A High-Frequency Doppler Feature in the Power Spectra of Simulated Grmhd Black Hole Accretion Disks

Citation

Wellons, Sarah, Yucong Zhu, Dimitrios Psaltis, Ramesh Narayan, and Jeffrey E. McClintock. 2014. "A High-Frequency Doppler Feature in the Power Spectra of Simulated Grmhd Black Hole Accretion Disks." *The Astrophysical Journal* 785 (2) (April 4): 142. doi:10.1088/0004-637x/785/2/142. <http://dx.doi.org/10.1088/0004-637X/785/2/142>.

Published Version

doi:10.1088/0004-637X/785/2/142

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:27802020>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

A HIGH-FREQUENCY DOPPLER FEATURE IN THE POWER SPECTRA OF SIMULATED GRMHD BLACK HOLE ACCRETION DISKS

SARAH WELLONS¹, YUCONG ZHU¹, DIMITRIOS PSALTIS², RAMESH NARAYAN¹, AND JEFFREY E. MCCLINTOCK¹

¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA, USA; swellons@cfa.harvard.edu

² Astronomy Department, University of Arizona, 933 North Cherry Avenue, Tucson, AZ, USA

Received 2013 December 11; accepted 2014 March 3; published 2014 April 4

ABSTRACT

Black hole binaries exhibit a wide range of variability phenomena, from large-scale state changes to broadband noise and quasi-periodic oscillations, but the physical nature of much of this variability is poorly understood. We examine the variability properties of three GRMHD simulations of thin accretion disks around black holes of varying spin, producing light curves and power spectra as would be seen by observers. We find that the simulated power spectra show a broad feature at high frequency, which increases in amplitude with the inclination of the observer. We show that this high-frequency feature is a product of the Doppler effect and that its location is a function of the mass and spin of the black hole. This Doppler feature demonstrates that power spectral properties of the accretion disk can be tied to, and potentially used to determine, physical properties of the black hole.

Key words: accretion, accretion disks – black hole physics – X-rays: binaries

Online-only material: color figures

1. INTRODUCTION

Accreting black holes in binary systems are variable on many timescales. In the long term, they are observed to undergo changes in state as the mass accretion rate changes, from a bright state dominated by soft thermal emission to a low-luminosity state dominated by a hard power-law spectrum (see Remillard & McClintock 2006 for a review). These large-scale changes are interpreted (Esin et al. 1997; Done et al. 2007; Narayan & McClintock 2008; Yuan & Narayan 2014) as the inner part of the accretion disk switching from an optically thick and geometrically thin disk (Shakura & Sunyaev 1973; Novikov & Thorne 1973) to an optically thin and geometrically thick disk, or ADAF (Narayan & Yi 1994, 1995; Abramowicz et al. 1995).

On shorter timescales, the power spectra of observed accreting black holes consist of broad band noise spanning a large range of frequencies between a low and high break frequency plus, in some cases, sharper, narrower features known as quasi-periodic oscillations (QPOs; see van der Klis 2006 for a review). Both components can be well-described by a series of Lorentzians (Belloni et al. 2002), but the physical mechanisms by which the Lorentzian components are produced are poorly understood.

Models for the broad band noise require a mechanism which produces variability over a large range of dynamical time- and length scales. The magneto-rotational instability (MRI; Balbus & Hawley 1991), for example, creates turbulence in the accretion flow and is believed to be responsible for angular momentum transport in accretion disks. This MHD turbulence does indeed generate fluctuations over a broad range of timescales (Krolik & Hawley 2002) and thus could potentially produce the sort of broad-band, power-law power spectrum that is typically observed. Proposed mechanisms for QPOs are many and varied, including (but not limited to) tilted accretion flows (Stella & Vietri 1998; Marković & Lamb 1998; Ingram & Done 2011), diskoseismic oscillations (Nowak & Wagoner 1993; Kato 2001; Silbergleit et al. 2001) and resonances in the accretion disk (Abramowicz et al. 2003).

Identifying features in the power spectrum of an accreting black hole could characterize special frequencies in the disk related to the black-hole mass or spin. For example, for a black hole of known mass, the identification of a power spectral feature with the orbital frequency at the innermost stable circular orbit (ISCO) of the accretion disk would indicate the black-hole spin a_* , since disks around black holes with higher spin have smaller ISCOs.

By examining the variability properties of simulations of accretion disks around stellar-mass black holes, we aim to shed light on the physical origins of some of these power spectral features and provide some tools for understanding observations. In this paper, we analyze general relativistic magnetohydrodynamic (GRMHD) disk simulations from an observer’s point of view. We examine how the variability properties change with the observer’s inclination and choice of photon energy band and identify power spectral features that may hint at black hole properties. In Section 2, we describe the simulations themselves as well as the methods by which we produce the observed lightcurves and power spectra. In Section 3, we investigate a high-frequency feature in the power spectrum, determine its origin, and show how it changes with spin, inclination, and energy. We discuss the implications for observers who seek to characterize black holes in Section 4, and conclude in Section 5.

Throughout, we use standard gravitational units and set $G = c = 1$. Thus, we express lengths in units of the black hole mass M , which corresponds to a physical unit of GM/c^2 . Similarly, our unit of time is also M , which corresponds to a physical unit of GM/c^3 .

2. METHODOLOGY

2.1. GRMHD Simulations

In this paper, we analyze simulations of accretion disks around black holes with three different spin parameters: $a_* \equiv a/M = (0, 0.7, 0.9)$. The simulations are performed using the three-dimensional GRMHD code HARM (Gammie et al. 2003; McKinney 2006; McKinney & Blandford 2009) in a Kerr

Table 1
GRMHD Simulation Parameters

a_*	n_r	n_θ	n_ϕ	$\dot{M}$ ($\dot{M}_{\text{Edd}}$)	Length ^a (M)	$f_{\text{QPO}}(1/M)$	$f_{\text{peak}}^b(1/M)$	$r_{\text{peak}}^c(M)$	$r_{\text{isco}}(M)$
0	256	64	64	0.7	40000	8×10^{-4}	9.5×10^{-3}	6.53	6.00
0.7	256	64	96	0.64	20000	5×10^{-4}	1.45×10^{-2}	4.68	3.39
0.9	256	64	96	0.72	30000	3×10^{-4}	1.57×10^{-2}	4.4	2.32

Notes. The dimensions of the simulation grid are given by the number of cells in each direction n_r , n_θ , and n_ϕ . All simulations have $h/r \approx 0.1$.

^a Length refers to the length of time in the simulation which is used for analysis, and does not include the first 20,000 M where the simulation has not reached inflow equilibrium.

^b f_{peak} refers to the average peak frequency of P_{Dopp} , calculated as described in Section 3.

^c r_{peak} is the radius whose orbital frequency corresponds to f_{peak} .

spacetime with the appropriate value of a_* . These are the same simulations analyzed by Zhu et al. (2012) and Kulkarni et al. (2011), but have been run longer. Simulation length and other parameters are shown in Table 1.

Each simulation is initialized with an orbiting torus of gas in hydrostatic equilibrium in the equatorial plane around the black hole. The gas is evolved according to the GRMHD equations and is driven to steady-state accretion via MHD turbulence that emerges as a result of the MRI. This initial part of the simulation before inflow equilibrium is reached is not considered in the analysis that follows. Additionally, the simulations are only converged (at a steady $\dot{M} \approx 0.7 \dot{M}_{\text{Edd}}$) out to 15–20 M, so we disregard the region outside the radius of convergence.

The code HARM does not include radiative transfer. In the present work, we use an ad hoc cooling prescription (Shafee et al. 2008; Penna et al. 2010) whereby we effectively set the thickness of the disk. (Section 2.3 describes how we reconstruct the emergent spectrum.) All three simulations discussed here are thin disks with $h/r \approx 0.1$.

2.2. Raytracing

In order to produce the light curves that would be seen by a distant observer, we need to trace the paths taken by photons that are emitted at the disk midplane and navigate the curved spacetime around the black hole. The strong gravity around the black hole deflects the paths of photons and produces a warped image for a distant observer; see Figure 1 for examples. To reproduce this, we construct an elliptical grid in the observer’s image plane and follow the paths of photons back to their places of origin, as described by Kulkarni et al. (2011). Initial conditions for the position vector of each ray are given by their location in the image plane, and initial conditions for the momentum vector are given by assuming that the observer is sufficiently far away from the black hole that all photons travel parallel to the line of sight.

The trajectory of the ray from each grid cell in the image plane is evolved along null geodesics in a Kerr spacetime until the ray strikes the plane of the disk or crosses the event horizon and is lost. Then, knowing the location at which the ray struck (or emerged from) the disk as well as the strike angle, we reconstruct the emergent spectrum.

2.3. Calculation of Emitted Spectrum

We use the stellar atmospheres code TLUSTY (Hubeny & Lanz 1995), which has been adapted to the problem of black hole accretion disks (see Davis et al. 2005; Davis & Hubeny 2006), to calculate the radiation spectrum emerging from each disk annulus. TLUSTY treats annuli in the accretion disk as one-dimensional plane-parallel atmospheres. Under the assumption

Figure 1. Raytraced image of a simulated accretion disk around a non-spinning black hole out to a radius $r = 25$ M, at several inclinations. Colors indicate bolometric intensity in arbitrary but consistent units. As inclination increases, the image becomes warped as photon geodesics are deflected by the strong gravity near the black hole. Also, the Doppler effect becomes sharper, boosting the left side of the disk where the orbital velocity points toward the observer.

that the atmosphere is in hydrostatic equilibrium, the code solves for both the radiation field and the vertical structure of the annulus in the local comoving frame of the gas.

Free-free, free-bound, and bound-free radiative processes are all included, as well as an angle averaged Kompaneets treatment of Compton scattering (Hubeny et al. 2001). The physics of magnetic pressure support, convection, and heat conduction are, however, neglected. Deviations from local thermodynamic equilibrium are allowed and the level populations for H, He, C, N, O, Ne, Mg, Si, S, Ar, Ca, and Fe are explicitly computed via relevant rate equations (under the assumption of solar abundances). Only the lowest ionization energy level is considered for each ionization state, except for H and He where we include the lowest 9 and 4 levels, respectively.

With a final assumption that the local viscous heating is proportional to gas density, TLUSTY allows the full range of disk annuli to be encapsulated with just three parameters (Davis & Hubeny 2006): (1) the vertical column mass Σ , (2) the vertical tidal gravity $Q = g_\perp/z$ where $g_\perp$ is the vertical component of the gravitational acceleration and z is height above the midplane, and (3) the total emerging flux $F = \sigma T_{\text{eff}}^4$. With these three parameters and knowledge of the angle at which the ray emerges from the disk, the emergent spectrum is

fully determined. Note that TLUSTY computes the emerging intensity as a function of the angle of the ray with respect to the surface normal so that limb darkening is calculated self-consistently. (At very high inclinations where $\cos \theta_i \gg h/r$, however, this method would not be valid since the radii at which the ray intersects the midplane and the photosphere would differ significantly.)

Since the properties of each annulus are uniquely specified by only three parameters, we have constructed a large precomputed grid of annuli to span the full parameter range for stellar-mass black-hole accretion disks. In this work, we make use of a grid of over 40,000 TLUSTY models that were developed in Zhu et al. (2012). The use of a precomputed grid provides a significant computational time advantage; the overall spectra of many different disk models can be quickly computed by interpolating models from this grid. When the atmosphere is optically thick, the spectra produced by these annuli will resemble a blackbody (plus an extended non-Wien tail from non-saturated Comptonization), but we are not limited to cases of high optical depth—the code has the advantage of handling optically thin atmospheres as well.

The strike angle comes from the raytracing as described above and Q is simply a function of r ; the other parameters must be retrieved from the simulation. Because the simulation is three-dimensional, while the procedure described above requires a flat disk, T_{eff} and Σ are collapsed to the midplane along z . The effective temperature at a point on the two-dimensional disk is calculated by finding the total flux emitted within a vertical column. Similarly, the column density is calculated by summing up the mass in all the cells in that column; see Zhu et al. (2012) for details.

Having calculated these parameters for each midplane cell in the simulation, we retrieve the spectrum emitted in the fluid frame at the location where each ray strikes the disk using the appropriate pre-computed TLUSTY model. Finally, we calculate the observed spectrum by applying a redshift which is determined by the radius, strike angle, and orbital velocity as described by Kulkarni et al. 2011. This redshift accounts both for gravitational redshift and the Doppler shift.

The effect of the Doppler shift is visible in Figure 1. A highly inclined observer for whom the disk is rotating counter-clockwise will see bright, Doppler-boosted emission on the left where a large component of the fluid velocity is directed toward the line of sight. Conversely, on the right, emission is boosted away from the observer and the disk is dim.

We have tested the method described in Sections 2.2 and 2.3 and shown that it accurately recovers the emission produced by a Novikov–Thorne model disk; see Appendix A for details.

2.4. Light Curves and Power Spectra

The observed flux at a particular time is calculated by integrating over the disk in the image plane using the spectra retrieved as described in Section 2.3. The light curves shown in this paper do not take into account the differences in light travel times between photons emitted at different locations on the disk. Similar to Dexter et al. (2010), we found that including the delays has a negligible effect on the light curves and power spectra for a $10 M_{\odot}$ black hole, so we ignore them to make the computation more straightforward.

Figure 2 shows sample light curves at different photon energy bands. A band which is in the Rayleigh–Jeans regime for the temperatures in the disk (such as the 1–2 keV panels at the top) varies smoothly over a broad range of radii, whereas a

Figure 2. Light curves and images seen by an observer face-on to a simulated accretion disk around a non-spinning black hole out to a radius $r = 25 M$, in several photon energy bands. Units of flux and intensity are arbitrary but consistent within this Figure and with Figure 1. As the band moves to higher photon energies, the lightcurves and images are dominated by fluctuations at small radii.

higher-energy band in the Wien regime (lower panels) is dominated by fluctuations in the hottest, innermost radii.

At high energies, the spectral energy distribution is steep, so a small change in temperature results in a large change in intensity; this causes much larger variability in the light curves.

The primary tool used to characterize variability in X-ray binary systems is the power spectrum

$$P(f) = |F(f)|^2, \quad (1)$$

where $F(f)$ is the Fourier transform of the lightcurve. Following this practice, we use a fast Fourier transform method to compute the discrete Fourier transform and power spectrum of our light curves, which have temporal resolution $\Delta t = 5 M$ (the output cadence of the simulation, which corresponds to $\Delta t \approx 250 \mu s$ for a $10 M_{\odot}$ black hole). We are able to substantially reduce the noise in the power spectrum by segmenting the light curves into $N_s = 2-4$ pieces before taking the FFT and averaging the resulting power spectra together. We then rebin the power spectra in frequency in such a way that each bin has equal size in $\log f$, averaging the power spectral components in each bin. Then the error in the power spectrum is estimated to be

$$\sigma_S(f) = \frac{P(f)}{\sqrt{N_s N_{\text{bin}}}}, \quad (2)$$

where N_{bin} is the number of FFT bins contributing to each power spectrum frequency. Typical values for N_{bin} within a single power spectrum range from a few at the low-frequency end to ~ 200 at the high-frequency end (due to the logarithmic spacing of the bins), such that the uncertainty in the calculated power spectrum ranges from $\sim 35\%$ to 3% .

Figure 3. Power spectra of bolometric lightcurves seen by an observer at varying inclination to the accretion disk of a black hole with $a_* = 0$. The power spectra are well-described by one or more underlying power laws plus an apparent QPO around $10^{-3}/M$ and an additional broad feature from $(10^{-3}-10^{-2})/M$ which is stronger at higher inclinations. For a $10 M_\odot$ black hole, a frequency of $10^{-2}/M$ corresponds to about 200 Hz. The orbital frequency at the ISCO is plotted as a vertical dashed line; note that the broad feature is centered at a lower frequency, which suggests that the power originates at radii outside the ISCO.

In addition, we suppress leakage between frequencies by applying a Welch window to each segment,

$$w_k(t) = 1 - \left(\frac{k - N/2}{N/2} \right)^2, \quad (3)$$

where N is the number of points in each segment and k is the time index. Because the Welch window de-weights the points at the ends of the segments, it effectively removes the information they carry from the power spectrum. To compensate for this, we take an additional $N_s - 1$ segments which are offset from the original segments by half their length, and average these into the power spectra as well so that no information is lost (Press et al. 2007). Examples of the resulting power spectra are shown in Figure 3.

3. RESULTS

3.1. Simulation Power Spectra

The power spectra we obtain from our simulated lightcurves are characterized by three components, which can be seen for the non-spinning case in Figure 3.

1. An underlying power law.
2. A QPO at low frequencies, $\lesssim 10^{-3}/M$ (~ 10 Hz for a $10 M_\odot$ black hole).
3. A broad feature contributing extra power at higher frequencies, breaking off around $10^{-2}/M$ (~ 200 Hz for a $10 M_\odot$ black hole).

As discussed in Section 1, the underlying power law $P(f) \propto f^{-\alpha}$ may be representative of the constant fluctuations caused by MHD turbulence associated with the MRI. The power law index from the simulations ($\alpha \approx 3$) is steeper than typically observed or predicted ($\alpha = 2$), and the process producing it remains

Figure 4. Top panel: power spectra of bolometric light curves seen by observers who are face-on (0° , blue) and at high inclination (75° , red) to the accretion disk of a black hole with $a_* = 0.9$. Bottom panels: the same power spectra, broken down by radius of origin of the photons. At high inclination, each annulus contributes additional power at its orbital frequency (black line), which results in the broad high-frequency feature in the combined power spectrum.

to be precisely characterized. (For a discussion of several possible mathematical models for the underlying variability, see Pecháček et al. 2008, who also find a high-frequency feature similar to ours.)

The QPO is present for all spins and at all inclinations. It also remains in approximately the same location throughout, though it moves slightly with spin from ~ 16 Hz for $a_* = 0$ to ~ 10 Hz for $a_* = 0.7$ and ~ 6 Hz for $a_* = 0.9$. In this paper we focus on understanding the origin of the high-frequency feature and defer discussion of the QPO to a future paper.

The effect on the broad high-frequency feature of changing the observer's inclination is evident in Figure 3. As inclination increases, the feature becomes more and more prominent. These changes with inclination strongly suggest that the Doppler effect is of great importance for this feature—as the observer becomes more inclined with respect to the disk, the component of orbital motion along the line of sight increases, strengthening the boost supplied by the Doppler effect. As a result, at high inclinations the radii closest to the black hole possessing the highest orbital velocities are emphasized.

An examination of the power spectra broken down by radius confirms this emphasis on smaller radii. As seen in Figure 4, at high inclination each annulus contributes additional power at its own orbital frequency (the black line on the bottom two panels). This excess is not present for a face-on observer. Physically, this can be understood to mean that there are structures in the accretion disk that survive for at least one orbital period. As a feature passes through the left side of the disk, it experiences a Doppler boost and appears brighter, then dims as it passes on the right, and brightens again as it comes back around, repeating this process as long as the structure is present and thereby producing

Figure 5. Power spectra of light curves seen by an observer at 55° inclination in several energy bands for each of the three spins. As the band moves to higher energies, the high-frequency Doppler feature becomes more pronounced. Despite its origin in the thermal disk, the Doppler feature is noticeably present at high energies due to the steepness of the high-energy tail. (The different curves in each panel have been shifted vertically for clarity.)

periodicity at the orbital frequency. Based on this explanation, we refer to the resulting feature in the power spectrum as the “Doppler feature.”

Because the orbital frequency changes as a function of radius, the Doppler feature is necessarily broad, extending up to the orbital frequency of the radius which dominates the variability. This radius does not appear to change with inclination or the energy band in which the disk is observed (except perhaps at the highest spin), though these two quantities do certainly affect how pronounced the feature is. As can be seen in Figure 5, the Doppler feature is present at all energies, both soft and hard. Moving to higher energies increases the prominence of the feature, since the spectrum is steep at those energies and a small change in temperature produces a larger change in intensity than in the lower-energy bands where the spectrum is shallower.

The location of the Doppler feature does change with spin, however, moving to higher frequencies as spin increases (see Table 1). At higher spins, circular orbits are stable closer to the black hole, so the inner edge of the accretion disk reaches smaller, hotter, and faster-moving radii. Accordingly, the higher orbital frequencies of these radii are visible in the power spectrum. A priori, one might expect the Doppler feature to be dominated by the orbital frequency of the ISCO, but this is not the case. In fact, the feature tends to peak or break at frequencies corresponding to radii up to ~ 2 M outside the ISCO, implying that there is another location in the disk which is dominant, and necessitating a more detailed explanation.

The location of the break is affected not only by the strength of the Doppler effect, which increases toward smaller radii, but also by the emissivity of the disk, which peaks further out. In the following sections, we model the shape of the Doppler feature as a function of spin, inclination, and temperature profile, and show how its break lies between the orbital frequencies at the ISCO and at the peak of emissivity.

3.2. A Simplified Model of the Doppler Feature

As a hotspot travels around the disk, its intensity as seen by an inclined observer is modulated by the Doppler effect. Of course, for a smooth, non-fluctuating disk, this modulation will not produce any variability; some background fluctuations (such as those provided by the MHD turbulence) are required. Thus the strength of the signal from the Doppler effect for a given radius will depend on (1) the strength of the underlying

Figure 6. Distributions of temperature within an annulus from the simulations, for annuli spanning 2–10 M for each of the three spins. The distributions are normalized against and centered on the median temperature of each annulus. The distributions maintain a similar profile over all radii, with a characteristic half-width half-max $\delta T/T \sim 0.1$. Different annuli are represented by different colors.

(A color version of this figure is available in the online journal.)

fluctuation, $\delta F(r)$, (2) the orbital frequency $\Omega(r)$, and (3) the observer’s inclination. The observer’s inclination is known and the orbital frequency is easily determined analytically, but $\delta F(r)$ is not as obvious.

In the simulations, we have found that fluctuations in temperature within an annulus are approximately normally distributed, with a half-width half-max $\delta T/T \equiv \delta_T \sim 0.1$ which is consistent across all radii. (See Figure 6.) This means that the typical size of fluctuations in a given annulus scales with the typical flux at that radius. Given an orbital frequency $\Omega(r)$ and temperature profile $T(r)$, we can therefore estimate the Doppler power as a function of radius for an observer at any inclination.

In the absence of any lensing effects (an approximation which is obviously crude, but which simplifies the analysis), the total flux for an observer at an angle θ_i to the disk is

$$F(E) = \frac{\cos \theta_i}{D^2} \int r dr \int I(E, r, \phi, t) d\phi, \quad (4)$$

where D is the distance to the observer and $I(E, r, \phi, t)$ is the intensity emitted at the observer’s photon energy E at the location (r, ϕ) on the disk midplane at time t . Since we have determined from the simulations that the typical scale of fluctuations is independent of radius, we treat I as the sum of constant and time-varying parts

$$I(E, r, \phi, t) = I(E, r)[1 + \delta_I \cos(\phi - \Omega t)], \quad (5)$$

where the second term represents a fluctuation of scale δ_I which orbits in the disk with frequency $\Omega(r)$.

To apply a known disk emission profile I_e (all subscripts e indicate the frame of an observer comoving with the fluid at the midplane), we must convert between the frames of our stationary, distant observer and our comoving, midplane observer. The photon energies in the two frames, E and E_e , are related as follows,

$$\frac{E_e}{E} = \frac{h\nu_e}{h\nu} = 1 + z, \quad (6)$$

where h is the Planck constant, ν and ν_e are the frequencies of the photon in the two frames, and $1 + z$ is the redshift factor. Then since I/ν^3 is invariant,

$$I(E, r) = I_e(E_e, r) \left(\frac{E}{E_e} \right)^3 = \frac{I_e[(1+z)E, r]}{(1+z)^3}. \quad (7)$$

To return to the distant observer’s energy E we take a first-order Taylor expansion of the emission profile so that

$$I(E, r) = \left[I_e(E, r) + \frac{dI_e}{dE} \Big|_E zE \right] \frac{1}{(1+z)^3}. \quad (8)$$

Figure 7. Upper panels: bolometric power spectra from the simulations for each of the three spins, at several inclinations. Middle panels: estimate of the power from Doppler boosting of intrinsic fluctuations, P_{Dopp} , for each of the spins and inclinations in the first row. (For an observer at 0° inclination, there is no Doppler power.) Lower panels: sum of the simulation power spectra at 0° inclination and the estimated additional power from Doppler boosting for inclined observers. Note the good agreement of these semi-analytic power spectra with the simulation results shown in the upper panels.

(A color version of this figure is available in the online journal.)

The redshift factor $1 + z$, which includes both gravitational redshift and Doppler shifts, is calculated for the Kerr metric as a function of r , ϕ , a_* , and θ_i in Appendix B.

Combining Equations (4)–(8), we find that the time-varying component of the flux F_{var} changes with r as

$$\frac{dF_{\text{var}}}{dr} \propto r \int d\phi \left[I_e(E, r) + \frac{dI_e}{dE} \Big|_E zE \right] \frac{\cos[\phi - \Omega(r)t]}{(1+z)^3}. \quad (9)$$

Given a choice of I_e , this integral is fully specified and may be evaluated numerically. In the following analysis, we use a blackbody, $I_e(E, r) = B[E, T(r)]$, so that all that remains is a choice of temperature profile. For direct comparison with simulation results, we use the temperature profiles recovered from the simulation, but any reasonable profile (e.g., a Novikov–Thorne profile) will provide similar results.

The variable part of the flux $F_{\text{var}}(t)$ at a radius r follows a sinusoidal pattern as the fluctuation travels around the disk, providing power at the orbital frequency $\Omega(r)$, with the amount of power scaling with the amplitude of the sinusoid. The relevant quantity for comparison with the power spectrum $\nu P(\nu)$ is the amplitude of $(dF_{\text{var}}/d \ln \Omega)^2$, which can be derived from dF_{var}/dr as

$$\frac{dF_{\text{var}}}{d \ln \Omega} = \Omega \frac{dF_{\text{var}}}{dr} \frac{dr}{d\Omega}. \quad (10)$$

Finally, we note that in the simulation, the intrinsic variability present in the disk is more complicated than the simple cosine assumed above. The fluctuations are not constant nor are they

of a single size; rather, they have a range of sizes as shown in Figure 6 and appear and decay over many timescales. Thus, for a face-on observer, this intrinsic variability appears in the power spectrum not as a single delta function, but as a power law over a broad range of frequencies. For inclined observers, this variability will be modulated (or multiplied) by the boosts supplied by the Doppler effect, and in Fourier space, this multiplication becomes a convolution. Before making the comparison with simulations, therefore, we convolve our estimate of the Doppler power with a power law extending to higher frequencies to produce a final estimate of the shape of the Doppler feature P_{Dopp} :

$$P_{\text{Dopp}} = \left(\frac{dF_{\text{var}}}{d \ln \Omega} \right)^2 \otimes f^{-\alpha} \quad (11)$$

where α is the power law exponent at high frequencies that we fit to the simulated power spectrum for a face-on observer.

The result of this estimation can be seen in Figure 7. In each column, the top panel shows the power spectrum calculated directly from the simulation. In the bottom two sets of panels, we try to replicate the behavior of the power spectra analytically as described above—the central panels show our estimate of the Doppler feature P_{Dopp} while the bottom panels combine this estimate with the power spectrum for a face-on observer. P_{Dopp} peaks at frequencies consistent with the breaks that appear in the simulations. Because of the strong dependence on redshift $1 + z$ (which is a function of inclination θ_i) it is stronger at higher inclination, and at high spin the location of the peak changes

Figure 8. Additional power at the location of the peak of the Doppler feature for an inclined observer, normalized against the power for a face-on observer. Observers at higher inclinations have a higher fraction of the orbital motion along the line of sight, and observe more variability power because of the Doppler effect.

slightly with inclination as well. The top and bottom sets of panels are in good agreement, indicating that we have identified the origin of the high-frequency feature. In the following section, we discuss how this may apply to observations.

4. DISCUSSION

Observations of accreting black holes are typically sorted into “states,” determined by their X-ray flux and spectral hardness. The canonical states are the “thermal,” “hard,” “quiescent,” and “steep power law” states (Remillard & McClintock 2006), though there exist many intermediate forms as well (Klein-Wolt & van der Klis 2008). Our simulated disks most closely correspond to accreting black holes in the thermal or steep power law state, as they are dominated by thin accretion disks at high accretion rates ($\approx 70\% \dot{M}_{\text{Edd}}$) and thus high luminosities. In all states, the power spectra can be described by the superposition of several Lorentzian components (Belloni et al. 2002), bounded by a low-frequency component at one end and two high-frequency components on the other. QPOs are represented as Lorentzians with a high quality factor Q and are sharply peaked.

Because we have only 40,000 M (2 s for a $10 M_{\odot}$ black hole) of simulation data at best, we can only compare to the high-frequency end of the observed power spectra. In particular, our highest-frequency feature ought to correspond to the highest-frequency component for observed black hole binaries. We have seen in Section 3 that the power spectra from our simulated disks have a feature at high frequencies which is induced by the orbital motion of structures in the disk and highlighted by the Doppler effect. The strength and location of this Doppler feature is affected by four quantities: the inclination of the observer, the energy band used by the observer, the black hole mass, and the black hole spin.

As discussed in Section 3, the observer’s inclination with respect to the disk as well as the energy band chosen by the observer each affect the strength of the Doppler feature in the power spectrum. As the feature is a direct result of the Doppler effect, a more inclined observer has a larger component of the

Figure 9. Frequency at which the estimated shape of the Doppler feature, P_{Dopp} , peaks as a function of spin for several inclinations. Solid lines correspond to the temperature profiles from the simulations, and dashed lines to Novikov–Thorne profiles (Page & Thorne 1974). Regardless of the specific $T(r)$, the peak shifts to higher frequencies as black hole spin increases. The orbital frequencies at the ISCO and at the peak of $dL/d\ln r$ (dotted lines) are plotted for comparison. Note that the peak of P_{Dopp} always lies between the two, indicating that its location is a compromise between the strength of the Doppler effect and the luminosity profile of the disk.

(A color version of this figure is available in the online journal.)

fluid velocity along the line of sight, a higher Doppler factor, and thus a stronger power spectral feature (Figure 8). Moving to high energies also enhances the feature: past the peak of the blackbody, the radiation spectrum becomes steeper so that the shifts in frequency provided by the Doppler effect produce a larger shift in intensity than in bands where the radiation spectrum is shallower (Figure 5).

We have shown in Figure 7 that the location of the Doppler feature agrees with the peak of the function $P_{\text{Dopp}}(\nu)$ derived in Section 3.2, which estimates the amplitude of Doppler modulations of hotspots produced by the temperature fluctuations present in the disk and their contribution to the power at orbital frequencies $\Omega(r)$. The shape of this function depends on the spin of the black hole, a_* , as well as the temperature profile of the disk, $T(r)$. At higher spins, the inner edge of the accretion disk extends closer to the black hole, providing access to the higher orbital frequencies at those radii and shifting the peak of P_{Dopp} to higher frequencies. This effect is present regardless of the temperature profile chosen. Figure 9 shows the location of the peak for the temperature profiles measured from the simulation $T_{\text{sim}}(r)$ as well as an analytical $T_{\text{NT}}(r)$ from Page & Thorne (1974), and both choices of temperature profile show the peak moving to higher frequencies with increasing black-hole spin as the accretion disk extends to smaller radii.

Observationally, the location of the highest-frequency Lorentzian component has been seen to shift to higher frequencies as accreting black holes undergo state transitions from low to high luminosity (Klein-Wolt & van der Klis 2008), which is consistent with the idea that it represents some dominant radius that moves closer to the black hole as luminosity increases and the truncation radius marking the transition between the thin and thick disks moves inward. At the highest luminosities where the accretion disk is expected to extend all the way to the ISCO, therefore, one could potentially use the location of this high-frequency feature to estimate spin in cases where the black hole mass is known.

One interesting implication of the results shown in Figure 9 is that the variation of the Doppler feature frequency with black hole spin a_* , as derived from the simulations (solid lines), is fairly shallow. This is very different from the much steeper variation predicted by the Novikov–Thorne model (dashed lines) or the orbital frequency at the ISCO, $\Omega_K(r_{\text{isco}})$. Models of high-frequency QPOs based on dynamical frequencies in the disk (Psaltis 2004) typically predict steep variations akin to the latter. If the Doppler feature described in this paper is relevant to observed QPOs, then it is clear that simulation data will play an increasingly important role in interpreting observations.

The thin accretion disk simulations presented in this paper represent the state-of-the-art. Nevertheless, they omit some important physics. Notably, the code we used does not include radiative transfer. Instead, radiative cooling is handled via a toy prescription (Shafee et al. 2008; Penna et al. 2010) which is designed to keep the disk geometrically thin. New codes are currently being developed that can handle radiation GRMHD (Sądowski et al. 2013, 2014). Simulations in the future with these codes will provide more reliable profiles of $T_{\text{sim}}(r)$ and hence more accurate estimates of the Doppler feature frequency.

The origin of the QPO which appears in our power spectra at lower frequencies is as yet undetermined. It moves to slightly lower frequency as the spin of the black hole increases, which is somewhat unexpected. It may be that the QPO is an artifact of the simulation, e.g., some characteristic frequency of the torus with which the simulations are initialized, but this remains to be seen. It is unlikely that the QPO corresponds to any real mode present in the disk, since the MRI tends to damp out these modes and there is no obvious mechanism to replenish them in the simulations.

5. CONCLUSION

We have shown that our simulated GRMHD accretion disks have some of the same variability properties observed in black hole binaries. These variability properties, in particular the feature observed at high frequencies, are a function of the physical properties of the system such as black-hole mass and spin.

The high-frequency feature originates from the orbital motion in the disk. Hotspots that survive for at least one orbital period are alternately Doppler boosted and dimmed for an inclined observer, producing periodicity and hence additional power at orbital frequencies. This Doppler feature peaks or breaks at the orbital frequency of the radius whose Doppler fluctuations are dominant, which is a function of spin (and, to a lesser degree, inclination). This feature is stronger at higher inclinations, where the Doppler effect is stronger, and at higher energies past the peak of the blackbody spectrum, where the SED is steep.

Thus, given the disk temperature profile (e.g., from simulations), an observer could potentially use the location of the highest-frequency Lorentzian component in the power spectrum to determine the spin of an accreting black hole with known mass. This procedure would be most plausible in the bright thermal or steep power law states, where the thin disk component is believed to extend all the way down to the ISCO.

While we were successful in reproducing some of the observed characteristics of accreting black holes and understanding the origin of the highest-frequency feature, our understanding of black hole binary variability is far from complete. In particular, the origin of the low-frequency QPO which appears in our power spectra is still unknown. It is not clear whether it is a consequence of some physical process or an artifact of the simulation. Varying the parameters of the simulation itself

(e.g., the location of the initial torus) may reveal whether it is a physical or computational effect. Additionally, increasing the length of the simulations would provide more low-frequency information in the power spectra and help to better characterize this feature.

We will also aim to better describe these accreting black hole systems in the future by introducing more realistic physics into the simulations. The observed profile of emission would be significantly improved by incorporating full radiative transfer into the simulations, rather than employing our current method of using an ad-hoc cooling prescription and doing radiative transfer in post-processing. This step is crucial to fully understanding the radiative processes at work. These goals are very ambitious, but would represent a significant improvement in our ability to faithfully simulate accreting black holes.

S.W. is supported by the National Science Foundation Graduate Research Fellowship under grant number DGE1144152. Resources supporting this work were provided by the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575, and the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center. This material is based upon work supported by National Science Foundation grant number AST1312651 to R.N. D.P. was partially supported by NSF CAREER award AST-0746549.

APPENDIX A

TEST WITH NOVIKOV–THORNE DISKS

To test the raytracing code, we wish to see whether it can accurately recover all of the energy that is emitted at the disk midplane using distant observers scattered around the sky. For an accretion disk with a Novikov–Thorne flux profile $F(r)$ (Novikov & Thorne 1973; Page & Thorne 1974), we know analytically that the luminosity per radius as observed at infinity is given by

$$\frac{dL}{dr} = 4\pi r u_r(r) F(r), \quad (\text{A1})$$

where the factor u_r , i.e., the covariant time component of the four-velocity, accounts for gravitational redshift. The goal of this Appendix is to reproduce this analytical profile numerically by raytracing back to a Novikov–Thorne disk to retrieve the derivative dL/dr at a particular location (θ, ϕ) on the sky, then integrating over observers at all locations on the sky to retrieve the universal derivative dL/dr . This is a comprehensive test of our raytracing code.

To compute the derivative dL/dr for a single observer, we must integrate over the observer's image plane. In our raytracing code, the image plane is an elliptical grid in radial and angular coordinates (b, β) which is squeezed by a factor of $\cos \theta_i$. The total flux in the image plane at that inclination is therefore

$$F \propto \cos \theta_i \int I(b, \beta) b db d\beta \propto \cos \theta_i \int I(b, \beta) b^2 d(\log b) d\beta, \quad (\text{A2})$$

where $I(b, \beta)$ is the intensity at infinity of the ray through the point with coordinates (b, β) as recovered from the accretion disk by raytracing. Each angle β has its own flux profile along the image plane radius b ,

$$\frac{dF}{d \log b}(b, \beta) \propto \cos \theta_i I(b, \beta) b^2 d\beta, \quad (\text{A3})$$

Figure 10. Left: the fraction of energy emitted at a radius r which escapes to infinity (f_{esc}), returns to the midplane (f_{ret}), or is swallowed by the black hole (f_{sw1}) for three disks with Novikov–Thorne orbital velocity profiles. The lines terminate at the photon orbit, where the orbital velocity is no longer defined, for each spin. The locations of the innermost stable circular orbits are marked with a circle. Right: luminosity profiles of Novikov–Thorne disks, computed from the flux profile at the surface both analytically (solid lines) and by raytracing (points).

(A color version of this figure is available in the online journal.)

or, converting between image plane radius b and disk plane radius r ,

$$\frac{dF}{dr}(r, \beta) \propto \cos \theta_i I[b(r, \beta), \beta] b(r)^2 \frac{d \log b}{dr} d\beta. \quad (\text{A4})$$

(Note that $d \log b / dr$ can be different for each β .) Then the radial profile over the whole image plane, summed over all angles, is

$$\frac{dF}{dr}(r) \propto \cos \theta_i \sum_{\beta} I[b(r, \beta), \beta] b(r)^2 \frac{d \log b}{dr}. \quad (\text{A5})$$

We calculate the derivative $d \log b / dr$ as a first-order numerical derivative along each angular coordinate β , and interpolate the $dF/dr(r, \beta)$ derivative to a common set of radii before summation.

Then to integrate over the full 4π sphere, we perform raytracing for observers at $\theta_i = [2^\circ.5, 7^\circ.5, 12^\circ.5, \dots, 90^\circ]$, so that the total luminosity of the disk is

$$L \propto \int F(\theta_i) \sin \theta_i d\theta_i d\phi \propto \sum_{\theta_i} F(\theta_i) \sin \theta_i, \quad (\text{A6})$$

where $F(\theta_i)$ is the flux observed at an inclination θ_i . For each of these inclinations, we find the flux profile dF/dr as described above; then the total observed luminosity profile is

$$\frac{dL}{dr} \propto \sum_{\theta_i} \frac{dF}{dr}(r, \theta_i) \sin \theta_i. \quad (\text{A7})$$

One additional correction is needed before we can expect our raytraced dL/dr to match the analytical profile. The analytical profile takes the flux produced by the disk and transforms it for an observer at infinity; it does not account for the fact that some photons never reach infinity but are rather swallowed by the black hole or fall back onto the disk. These photons will be “missing” from our raytracing solution. To account for this, we find the fraction of energy that escapes to infinity by running the raytracing code in the reverse direction, starting on the disk and following the paths of photons outward.

Beginning in the rest frame of the fluid at a radius r , we distribute rays at equal intervals in $\Delta\phi_{\text{com}}$ and $\Delta\theta_{\text{com}}$ (fluid-frame coordinates centered at r) and assign them energy proportional to their solid angle $d\Omega = \sin \theta_{\text{com}} d\theta_{\text{com}} d\phi_{\text{com}}$ times a factor of $\cos \theta_{\text{lab}}$ since we are representing a surface along the midplane. These rays are then transformed to disk coordinates in the lab frame assuming a Novikov–Thorne orbital velocity profile and integrated outward along null geodesics.

Raytracing is halted when one of three conditions is met: (1) the ray reaches the event horizon, (2) the ray passes through the midplane, or (3) the ray reaches some large radius and escapes. We assign the energy carried by the rays to these three bins and divide by the total energy to get the three fractions $f_{\text{sw1}}(r)$, $f_{\text{ret}}(r)$, and $f_{\text{esc}}(r)$ which describe the fraction of energy produced at a radius r which is swallowed by the black hole, returns to the disk midplane, or escapes to infinity respectively. These fractions can be seen for disks with spin parameters $a_* = 0, 0.7$, and 0.9 (the same spins as our GRMHD simulations) in Figure 10.

Rays which meet the third condition and escape to infinity are the only ones accessible by the reverse-raytracing method. Thus, we expect that the analytical and raytraced derivatives dL/dr should differ by a factor of $f_{\text{esc}}(r)$ so that

$$\frac{dL}{dr} \propto \frac{1}{f_{\text{esc}}(r)} \sum_{\theta_i} \frac{dF}{dr}(r, \theta_i) \sin \theta_i. \quad (\text{A8})$$

This quantity is plotted in Figure 10, along with the analytical dL/dr . A comparison of the two shows very good agreement, which demonstrates that our method captures the emission produced at the disk surface with excellent accuracy.

APPENDIX B

DERIVATION OF REDSHIFT IN KERR METRIC

The redshift we seek to derive is the shift in frequency,

$$1 + z = \frac{\nu_{\text{em}}}{\nu_{\text{obs}}} = \frac{(g_{\mu\nu} u^\mu k^\nu)_{\text{em}}}{(g_{\mu\nu} u^\mu k^\nu)_{\text{obs}}}, \quad (\text{B1})$$

for a photon with wave vector k^μ between an observer corotating with the disk with some circular four-velocity $u_{\text{em}}^\mu =$

$(u^t, 0, 0, u^\phi)$, and a distant observer with $u_{\text{obs}}^\mu = (1, 0, 0, 0)$. For a spinning black hole, the metric tensor $g_{\mu\nu}$ is given by the Kerr metric:

$$ds^2 = -\left(1 - \frac{2Mr}{\rho^2}\right)dt^2 - \frac{2Mar \sin^2 \theta}{\rho^2}(dt d\phi + d\phi dt) + \frac{\rho^2}{\Delta}dr^2 + \rho^2 d\theta^2 + \frac{\sin^2 \theta}{\rho^2}[(r^2 + a^2)^2 - a^2 \Delta \sin^2 \theta]d\phi^2 \quad (\text{B2})$$

where (t, r, θ, ϕ) are the Boyer–Lindquist coordinates, $a = J/M$ is the angular momentum per unit mass of the black hole, M is the black hole mass, and Δ and ρ are given by

$$\Delta = r^2 - 2Mr + a^2, \quad (\text{B3})$$

$$\rho^2 = r^2 + a^2 \cos^2 \theta. \quad (\text{B4})$$

After expanding the sums in Equation (B1) using the Kerr metric, the redshift is

$$1 + z = \frac{(g_{tt}u^t k^t + g_{\phi\phi}u^\phi k^\phi + g_{t\phi}u^t k^\phi + g_{\phi t}u^\phi k^t)_{\text{em}}}{(g_{tt}k^t + g_{t\phi}k^\phi)_{\text{obs}}}. \quad (\text{B5})$$

This can be simplified by noting that the Kerr spacetime has two Killing vectors, $\xi^\mu = (1, 0, 0, 0)$ and $\eta^\mu = (0, 0, 0, 1)$, leading to two conserved quantities

$$\varepsilon_p = -g_{\mu\nu}k^\mu \xi^\nu = -g_{tt}k^t - g_{t\phi}k^\phi, \quad (\text{B6})$$

$$l_p = g_{\mu\nu}k^\mu \eta^\nu = g_{\phi t}k^t + g_{\phi\phi}k^\phi, \quad (\text{B7})$$

corresponding to energy and angular momentum respectively. Then the expression for redshift reduces to

$$1 + z = u^t - u^\phi \frac{l_p}{\varepsilon_p} \quad (\text{B8})$$

where all terms now correspond to the observer at the midplane.

Making the approximation that the trajectories of photons are not affected by the curved spacetime, the spatial wave vector $\mathbf{k}$ from the midplane to an observer at inclination θ_i is

$$\mathbf{k} = |k|(\sin \theta_i \sin \phi \hat{r} + \cos \theta_i \hat{\theta} + \sin \theta_i \cos \phi \hat{\phi}) \quad (\text{B9})$$

in spherical polar coordinates. The photon four-momentum $p^\mu = \hbar k^\mu$ has the property $p_\mu p^\mu = 0$. Thus it is also true that

$$k_\mu k^\mu = g_{\mu\nu}k^\mu k^\nu = g_{tt}k^t k^t + g_{rr}k^r k^r + g_{\theta\theta}k^\theta k^\theta + g_{\phi\phi}k^\phi k^\phi + 2g_{t\phi}k^t k^\phi = 0. \quad (\text{B10})$$

We can remove k^t from this expression by substituting $k^t = -(\varepsilon_p + g_{t\phi}k^\phi)/g_{tt}$ from the energy equation; then multiplying through by g_{tt} yields

$$(\varepsilon_p + g_{t\phi}k^\phi)^2 + g_{tt}g_{rr}k^r k^r + g_{tt}g_{\theta\theta}k^\theta k^\theta + g_{tt}g_{\phi\phi}k^\phi k^\phi - 2g_{t\phi}(\varepsilon_p + g_{t\phi}k^\phi)k^\phi = 0. \quad (\text{B11})$$

The wave vector components k^μ are given by $k^\mu = \mathbf{k} \cdot \hat{e}^\mu$, where $\hat{e}^\mu$ are the dual basis vectors for the Kerr metric. Once the k^μ are inserted, the expression has a solution of the form $|k| = \kappa \varepsilon_p$,

where κ is a function of r , ϕ , a_* , and θ_i . The ratio l_p/ε_p is, therefore, given by

$$\frac{l_p}{\varepsilon_p} = \frac{g_{\phi t}k^t + g_{\phi\phi}k^\phi}{\varepsilon_p} \quad (\text{B12})$$

$$= -\frac{g_{t\phi}}{g_{tt}}(1 + g_{t\phi}\kappa \hat{e}^{\phi\phi} \sin \theta_i \cos \phi) + g_{\phi\phi}\kappa \hat{e}^{\phi\phi} \sin \theta_i \cos \phi \quad (\text{B13})$$

$$= \kappa \hat{e}^{\phi\phi} \sin \theta_i \cos \phi \left[g_{\phi\phi} - \frac{(g_{t\phi})^2}{g_{tt}} \right] - \frac{g_{t\phi}}{g_{tt}}, \quad (\text{B14})$$

and the final expression for redshift is

$$1 + z = \left[u^t + u^\phi \frac{g_{t\phi}}{g_{tt}} \right] - u^\phi \kappa \hat{e}^{\phi\phi} \sin \theta_i \cos \phi \left[g_{\phi\phi} - \frac{(g_{t\phi})^2}{g_{tt}} \right] \quad (\text{B15})$$

$$= \frac{u_t}{g_{tt}} - u^\phi \kappa \hat{e}^{\phi\phi} \sin \theta_i \cos \phi \left[g_{\phi\phi} - \frac{(g_{t\phi})^2}{g_{tt}} \right]. \quad (\text{B16})$$

REFERENCES

- Abramowicz, M. A., Chen, X., & Taam, R. E. 1995, *ApJ*, **452**, 379
 Abramowicz, M. A., Karas, V., Kluzniak, W., Lee, W. H., & Rebusco, P. 2003, *PASJ*, **55**, 467
 Balbus, S. A., & Hawley, J. F. 1991, *ApJ*, **376**, 214
 Belloni, T., Psaltis, D., & van der Klis, M. 2002, *ApJ*, **572**, 392
 Davis, S. W., Blaes, O. M., Hubeny, I., & Turner, N. J. 2005, *ApJ*, **621**, 372
 Davis, S. W., & Hubeny, I. 2006, *ApJS*, **164**, 530
 Dexter, J., Agol, E., Fragile, P. C., & McKinney, J. C. 2010, *ApJ*, **717**, 1092
 Done, C., Gierliński, M., & Kubota, A. 2007, *A&ARv*, **15**, 1
 Esin, A. A., McClintock, J. E., & Narayan, R. 1997, *ApJ*, **489**, 865
 Gammie, C. F., McKinney, J. C., & Tóth, G. 2003, *ApJ*, **589**, 444
 Hubeny, I., Blaes, O., Krolik, J. H., & Agol, E. 2001, *ApJ*, **559**, 680
 Hubeny, I., & Lanz, T. 1995, *ApJ*, **439**, 875
 Ingram, A., & Done, C. 2011, *MNRAS*, **415**, 2323
 Kato, S. 2001, *PASJ*, **53**, L37
 Klein-Wolt, M., & van der Klis, M. 2008, *ApJ*, **675**, 1407
 Krolik, J. H., & Hawley, J. F. 2002, *ApJ*, **573**, 754
 Kulkarni, A. K., Penna, R. F., Shcherbakov, R. V., et al. 2011, *MNRAS*, **414**, 1183
 Marković, D., & Lamb, F. K. 1998, *ApJ*, **507**, 316
 McKinney, J. C. 2006, *MNRAS*, **368**, 1561
 McKinney, J. C., & Blandford, R. D. 2009, *MNRAS*, **394**, L126
 Narayan, R., & McClintock, J. E. 2008, *NewAR*, **51**, 733
 Narayan, R., & Yi, I. 1994, *ApJL*, **428**, L13
 Narayan, R., & Yi, I. 1995, *ApJ*, **452**, 710
 Novikov, I. D., & Thorne, K. S. 1973, in *Black Holes (Les Astres Occlus)*, ed. C. Dewitt & B. S. Dewitt (Paris: Gordon & Breach), 343
 Nowak, M. A., & Wagoner, R. V. 1993, *ApJ*, **418**, 187
 Page, D. N., & Thorne, K. S. 1974, *ApJ*, **191**, 499
 Pecháček, T., Karas, V., & Czerny, B. 2008, *A&A*, **487**, 815
 Penna, R. F., McKinney, J. C., Narayan, R., et al. 2010, *MNRAS*, **408**, 752
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, *Numerical Recipes: The Art of Scientific Computing* (3rd ed.; New York: Cambridge Univ. Press)
 Psaltis, D. 2004, in *AIP Conf. Proc.* 714, X-ray Timing 2003: Rossi and Beyond, ed. P. Kaaret, F. K. Lamb, & J. H. Swank (Melville, NY: AIP), 29
 Remillard, R. A., & McClintock, J. E. 2006, *ARA&A*, **44**, 49
 Sądowski, A., Narayan, R., McKinney, J. C., & Tchekhovskoy, A. 2014, *MNRAS*, **439**, 503
 Sądowski, A., Narayan, R., Tchekhovskoy, A., & Zhu, Y. 2013, *MNRAS*, **429**, 3533
 Shafee, R., McKinney, J. C., Narayan, R., et al. 2008, *ApJL*, **687**, L25
 Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, **24**, 337
 Silbergleit, A. S., Wagoner, R. V., & Ortega-Rodríguez, M. 2001, *ApJ*, **548**, 335
 Stella, L., & Vietri, M. 1998, *ApJL*, **492**, L59
 van der Klis, M. 2006, *Rapid X-ray Variability*, ed. W. H. G. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 39
 Yuan, F., & Narayan, R. 2014, *ARA&A*, in press
 Zhu, Y., Davis, S. W., Narayan, R., et al. 2012, *MNRAS*, **424**, 2504