\kappa}}{x^{\lambda+2}} \frac{c}{R_{\rm g}}, \quad (3)$$

such that $f_{\rm visc} \propto r^{\zeta-2}$ for large r and $f_{\rm visc} \propto r^{-(\lambda+2)}$ for small r.

2.2 Propagating mass accretion rate fluctuations

As in ID11 (and AU06), we start by splitting the flow up into N annuli of width ${\rm d}r_n$ such that $r_1 = r_{\rm o}$ (the truncation radius) and $r_N = r_{\rm i} + {\rm d}r_n \approx r_{\rm i}$ (the inner radius of the flow). We assume that the power spectrum of variability generated in mass accretion rate at the nth annulus is given by a zero-centred Lorentzian cutting off at the local viscous frequency:

$$|\tilde{\dot{m}}(r_n, f)|^2 \propto \frac{1}{1+[f/f_{\rm visc}(r_n)]^2}, \quad (4)$$

where a tilde denotes a Fourier transform and $f_{\rm visc}(r_n)$ is derived from equation (3).

We use the method of Timmer & Koenig (1995) to generate mass accretion rate fluctuations, $\dot{m}(r_n, t)$, which satisfy equation (4). These are normalized to have a mean of unity and fractional variability $\sigma/I = F_{\rm var}/\sqrt{N_{\rm dec}}$, where, unlike ID11, $F_{\rm var}$ and $N_{\rm dec}$ are the fractional variability and number of annuli per decade in *viscous frequency* rather than radius. These two descriptions are exactly equivalent where $f_{\rm visc}$ is a power-law function of radius as in ID11, as ${\rm d}f/f = {\rm d}r/r$. However, the more physical smoothly broken power-law form for $f_{\rm visc}$ does not retain this property. We choose to parametrize the noise power in terms of ${\rm d}f/f$ (see below).

The mass accretion rate through the outer annulus is given by $\dot{M}(r_1, t) = \dot{M}_0 \dot{m}(r_1, t)$. Variability is generated in every other annulus according to equation (4), but this is also accompanied by the noise from the outer regions of the flow which propagates inwards. Thus, the mass accretion rate at the nth annulus is given by

$$\dot{M}(r_n, t) = \dot{M}(r_{n-1}, t - t_{\rm lag}) \dot{m}(r_n, t), \quad (5)$$

where $t_{\rm lag} = -R_{\rm g}{\rm d}r_n/v_r(r_n) = {\rm d}r_n/[r_n f_{\rm visc}(r_n)]$ is the propagation time across the nth annulus and $v_r(r_n) = -R_{\rm g} r_n f_{\rm visc}(r_n)$ is the infall velocity.

To convert these mass accretion rate fluctuations into a light curve, we assume that the luminosity emitted from the nth annulus is given by

$${\rm d}L(r_n, t) = \eta/2 \dot{M}(r_n, t) c^2 \epsilon(r_n) r_n {\rm d}r_n, \quad (6)$$

where the (dimensionless) emissivity is given by

$$\epsilon(r_n) = \epsilon_0 r_n^{-\gamma} b(r_n), \quad (7)$$

and γ is the emissivity index, η the accretion efficiency, $b(r)$ the boundary condition and ϵ_0 is a normalization constant. In ID11,

we considered two boundary conditions: the 'stress-free' boundary condition $b(r) = 3(1 - \sqrt{r_n/r_i})$ and the 'stressed' boundary condition $b(r) = 1$. Here, however, we make the intuitive and physical assumption that the boundary condition is set by the surface density such that $b(r) \propto \overline{\Sigma(r_n, t)}$, where $\overline{\Sigma(r_n, t)}$ is the *time-averaged* surface density. This allows the model to link the emission with the amount of material in a particular annulus.

The fluctuating mass accretion rate will also have an effect on precession because mass conservation needs to hold on short time-scales as well as long time-scales, which gives $\dot{M}(r_n, t) \propto \Sigma(r_n, t) 2\pi r^2 f_{\rm visc}$. This means that the surface density at time t is given by

$$\Sigma(r_n, t) = \frac{\Sigma_0 \dot{M}(r_n, t)}{cR_{\rm g}} \frac{x^{\lambda}}{(1 + x^{\kappa})^{(\zeta+\lambda)/\kappa}}, \quad (8)$$

which trivially averages to equation (2) on long time-scales. Because the surface density sets the precession frequency (equation 1), we see that the fluctuations in mass accretion rate cause the precession frequency to vary, thus allowing the model to predict a QPO rather than a purely periodic oscillation.

2.3 Surface density profile

In ID11, we parametrized the viscous frequency with a power law. The fiducial model parameters therein gave $f_{\rm visc} = 0.03 r^{-0.5} f_{\rm k}$, corresponding to a surface density profile $\Sigma(r) \propto r^0$ between the inner and outer radii which were set to $r_{\rm i} = 2.5$ and $r_{\rm o} = 50$, respectively. By comparison, the GRMHD titled flow simulations of Fragile & Meier (2009) also give $\Sigma(r) \propto r^0$ at large radii, but then smoothly break at the bending wave radius to a much steeper dependence. The most relevant simulation to this paper is the case with $a_* = 0.5$ as this is likely closest to the spin of XTE J1550−564 (e.g. Davis, Done & Blaes 2006; Steiner et al. 2011). This has surface density parameters (equation 2; see fig. 4 in IDF09) of $r_{\rm bw} = 8.1$ (corresponding to $h/r = 0.21$), $\kappa = 5$, $\lambda = 7.6$ and $\zeta = 0$.

In Figs 1(a) and (b), we plot these two different surface density prescriptions and their resulting viscous frequencies, with the power law shown by the black solid line and the broken power law shown by the red dashed line. In the case of the broken power law, we choose the normalization $\Sigma_0 = 33.3$ to ensure that both assumptions become consistent with one another at large radii. Fig. 1(c) shows the PSD resulting from the two different prescriptions. The new (and more physically realistic) surface density prescription predicts much less noise at high frequencies than the previous model, where the surface density remained constant down to the innermost radii. This distinction is even more striking as our assumption that the MRI noise power is constant per decade in frequency (as opposed to radius) means that the regions with $r \ll r_{\rm bw}$ have more MRI power generated per decade (in radial extent) than those with $r \gg r_{\rm bw}$. Yet even this additional power at small radii is not sufficient to give enough high-frequency power to match that seen in the data.

To retrieve sufficient high-frequency power requires $\lambda = 1$ rather than 7.6 [green dot–dashed lines in Figs 1(a) and (b)]. This gives a more gradual drop-off in surface density, leading to a less severe transition in viscous frequency at the bending wave radius and hence more high-frequency power [green points in Fig 1(c)]. We discuss the physical implications of this in more detail in Section 5. For now, however, we use $\lambda = 1$ for our fiducial model.

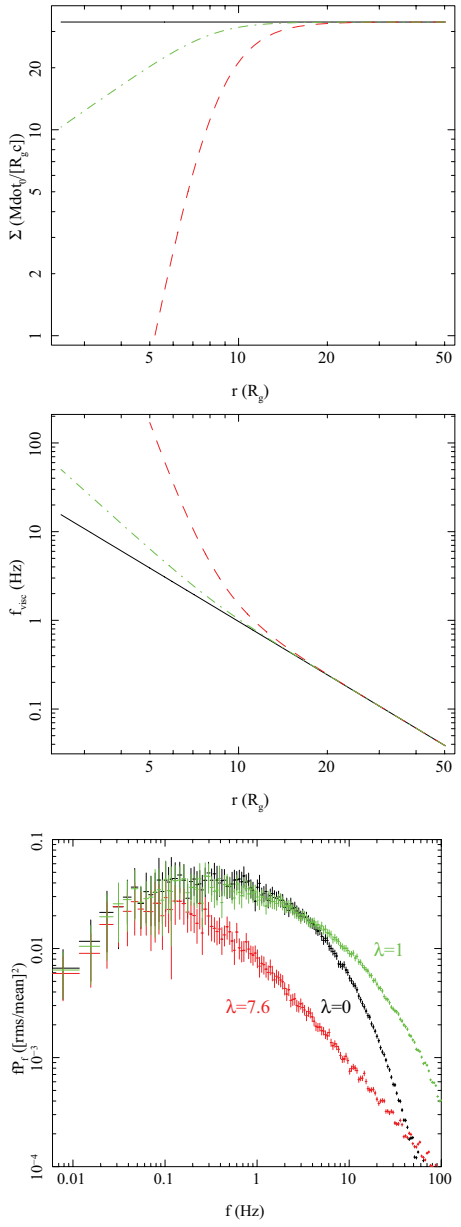

Figure 1. Top (a): surface density as a function of radius for the fiducial model parameters in ID11 (solid black line), simulations of a misaligned accretion flow around a $10\,{\rm M}_{\odot}$ black hole with $a_* = 0.5$ (red dashed line) and the fiducial model parameters we choose for this paper (green dot–dashed line). The red dashed line is calculated using equation (2) with $\lambda = 7.6, \kappa = 5, \zeta = 0$ and $r_{\rm bw} = 8.08$ (the parameters which best fit the simulation data). For the dot–dashed green line, $\lambda = 1$ with all other parameters the same. Middle (b): the viscous frequency as a function of radius resulting from assuming the surface density to be given by the corresponding line in the top panel. Bottom (c): the PSD predicted using the surface density given by corresponding lines in the top panel.

2.4 The fiducial model

Following the discussion in Section 2.1, we use model parameters $\Sigma_0 = 33.3$, $r_{\rm bw} = 8.1$, $\kappa = 5$, $\lambda = 1$ and $\zeta = 0$. We also set $r_{\rm i} = 2$ and $\gamma = 4$ but note that the new assumptions for surface density coupled with the new boundary condition mean the model is now much less sensitive to the parameter $r_{\rm i}$ than its predecessor in ID11. Fig. 2(a) shows a 10-s segment of the light curve created using these assumptions and with $r_{\rm o} = 50$. We use $N_{\rm dec} = 15$ (i.e. 15 annuli per

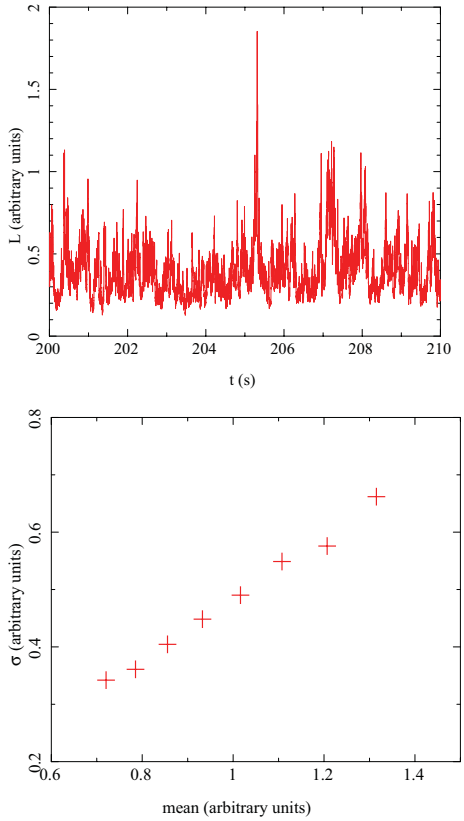

Figure 2. Top (a): a 10-s segment of the light curve calculated using the fiducial model parameters and $r_o = 50$. Bottom (b): the σ–flux relation for the above light curve. We see this is linear as is seen in the data.

decade in viscous frequency) with 2^{22} time-steps, giving a duration of 4096 s (similar to a typical *RXTE* observation) for a time bin of $dt = 9.7 \times 10^{-4}$. Fig. 2(b) confirms that this light curve has the linear σ–flux relation implied by its skewed nature. The PSD of this light curve is represented by the red points in Fig. 1(c).

We calculate the QPO as in ID11, but we briefly summarize this here for completeness. The QPO fundamental frequency is set to the average (over the 4096 s duration) precession frequency calculated from equation (1). In principle, we can calculate the width of the QPO from the fluctuations in frequency which result from fluctuations in surface density. However, these only set a lower limit to the width of the QPO since it can also be broadened by other processes (ID11), so we leave this as a free parameter. We can in principle predict the harmonic structure in the QPO light curve by a full Comptonization calculation of the angle-dependent emission from a precessing hot flow (Ingram, Done & Zycki, in preparation). Until then, we simply allow the normalizations of the harmonics to be a free parameter but fix their width so that they have the same quality factor as the fundamental (apart from the subharmonic which is free; Rao et al. 2010). We use the method of Timmer & Koenig (1995) to generate a light curve from these narrow QPO Lorentzians and add this to the light curve already created for the broad-band noise.

Fig. 3 shows the full PSD given by the fiducial model parameters with $r_o = 50$ (black), 20 (red) and 10 (green). For clarity, we set the normalizations of the QPO harmonics to zero, and increase the normalization and quality factor of the fundamental as r_o decreases to match with the data. This captures the essence of the observed evolution of the PSD in terms of a decreasing truncation radius.

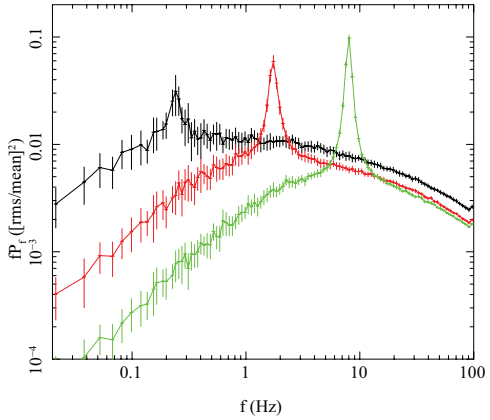

Figure 3. The predicted PSD for the fiducial model parameters with $r_o = 50$ (black), 20 (red) and 10 (green). For clarity, the normalization of the fundamental is set to increase as r_o reduces and all the other QPO components are normalized to zero.

3 FITTING TO DATA

The main problem with power spectral fitting is that we can only calculate the power spectrum of an infinitely long light curve which, of course, can neither be observed nor simulated! The discreet version of the power spectrum is the periodogram which provides an estimate of the power spectrum, but a very poor one with large, non-Gaussian errors. There are two main techniques for smoothing the periodogram in order to improve the statistics (i.e. smaller and more Gaussian errors). First, the light curve can be split up into shorter segments with the PSD estimate given by averaging the periodogram over all segments. Secondly (or additionally), the linearly sampled PSD can be averaged into logarithmic/geometrically spaced frequency intervals (e.g. van der Klis 1989).

In ID11, we averaged the logarithm of the periodogram in order to get a PSD estimate which, for active galactic nuclei or any red noise power spectra, gives an unbiased estimate with errors close to Gaussian for a small amount of smoothing compared with averaging the periodogram linearly (Papadakis & Lawrence 1993). However, the BHB power spectra are better described by band limited noise, with high $d^2P(f)/df^2$. In this case, averaging the logarithmic periodogram gives a biased estimator [because the bias is proportional to $d^2P(f)/df^2$; Papadakis & Lawrence 1993]. Because of this, we experienced some statistical difficulties, the most obvious being apparent overfitting with reduced χ^2 falling well below unity (indicating a good fit) for parameters which have a high rejection probability. Here, instead we average the periodograms linearly and test to see how reliable χ^2 is as a measure of goodness of fit in this case. This also has the added advantage that we can compare to the white-noise-subtracted PSD from the data, whereas, because logarithmic power must be positive definite, we previously had to add white noise to the model.

We split each simulated model light curve of 4096 s into 32 segments each of duration 128 s, so the lowest frequency we can resolve is 1/128 Hz. We average these at each linearly spaced frequency bin, j, to get $P_{\rm mod}(f_j)$ and the dispersion around this mean $\Delta P_{\rm mod}(f_j)$. We then re-bin the result in the same way as the data in geometrically separated frequency bins, f_J, so that the model power is $P_{\rm mod}(f_J) = \left[\sum P_{\rm mod}(f_j)/\Delta P^2_{\rm mod}(f_j)\right] \Delta P^2_{\rm mod}(f_J)$ and its propagated error is $\Delta P_{\rm mod}(f_J) = \left[\sum(1/\Delta P^2_{\rm mod}(f_j))\right]^{-1/2}$, where the sums are taken over all the f_j which fall inside the bin width of f_J. The lowest geometric frequency bin typically only contains one

linear frequency bin, so is an average over only 32 points. While this would not be Gaussianly distributed for red noise (Papadakis & Lawrence 1993), our noise power spectra at these low frequencies are typically white, so this is approximately Gaussian. At higher frequencies, the PSD is made from averaging many more estimates, so is again approximately Gaussian. Thus, we should get a good estimate for the best fit by minimizing χ^2 between the model and data using

$$\chi^2 = \sum_J \frac{[P_{\rm mod}(f_J) - P_{\rm obs}(f_J)]^2}{\Delta P^2_{\rm mod}(f_J) + \Delta P^2_{\rm obs}(f_J)}, \quad (9)$$

where $P_{\rm obs}(f_J)$ is the PSD of the data, and $\Delta P_{\rm obs}(f_J)$ is the propagated error from the (linear) averaging of the multiple segments and geometric re-binning. For the data, the number of segments depends on the observation length rather than being fixed as in the model. For our light curves, this gives 41, 26, 13, 14 and 14 segments for observations 1–5, respectively. The rather small numbers of segments in observations 3–5 mean that the errors on the data at the lowest frequencies may only be approximately Gaussian, but the total error is given by the sum of those from *both* the model and the data, so the fit statistic should remain close to a χ^2 distribution, so minimizing χ^2 should indeed return the best fit.

We have coded the entire model into XSPEC for public release as PROPFLUC, described in detail in Appendix A.

4 EXAMPLE FITS TO XTE J1550−564

We use the same data as in ID11 so that we can directly compare results, i.e. *RXTE* data from the 1998 rise to outburst of XTE J1550−564 (Sobczak et al. 2000; Wilson & Done 2001; Remillard et al. 2002; Altamirano 2008; Rao et al. 2010) from Obs IDs: 30188-06-03-00, 30188-06-01-00, 30188-06-01-03, 30188-06-05-00 and 30188-06-11-00 (hereafter observations 1–5, respectively). We consider energy channels 36–71 (i.e. 10–20 keV) in order to avoid any direct contamination from the disc emission.

We fit each observed PSD to derive the parameters of the smoothly broken power-law surface density. We assume that the shape of the surface density stays constant across all data sets, but its normalization Σ_0 can change. We also allow the bending wave radius to be a free parameter, $r_{\rm bw} = 3(h/r)^{-4/5}a_*^{2/5}$ (where h/r is the scaleheight of the flow). As we have fixed the spin, the best-fitting value of $r_{\rm bw}$ gives us an estimate of the scaleheight of the flow which may change through the transition due to the increase in seed photons from the disc cooling the flow. The inner radius of the flow is tied across all the data sets, and we fit for $r_{\rm o}$. The remaining free parameters which determine the broad-band noise are the level of MRI fluctuations generated over each decade in frequency, $F_{\rm var}$, and the emissivity index, γ (held constant across all five observations).

While XSPEC can fit the model to the five PSD simultaneously, this is very slow. Instead, we used trial and error to set values of the parameters which are tied across all the data sets and then fix these to fit the remaining parameters for each PSD individually.

4.1 Fit results

The data and best-fitting model PSD are shown in Fig. 4. These give a reasonable reduced χ^2 value of 1.09 (764.6 for 704 degrees of freedom), unlike ID11. We check the goodness of fit a posteriori by calculating the rejection probability (Uttley, McHardy & Papadakis 2002; Markowitz et al. 2003). We do this by taking the minimized χ^2 value which tells us about the agreement between the PSD

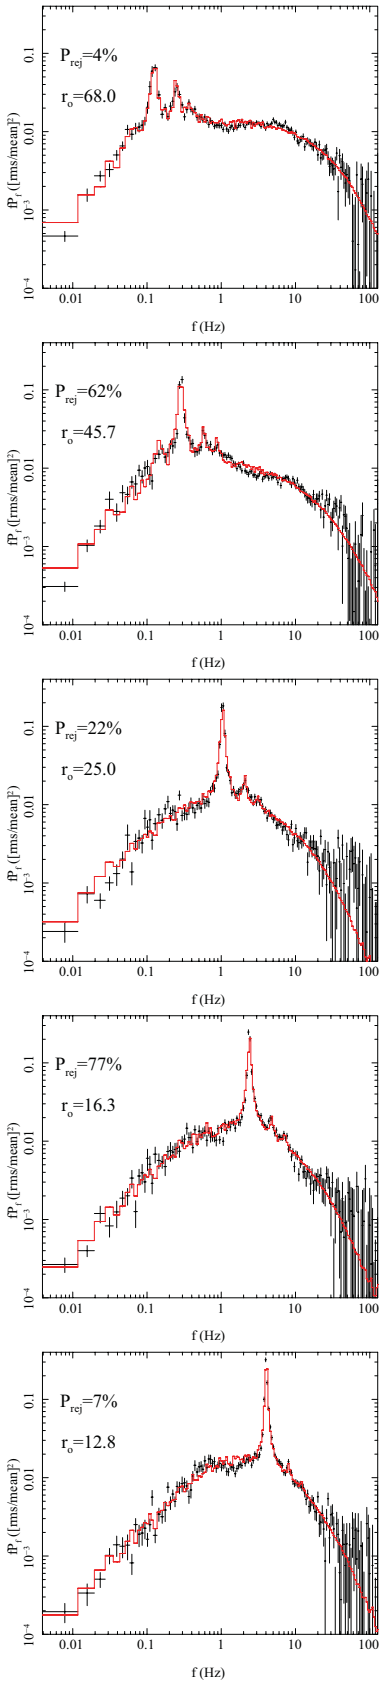

Figure 4. Best-fitting PSDs along with data points for observations 1–5 (top to bottom, respectively). The rejection probability, $P_{\rm rej}$, and truncation radius, $r_{\rm o}$, are included in each plot. The rest of the best-fitting physical parameters are included in Table 1.

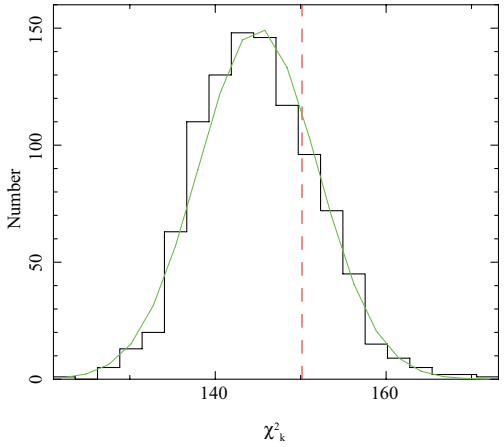

Figure 5. Distribution of χ^2_k values calculated using the best-fitting model parameters for observation 4. The green line illustrates that this is a nearly Gaussian distribution. The red dashed line picks out the χ^2 value for this observation and we see that although it is larger than the mean χ^2_k value, it still lies believably within the distribution, meaning we can be confident that the model fits.

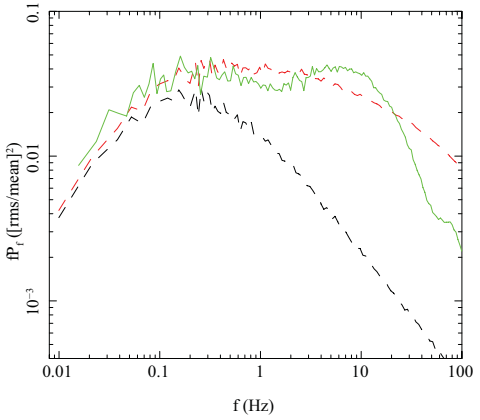

Figure 6. The red dashed line is the PSD predicted using the fiducial model parameters with $r_o = 50$ (i.e. $\lambda = 1$) whereas the black dashed line is for $\lambda = 7.6$ with all other parameters the same. For the green solid line, we use the same parameters as used for the black line but we have changed the model by assuming the annulus containing r_{bw} to be more variable than the other annuli. We see we can recover the amount of high-frequency power required to match the observations using this assumption.

estimate for the observed data, $P_{obs}(f)$, and that for one particular realization of the model, $P_{mod}(f)$. We then simulate many more (1000) realizations with the same model parameters in order to calculate many values of

$$\chi^2_k = \sum_f \frac{[P_{mod}(f) - P_k(f)]^2}{\Delta P^2_{mod} + \Delta P^2_k}, \quad (10)$$

where $P_k(f)$ is the PSD estimate for the kth realization. The rejection probability, P_{rej}, is given by the percentile of χ^2_k values which are smaller than χ^2. This does not assume that the errors are Gaussian, so is more strictly accurate as it assesses the likelihood that $P_{obs}(f)$ does not belong to the distribution which $P_{mod}(f)$ and each $P_k(f)$ belong to. We derive $P_{rej} = 4$, 62, 22, 77 and 7 per cent for observations 1–5, respectively. The lowest values of P_{rej} obviously imply a very good fit but even the higher values are still acceptable. Fig. 5 shows the distribution of χ^2_k values from the P_{rej} calculation (black stepped line) using the best-fitting parameters for observation 4. The red dashed line shows the χ^2 value for this observation and we see that although it is larger than most χ^2_k values, it still lies believably within the distribution. We also plot (green solid line) a Gaussian with the same mean, standard deviation and normalization as the distribution and we see very good agreement between the two. This confirms that the PSD estimate we use does indeed give (approximately) Gaussian errors and therefore χ^2 is a reliable measure of goodness of fit.

Table 1 shows all of the best-fitting physical parameters. Some of these parameters are very similar to those derived from the previous model fits in ID11, e.g. the truncation radius moves from $r_o = 68$ to 13, while F_{var} increases throughout the transition. However, our new parametrization means that we can directly explore the change in bending wave radius, r_{bw}, and normalized surface density Σ_0. The bending wave radius increases, implying that the flow scaleheight, h/r, is collapsing. This makes sense physically as the decreasing truncation radius means that the flow is cooled by an increasing number of seed photons, so the electron temperature decreases. The spectra also show that the optical depth increases (as is also implied by the increasing normalized surface density). This increases the coupling between electrons and ions, so the ion temperature also decreases (Malzac & Belmont 2009). The flow is held up (at least partly) by ion pressure, so the scaleheight of the flow collapses.

5 DISCUSSION

We have improved upon the model of ID11 by including a surface density profile which has the same shape as predicted by GRMHD simulations. We obtain an excellent fit to data for five observations, and the evolution of the best-fitting parameters is self-consistent. However, we require the surface density interior to the bending wave radius to drop-off as r^λ with $\lambda \sim 1$, where as the simulations predict $\lambda \approx 7$ (see Fig. 1). The most likely reason for this apparent discrepancy is that the torque created by the misalignment between flow and black hole angular momenta not only creates a drop-off in surface density, but also generates extra turbulence which we do not account for in our model. Because the surface density sets the emissivity, we can still reproduce the observations by overpredicting the surface density at small r to compensate for underpredicting the intrinsic variability. In Fig. 6, we replot the predicted PSD for $\lambda = 7.6$ (dashed black line) and $\lambda = 1$ (dashed red line) without

Table 1. Best-fitting physical parameters for observations 1–5. A $\equiv$ symbol indicates that the parameter has been fixed.

Obs	Σ_0	ζ	λ	κ	r_i	r_o	$h/r\,(r_{bw})$	F_{var}	γ
1	5.43					68.0	0.41(4.6)	0.32	
2	10.48					45.7	0.27(6.5)	0.31	
3	21.73	$\equiv 0$	$\equiv 0.9$	$\equiv 3.0$	$\equiv 3.3$	25.0	0.21(8.0)	0.36	5.28
4	30.03					16.3	0.13(12.03)	0.43	
5	30.36					12.8	0.12(12.1)	0.48	

errors for clarity. For the green solid line, also plotted without errors, we have changed the model slightly. We again set $\lambda = 7.6$ but now the fractional variability in the annulus containing $r_{\rm bw}$ is higher (by a factor of 10) than that at all other annuli so as to approximate the additional turbulence created by the bending waves. We see that it is possible to qualitatively reproduce the shape of the broad-band noise using a surface density profile consistent with simulations if we include this extra assumption.

It is interesting that the green line in Fig. 6 does not have a flat top between low- and high-frequency breaks as the model generally predicts, but rather has a 'bump' at $\sim$7 Hz and another at $\sim$0.15 Hz. There are actually many observations of bumpy power spectra such as this which cannot be well described by the model in its current state (e.g. Axelsson, Borgonovo & Larsson 2006; Wilkinson & Uttley 2009). It therefore looks likely that the variability generated by the MRI is not as uniform as we naively assume, and actually some regions produce more variability than others, thus giving rise to a bumpy power spectrum such as the green line in Fig. 6.

6 CONCLUSIONS

We have made some improvements to a model designed to predict the power spectral behaviour of BHBs in the context of a truncated disc/hot inner flow geometry which can also explain the energy spectral evolution. The model uses simple, theoretically motivated assumptions in order to reproduce the shape and evolution of the broad-band noise with the extra requirement that the QPO is generated by Lense–Thirring precession of the entire hot inner flow. Because the model now assumes a surface density profile consistent with that predicted by GRMHD simulations, we can now gain more physical insight from the evolution of best-fitting parameters which reproduce the observed evolution of the PSD. A coherent picture is now emerging: as the truncation radius, $r_{\rm o}$, moves inwards, the increased number of seed photons incident on the flow cools it, thus reducing both the electron and ion temperatures, $T_{\rm e}$ and $T_{\rm i}$, respectively. The Comptonized emission from the flow is therefore softer and, in addition to this, the lower ion temperature gives rise to a lower pressure, meaning that the scaleheight of the flow, h/r, should collapse. The bending wave radius, which sets the shape of the surface density, is given by $r_{\rm bw} = 3(h/r)^{-4/5}a_*^{2/5}$ and therefore increases as h/r collapses. Also, because the volume of the flow is reducing, the surface density must increase and, by mass conservation, the infall velocity decreases. When we fit the model to five observations of XTE J1550−564, we see all of these trends: $r_{\rm o}$ reduces and $r_{\rm bw}$ increases as does Σ_0, the normalization of the surface density (and also the inverse of the normalization of the infall velocity).

It is also worth re-iterating some other successes of the model first addressed in ID11. The truncated disc geometry fundamentally predicts an inhomogeneous electron temperature across the radial extent of the flow and therefore the Comptonized emission from the flow should actually be inhomogeneous also. Specifically, emission from large radii should be softer than that from small radii because the outer regions of the flow can see more seed photons than the inner regions. This is quite a subtle effect and, due to the degenerate nature of spectral fitting, it is only now becoming possible to constrain inhomogeneous emission from the SED alone (Makishima et al. 2008; Takahashi et al. 2008; Kawabata & Mineshige 2010). However, it has long been possible to observe this effect using the technique of frequency-resolved spectroscopy, which involves constraining the SED for a given temporal frequency range. Revnivtsev, Gilfanov & Churazov (1999) show that the SED of the fast variability is clearly harder than that of the slow variability. This makes sense in the truncated disc geometry because the fast variability predominantly comes from the inner regions which have a harder spectrum and the slow variability predominantly comes from the outer regions which have a softer spectrum. Also, a greater number of photons emitted from the outer regions of the flow will reflect off the disc than that from the inner regions, meaning that we should expect a greater reflection fraction for the slow variability than for the fast variability. This effect is also observed: Revnivtsev et al. (1999) show that the SED of the slow variability displays much stronger reflection features than that of the fast variability. The frequency-dependent time lags between hard and soft energy bands (with hard lagging soft; Miyamoto & Kitamoto 1989; Nowak et al. 1999) also follow directly from the idea of an inhomogeneous spectrum because slow variability generated in the outer regions modulates the soft spectrum immediately but takes time to propagate down to the inner regions where it modulates the hard spectrum.

Taking all of this into consideration, it seems apparent that the model has the capability to explain most of the spectral variability properties of BHBs. In a future paper (Ingram, Done & Zycki, in preparation), we will include an energy dependence in order to explicitly compare the model predictions for different properties such as the PSD, the SED, the lag spectrum and the frequency-resolved SED simultaneously. It is important to note that no other geometry can currently explain anything like the range of different observational properties that it is possible to explain with the truncated disc model.

However, although we believe the *trends* in best-fitting parameter values to be reliable, their *absolute* values should not be taken too seriously. This is because there are a few complexities not currently included in the model. For example, we currently effectively assume that the disc is stable, which is not true, at least in the low/hard state (Wilkinson & Uttley 2009). Although we only consider energies at which the Comptonized emission dominates, the disc is feeding the flow and therefore disc variability should propagate to the flow and modulate the hard emission. This means that the lowest frequencies in the PSD are actually being generated in the disc and not in the flow, meaning we overpredict the truncation radius, $r_{\rm o}$. The main uncertainty associated with the model is that it is unclear exactly how the disc and flow couple together. Although the most likely truncation mechanism is evaporation via thermal conduction (e.g. Liu, Meyer & Meyer-Hofmeister 1997; Różańska & Czerny 2000; Mayer & Pringle 2007), the details of this process are still far from well understood and, in particular, numerical simulations of a truncated disc/hot inner flow configuration are far beyond current computing capabilities. Whatever the specific nature of the coupling, it seems very likely that the disc will exert a torque on the flow, especially in a region where the flow overlaps the disc, which would slow down precession. This means that $r_{\rm o}$ would need to be smaller in order for the model to reproduce both the QPO *and* the broad-band noise. For this reason, we see our best-fitting values of $r_{\rm o}$ as upper limits rather than definitive measurements.

Still, it is extremely encouraging that this model can produce an excellent fit to PSD data whilst also having the potential to qualitatively reproduce many other properties seen in the data.

ACKNOWLEDGMENTS

AI acknowledges the support of an STFC studentship. AI and CD thank Magnus Axelsson for useful conversations on the double-peaked PSD, and Chris Fragile, Kris Beckwith and Omer Blaes for physical intuition about the MRI.

REFERENCES

Altamirano D., 2008, PhD thesis, Universiteit van Amsterdam
Arévalo P., Uttley P., 2006, MNRAS, 367, 801 (AU06)
Arnaud K., Borkowski K. J., Harrington J. P., 1996, ApJ, 462, L75
Axelsson M., Borgonovo L., Larsson S., 2006, A&A, 452, 975
Bardeen J. M., Petterson J. A., 1975, ApJ, 195, L65
Beckwith K., Hawley J. F., Krolik J. H., 2008, MNRAS, 390, 21
Belloni T. M., 2010, in Comastri A., Angelini L., Cappi M., eds, Proc. AIP Conf. Vol. 1248, X-ray Astronomy 2009: Present Status, Multi-Wavelength Approach and Future Perspectives. Am. Inst. Phys., New York, p. 107
Belloni T., Psaltis D., van der Klis M., 2002, ApJ, 572, 392
Cabanac C., Henri G., Petrucci P.-O., Malzac J., Ferreira J., Belloni T. M., 2010, MNRAS, 404, 738
Davis S. W., Done C., Blaes O. M., 2006, ApJ, 647, 525
Dexter J., Fragile P. C., 2011, ApJ, 730, 36
Done C., Diaz Trigo M., 2010, MNRAS, 407, 2287
Done C., Gierliński M., Kubota A., 2007, A&AR, 15, 1 (DGK07)
Esin A. A., McClintock J. E., Narayan R., 1997, ApJ, 489, 865
Fragile P. C., 2009, ApJ, 706, L246
Fragile P. C., Meier D. L., 2009, ApJ, 693, 771
Fragile P. C., Mathews G. J., Wilson J. R., 2001, ApJ, 553, 955
Fragile P. C., Blaes O. M., Anninos P., Salmonson J. D., 2007, ApJ, 668, 417
Gierliński M., Nikołajuk M., Czerny B., 2008, MNRAS, 383, 741
Gierliński M., Done C., Page K., 2009, MNRAS, 392, 1106
Hawley J. F., Balbus S. A., 1991, ApJ, 376, 223
Ingram A., Done C., 2010, MNRAS, 405, 2447
Ingram A., Done C., 2011, MNRAS, 415, 2323 (ID11)
Ingram A., Done C., Fragile P. C., 2009, MNRAS, 397, L101 (IDF09)
Kawabata R., Mineshige S., 2010, PASJ, 62, 621
Klein-Wolt M., van der Klis M., 2008, ApJ, 675, 1407
Kotov O., Churazov E., Gilfanov M., 2001, MNRAS, 327, 799
Krolik J. H., Hawley J. F., 2002, ApJ, 573, 754
Kumar S., Pringle J. E., 1985, MNRAS, 213, 435
Lyubarskii Y. E., 1997, MNRAS, 292, 679
Makishima K. et al., 2008, PASJ, 60, 585
Malzac J., Belmont R., 2009, MNRAS, 392, 570
Markoff S., Nowak M. A., Wilms J., 2005, ApJ, 635, 1203
Markowitz A. et al., 2003, ApJ, 593, 96
Mayer M., Pringle J. E., 2007, MNRAS, 376, 435
Miller J. M., Homan J., Miniutti G., 2006, ApJ, 652, L113
Miyamoto S., Kitamoto S., 1989, Nat, 342, 773
Narayan R., Yi I., 1995, ApJ, 452, 710
Nowak M. A., Vaughan B. A., Wilms J., Dove J. B., Begelman M. C., 1999, ApJ, 510, 874
Papadakis I. E., Lawrence A., 1993, MNRAS, 261, 612
Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, Numerical Recipes in Fortran, 2nd edn. Cambridge Univ. Press, Cambridge
Psaltis D., Belloni T., van der Klis M., 1999, ApJ, 520, 262
Rao F., Belloni T., Stella L., Zhang S. N., Li T., 2010, ApJ, 714, 1065
Remillard R. A., Sobczak G. J., Muno M. P., McClintock J. E., 2002, ApJ, 564, 962
Revnivtsev M., Gilfanov M., Churazov E., 1999, A&A, 347, L23
Rodriguez J., Corbel S., Hannikainen D. C., Belloni T., Paizis A., Vilhu O., 2004, ApJ, 615, 416
Różańska A., Czerny B., 2000, MNRAS, 316, 473
Rykoff E. S., Miller J. M., Steeghs D., Torres M. A. P., 2007, ApJ, 666, 1129
Schnittman J. D., 2005, ApJ, 621, 940
Schnittman J. D., Homan J., Miller J. M., 2006, ApJ, 642, 420
Shakura N. I., Sunyaev R. A., 1973, A&A, 24, 337
Sobczak G. J., McClintock J. E., Remillard R. A., Cui W., Levine A. M., Morgan E. H., Orosz J. A., Bailyn C. D., 2000, ApJ, 544, 993
Sobolewska M. A., Życki P. T., 2006, MNRAS, 370, 405
Steiner J. F. et al., 2011, MNRAS, 416, 941
Stella L., Vietri M., 1998, ApJ, 492, L59
Takahashi H. et al., 2008, PASJ, 60, 69
Timmer J., Koenig M., 1995, A&A, 300, 707
Titarchuk L., Osherovich V., 1999, ApJ, 518, L95
Uttley P., McHardy I. M., 2001, MNRAS, 323, L26
Uttley P., McHardy I. M., Papadakis I. E., 2002, MNRAS, 332, 231
van der Klis M., 1989, ARA&A, 27, 517
van der Klis M., 2004, Advances Space Res., 34, 2646
Wagoner R. V., Silbergleit A. S., Ortega-Rodríguez M., 2001, ApJ, 559, L25
Wijnands R., van der Klis M., 1999, ApJ, 514, 939
Wilkinson T., Uttley P., 2009, MNRAS, 397, 666
Wilson C. D., Done C., 2001, MNRAS, 325, 167

APPENDIX A: USING PROPFLUC

We intend to release the model publicly as the XSPEC local model, PROPFLUC. Here we include some tips for anyone wanting to use the model.

A1 Data

We use powspec from XRONOS in order to create a power spectrum from the observed light curve. We set norm $= -2$, which means white noise will be subtracted, and choose the minimum light curve time-step, which is $\mathrm{d}t_{\rm obs} = 0.390625 \times 10^{-2}$ s for *RXTE* data. We set the number of time-steps per interval to $2^{15} = 32768$, meaning that the duration of an interval is $2^{15}\mathrm{d}t_{\rm obs} = 128$ s. This means that a periodogram will be calculated for each interval with minimum frequency 1/128 Hz and maximum (Nyquist) frequency $1/(2\mathrm{d}t_{\rm obs}) = 128$ Hz. The number of intervals per frame should be set to maximum so that powspec averages over as many intervals as the length of the observation allows, and we use a geometric rebinning with a constant factor of 1.045, resulting in 150 new bins. The resulting binned power spectrum can then be written to a data file in the form

$$f, \mathrm{d}f, P, \mathrm{d}P. \tag{A1}$$

XSPEC, however, is expecting to receive data in the form

$$E_{\rm min}, E_{\rm max}, F(E_{\rm max} - E_{\rm min}), \mathrm{d}F(E_{\rm max} - E_{\rm min}), \tag{A2}$$

where $E_{\rm min}$ and $E_{\rm max}$ are the lower and upper bounds of each energy bin and F is the flux. It is therefore necessary to create a data file with inputs

$$f_{\rm min}, f_{\rm max}, P(f_{\rm max} - f_{\rm min}), \mathrm{d}P(f_{\rm max} - f_{\rm min}), \tag{A3}$$

where $f_{\rm min}$ and $f_{\rm max}$ are the lower and upper bounds of each frequency bin. As f marks the centre of a bin and $\mathrm{d}f$ is defined such that $f_{\rm max} = f + \mathrm{d}f$ and $f_{\rm min} = f - \mathrm{d}f$, this equation can be rewritten as

$$f - \mathrm{d}f, f + \mathrm{d}f, 2P\mathrm{d}f, 2\mathrm{d}P\mathrm{d}f. \tag{A4}$$

We then use flx2xsp in order to convert this into a *.pha* file and also generate a diagonal response function. The data can now be loaded into XSPEC and, even though the axis on the plots is by default labelled as flux and energy, it is in fact reading in a power spectrum as a function of temporal frequency (i.e. the command ip euf will show frequency multiplied by power plotted against frequency for both data and model).

A2 Model

The model consists of a FORTRAN program, *propfluc.f*, and a data file *lmodel_pf.dat*. These two files are all that is needed to load the model using the local model functionality. The model has 18

Table A1. Summary of model parameters.

	Parameter	Comments	Value for obs 1
1	Sigma0	Normalization of surface density.	5.43
2	rbw	Bending wave radius – dictates where $\Sigma(r)$ breaks.	4.60
3	kappa	Dictates sharpness of the break.	$\equiv 3.0$
4	lambda	Dictates $\Sigma(r)$ for $r < r_{\rm bw}$.	$\equiv 0.9$
5	zeta	Dictates $\Sigma(r)$ for $r > r_{\rm bw}$.	$\equiv 0.0$
6	Fvar	Intrinsic amount of variability generated per decade in $f_{\rm visc}$.	0.32
7	fbmin	This is $f_{\rm visc}(r_{\rm o})$. It is much easier to set this instead of $r_{\rm o}$.	0.129
8	ri	Inner radius	$\equiv 3.3$
9	sig_qpo	QPO width (fundamental). Width of higher harmonics is tied to this.	0.0226
10	sig_subh	Width of the subharmonic. This can have a different Q value to the other harmonics.	0.0283
11	n_qpo	Normalization of fundamental (first harmonic).	0.244806
12	n_h	Normalization of second harmonic.	0.1706
13	n_3h	Normalization of third harmonic.	0.1018
14	n_subh	Normalization of subharmonic.	0.0967
15	em_in	Emissivity index (i.e. γ in the text).	5.281
16	dL	The model gives the option to generate a Gaussian error on each point of the simulated light curve, thus creating white noise. To match a typical *RXTE* observation, this needs to be $\mathrm{d}L \sim 0.8$; however, we recommend setting this to zero and using white-noise-subtracted data.	$\equiv 0.0$
17	nn	Sets the number of time-steps in the simulated light curve (i.e. the light curve has a total duration of $2^{nn}\mathrm{d}t$). This must be an integer because the model uses a fast Fourier transform algorithm (Press et al. 1992). The PSD estimate of the model must be calculated on the same interval as the data (128 s) and therefore the value of nn used dictates how many intervals are averaged over.	$\equiv 22.0$
18	Ndec	Sets the radial resolution. If this is particularly high, the code is very slow! $N_{\rm dec} = 15$ should be sufficient. The total number of annuli used is calculated from this.	$\equiv 15.0$

parameters, summarized in Table A1, plus XSPEC always includes a 19th normalization parameter which must be set to (and fixed at) unity. The simulated light curve is generated using a time-step of $\mathrm{d}t = \mathrm{d}t_{\rm obs}/4 = 9.76562\times10^{-4}$ s. It is important that this time-step is short because the Nyquist frequency must be higher than the highest frequency at which significant variability is generated. The final power spectrum is calculated using 2^{17} steps per interval, meaning that each interval is $2^{17}\mathrm{d}t = 2^{17}\mathrm{d}t_{\rm obs}/4 = 2^{15}\mathrm{d}t_{\rm obs} = 128$ s. The simulated power spectrum is then binned into the same frequency bins used for the observed power spectrum. For this reason, it is vital that the periodograms are calculated on the same interval (i.e. 128 s) for both model and data; the use of two different intervals could result in empty bins in the simulated power spectrum which does not help χ^2! In Table A1, we see that it is possible for the user to decide on the length of simulated light curve (parameter 17). Since the interval length is fixed, this dictates how many intervals the power spectrum is averaged over. We recommend $nn = 22$ (32 intervals) for fitting but this does make the code very slow. Preliminary fitting is best done with $nn = 20$ (eight intervals) as this is faster but provides a good enough PSD estimate to work with. It should be noted that this setting slightly underpredicts the power but it is a constant offset and so the best fit found using $nn = 20$ has a higher value for $F_{\rm var}$ than that found using $nn = 22$ but the other parameters are largely unaffected. The main advantage of using $nn = 22$ is that χ^2 gives a much more reliable estimate of goodness of fit.

The model is difficult to fit, partly because of the stochastic nature of the power spectrum and partly because of the complicated relationship between parameters. We recommend finding a good fit by eye first and fixing a few key parameters before fitting. We set XSPEC to calculate the gradient in χ^2 numerically rather than analytically and set the critical $\Delta\chi^2$ value to 0.1 rather than the default 0.01. Finally, the third column of Table A1 shows all of our best-fitting model parameters for observation 1, with a $\equiv$ symbol indicating that the parameter is fixed.

This paper has been typeset from a TEX/LATEX file prepared by the author.