The Astrophysical Journal, 873:71 (12pp), 2019 March 1 https://doi.org/10.3847/1538-4357/ab05d5

Predicting the X-Ray Spectra of Stellar-mass Black Holes from Simulations

Brooks E. Kinch[1], Jeremy D. Schnittman[2], Timothy R. Kallman[2], and Julian H. Krolik[1]
[1] Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA; kinch@jhu.edu, jhk@jhu.edu
[2] NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA; jeremy.d.schnittman@nasa.gov, timothy.r.kallman@nasa.gov
Received 2018 November 7; revised 2019 February 6; accepted 2019 February 9; published 2019 March 6

Abstract

We describe results from a new technique for the prediction of complete, self-consistent X-ray spectra from three-dimensional general relativistic magnetohydrodynamic (GRMHD) simulations of black hole accretion flows. Density and cooling rate data from a HARM3D GRMHD simulation are post-processed by an improved version of the Monte Carlo radiation transport code PANDURATA (in the corona) and the Feautrier solver PTRANSX (in the disk), with XSTAR subroutines. The codes are run in a sequential, iterative fashion to achieve globally energy-conserving and self-consistent radiation fields, temperature maps, and photoionization equilibria. The output is the X-ray spectrum as seen by a distant observer, including features, such as the Fe Kα emission line and corresponding K-edge absorption trough, due to disk reprocessing of coronal power. For the example cases we consider—a non-rotating 10 $M_{\odot}$ black hole with solar abundances, accreting at 0.01, 0.03, 0.1, or 0.3 Eddington—we find spectra resembling actual observations of stellar-mass black holes in the soft or steep power-law state: broad thermal peaks (at 1–3 keV), steep power laws extending to high energy ($\Gamma = 2.7$–4.5), and prominent, asymmetric Fe Kα emission lines with equivalent widths in the range 40–400 eV (larger EW at lower accretion rates). By starting with simulation data, we obviate the need for parameterized descriptions of the accretion flow geometry—no a priori specification of the corona's shape or flux, or the disk temperature or density, etc., is needed. Instead, we apply the relevant physical principles to simulation output using appropriate numerical techniques; this procedure allows us to calculate inclination-dependent spectra after choosing only a small number of physically meaningful parameters: black hole mass and spin, accretion rate, and elemental abundances.

Key words: accretion, accretion disks – black hole physics – line: formation – radiative transfer – X-rays: binaries

1. Introduction

Accreting black holes provide a unique opportunity to investigate both the strong-field regime of general relativity and the physics of accretion flows in extreme environments. Both active galactic nuclei (AGNs) and stellar-mass black hole binaries produce X-ray spectra with line and continuum features that convey information about the environment and spacetime geometry from which they originate. Relativistically broadened Fe Kα fluorescence lines are one of the key indicators that these systems do in fact contain black holes (Tanaka et al. 1995), and the thermal plus power-law continuum indicates the presence of disk and corona, respectively (Liang 1979; Haardt & Maraschi 1991). Indeed, studying the governing physics of accretion processes is tied to our ability to connect the underlying theory to observations. The quantitative information inferable from any spectrum is limited, however, by the templates to which the observation is compared.

To this end, we have developed a technique with which model spectra are computed directly from simulation data by applying the relevant physical principles while invoking almost no assumptions. The numerical machinery we describe here is an extension of that introduced in Kinch et al. (2016); as in that paper, we apply our method only to the case of non-rotating, stellar-mass black holes, but we otherwise greatly expand the predictive scope of our method by treating the X-ray emission lines *and* the continuum in a self-consistent, energy-conserving fashion. In addition, we explore the effects of varying the nominal accretion rate on the predicted spectrum.

A variety of methods are currently in use for modeling the spectra of black hole systems. Some are purely phenomenological—the continuum is fit with a multicolor disk blackbody (e.g., DISKBB, Mitsuda et al. 1984) plus a (typically broken) power law at high energy, and the Fe Kα emission from the disk surface is assumed to vary as a decreasing power law (or sometimes broken power law) in radius with a hard cutoff at the innermost stable circular orbit (ISCO) and another at some outer radius (e.g., RELLINE. Dauser et al. 2013; see Reynolds & Nowak 2003 for a discussion of these methods). More sophisticated techniques model the continuum with a single-zone Comptonization region and the disk reprocessed component (the Fe Kα line, the K-edge, and the Compton bump) by performing detailed radiative transfer and photoionization calculations within a sample section of the disk (e.g., the codes REFLIONX (Ross & Fabian 2005), XILLVER (García & Kallman 2010; García et al. 2011, 2013), and RELXILL (García et al. 2014)). At present *all* methods rely in some way upon a parameterized description of the black hole environment: at best, an idealized corona (often a "lamppost" point source or a single homogeneous region) emits a power-law spectrum (perhaps with a thermal cutoff) that illuminates a semi-infinite, blackbody-radiating disk, and this disk has a knife-edge cutoff precisely at the ISCO. When using such a model to, for example, extract spin measurements from spectral data (Reynolds 2013; Miller & Miller 2015), the accuracy of the measurement is limited by the accuracy of the assumed accretion flow geometry and associated coronal flux.

By starting with 3D GRMHD simulation data, we greatly reduce the number of assumptions needed to describe the accretion flow geometry. We therefore also reduce the number of free parameters needed to specify a resulting observable spectrum. We do not, for example, require a sharp cutoff in Fe Kα emission at the ISCO, nor do we specify the coronal

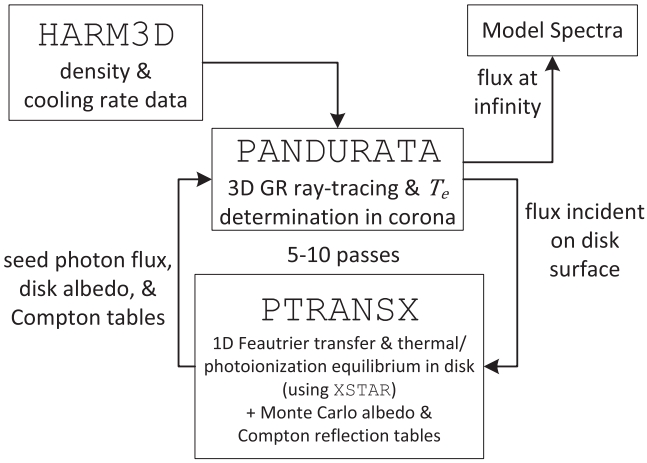

Figure 1. A schematic overview of the general procedure.

geometry (lamppost or otherwise) a priori; the density and temperature structure of the disk are not assumed in advance either. Instead, these are computed directly from the underlying physics, with the sole significant assumption being the equation of state employed by the simulation (as we describe below). As simulations improve, e.g., by the use of more realistic equations of state, our apparatus can easily be applied to their output as well. The resulting prediction of our method, the full inclination-dependent observable spectrum, is a function of a very small number of parameters, each physically meaningful: the black hole mass and spin, the nominal accretion rate, and the elemental abundances.

2. Method

Our procedure has three main components. First, an accreting black hole system is simulated using HARM3D (Noble et al. 2009). We take a three-dimensional snapshot of the fluid density, four-velocity, and dissipation (cooling) rate at a time when the simulation has achieved approximate inflow equilibrium (out to $r \sim 20\,M$). Using an injection rate of thermal seed photons computed by integrating the local dissipation rate within the disk's photosphere, the Monte Carlo radiation transport code PANDURATA (Schnittman & Krolik 2013) determines the radiation field consistent with the simulation data and thermal balance in the corona (Schnittman et al. 2013). With HARM3D's description of the disk structure and PANDURATA's calculation of the disk incident flux, PTRANSX computes the disk's reprocessed *outgoing* flux, requiring photoionization equilibrium and energy conservation everywhere within the disk (Kinch et al. 2016). This step yields a new guess for both the energy-dependent seed photon flux emerging from the disk surface and the spatial- and energy-dependent disk albedo—input for the next PANDURATA run. We cycle between PANDURATA (in the corona) and PTRANSX (in the disk) until a consistent picture of the *global* radiation field develops. This cycling is a significant improvement over the original PANDURATA method. To avoid confusion between the iterative procedures within each step and the outermost iterations between PANDURATA and PTRANSX, we refer to the latter as "passes." With one final round of relativistic ray-tracing, both reprocessed disk and coronal emission are transported to a distant observer in order to construct the complete predicted spectrum. The overall scheme is summarized in Figure 1.

2.1. Simulation Data—HARM3D

The density (ρ or n_e) and cooling rate ($\mathcal{L}$) data are from one snapshot of a HARM3D simulation, taken at a time when the disk is in a statistically steady state. HARM3D is a three-dimensional, intrinsically conservative general relativistic magnetohydrodynamic (GRMHD) code, with a cooling function designed to produce a geometrically thin disk. HARM3D solves a modified stress–energy conservation equation: in gravitationally bound gas above a target temperature T_*, the excess heat is radiated away on an orbital timescale; T_* is chosen so as to achieve a target aspect ratio (Noble et al. 2009). The specific simulation we use, "ThinHR" (Noble et al. 2011), has an aspect ratio $H_{\rm dens}/r = 0.06$ (where $H_{\rm dens}$ is the density-weighted scale height), and is still one of the best-resolved GRMHD disk simulations ever carried out (Hawley et al. 2013).

Translating the simulation data from "code units" to physical (cgs) units requires specification of the mass of the central black hole M, which sets the length scale and timescale ($1\,M = (M/M_\odot) \times 1.5 \times 10^5\,{\rm cm} = (M/M_\odot) \times 4.9 \times 10^{-6}\,{\rm s}$), and the accretion rate (in Eddington units) $\dot{m}$, which sets the scale for the density and cooling rate through (Schnittman et al. 2013)

$$\rho_{\rm cgs} = \rho_{\rm code}\frac{4\pi c^2}{\kappa G M}\frac{\dot{m}/\eta}{\dot{M}_{\rm code}}, \tag{1}$$

$$\mathcal{L}_{\rm cgs} = \mathcal{L}_{\rm code}\frac{4\pi c^7}{\kappa G^2 M^2}\frac{\dot{m}/\eta}{\dot{M}_{\rm code}}, \tag{2}$$

where $\kappa = 0.4\,{\rm cm^2\,g^{-1}}$ is the electron scattering opacity and $\eta = 0.061$ (>0.057, the value of Novikov & Thorne 1973) is the radiative efficiency found in that simulation (Noble et al. 2011). In this paper, we consider a $10\,M_\odot$ central black hole at four accretion rates, $\dot{m} = 0.01$, 0.03, 0.1, and 0.3.

With a known density structure (and a known spacetime geometry), surfaces of constant optical depth can be defined by integrating the electron scattering opacity along arcs of constant (r, ϕ), starting from the poles and continuing until the desired optical depth is reached. With the disk lying in the x–y plane, the collection of points $\Theta(r, \phi)$ that satisfy

$$\int_0^{\Theta(r,\phi)} \sigma_T n_e \sqrt{g_{\theta\theta}}\,d\theta = \tau \text{ or } \int_{\Theta(r,\phi)}^{\pi} \sigma_T n_e \sqrt{g_{\theta\theta}}\,d\theta = \tau \tag{3}$$

(where θ is the polar angle) defines the (upper and lower) surfaces of constant optical depth for the given value of τ. The natural choice of surfaces with which to divide the disk body from the corona are the $\tau = 1$ surfaces, which we call the disk photospheres and label $\Theta_{\rm top}$ and $\Theta_{\rm bot}$. At any given (r, ϕ), the region between these surfaces—if they exist—is the disk body; everywhere else is the corona. Which τ value to use for dividing the disk and corona is somewhat arbitrary. For our purposes, a division is needed such that the only significant cooling process in the corona is inverse Compton (IC) scattering, while all atomic processes (such as Fe Kα production) occur within the disk. Because the maximum local ratio of free–free power to net IC power that we compute post hoc in the corona is $\lesssim$4% (just outside $\tau = 1$), and we find significant Fe Kα production limited typically to $\tau > 1.5$, the $\tau = 1$ surface is a satisfactory choice for the photosphere.

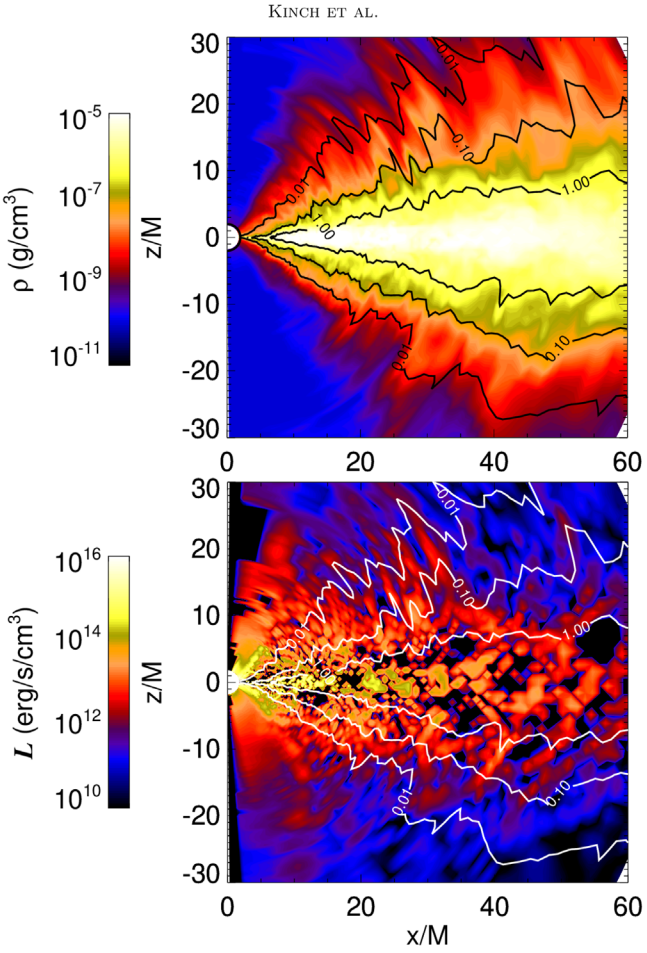

Figure 2. x–z slice of the HARM3D simulation data we use in this paper, scaled to $M = 10\,M_\odot$, $\dot{m} = 0.01$, with three surfaces of constant optical depth overlaid. Note the great difference (and rapid change) in density between the disk body and the corona (upper panel). The lower panel shows the local instantaneous cooling rate—not all fluid elements are radiatively cooling at each time step (see Schnittman et al. 2013 for a more extensive discussion of the cooling function). There is significant cooling in the corona, even where the density is very low.

Figure 2 shows a cross section of HARM3D density and cooling data, scaled to $10\,M_\odot$ and $\dot{m} = 0.01$, with several surfaces of constant optical depth overlaid.

2.2. Coronal Radiation Field—PANDURATA

For the initial pass, we assume that the total cooling at one (r, ϕ) arc within the disk body is radiated thermally at its photospheres. That is, at any (r, ϕ) for which the $\tau = 1$ surfaces exist, the flux outward at both $\Theta_{\rm top}$ and $\Theta_{\rm bot}$ is described as a hardened blackbody with effective temperature

$$\int_{\Theta_{\rm top}}^{\Theta_{\rm bot}} \mathcal{L}\sqrt{g_{\theta\theta}}\,d\theta = 2\sigma T_{\rm eff}^4. \quad (4)$$

These thermal seed photons are ray-traced through the corona by the Monte Carlo radiation transport and local temperature balance code PANDURATA. For all subsequent passes, PANDURATA uses PTRANSX's output seed photon spectra instead of the hardened blackbody. PANDURATA takes as input the density and cooling maps from HARM3D, as well as the seed photon emission at the disk photosphere, and outputs: (1) an electron temperature map of the corona; (2) the spectrum as seen by distant observers; and (3) the spectral shape and strength of the flux incident upon the (upper and lower) disk photospheres at each (r, ϕ).

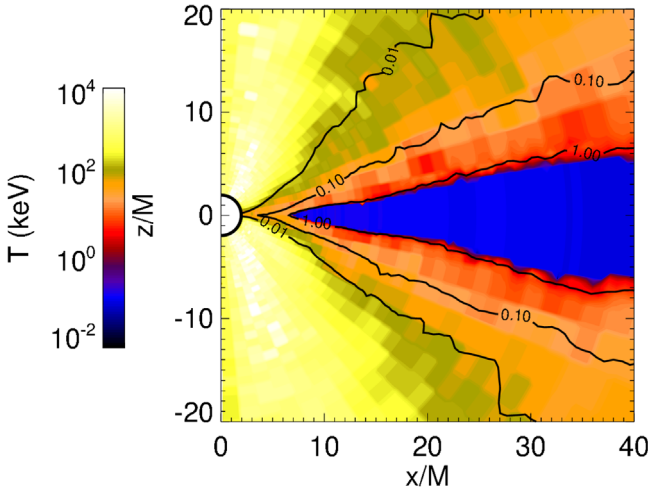

Figure 3. Electron temperature map produced by PANDURATA for the same slice of data shown in Figure 2 ($M = 10\,M_\odot$, $\dot{m} = 0.01$). Note that within the disk body (between the $\tau = 1$ surfaces) the electron temperature is shown as the constant value $T_{\rm eff}$; PANDURATA does not determine the electron temperature within the disk body. The visible "blockiness" is due to our division of the coronal volume into sectors.

The operation of PANDURATA's original version is described in detail, including series of tests to demonstrate the algorithm's validity, in Schnittman & Krolik (2013). In brief, the code simulates the trajectories and scattering of seed photons in the corona while solving for the electron temperature at each point in the corona by setting the net IC power equal to HARM3D's local cooling rate. Several modifications to PANDURATA were made for its use in this project. First, photon packet scattering off single electrons was replaced by an *ensemble* approach—when a photon packet scatters in the corona, the photon packet's spectrum is redistributed according to an angle-averaged energy redistribution function described below (see Section 2.3.2); its new direction, however, is determined as if it were a single photon scattering off a single electron whose particular velocity was selected from the Maxwell–Jüttner distribution. Second, the coronal volume is divided into sectors—a coarser grouping, compared to the underlying simulation grid, of ∼100 contiguous grid cells each—with the interior of each sector treated as having a single temperature; net IC power is assessed for the sector as a whole, and a sector's temperature is adjusted by way of a Newton–Raphson method until its net IC power equals its total internal cooling rate. These two changes to PANDURATA (compared to its description in Schnittman & Krolik 2013) allow faster determination of the coronal temperature map, now necessary since PANDURATA must re-determine the temperature map each pass. We have verified that modified PANDURATA produces the same output spectrum as unmodified PANDURATA. The final modification, however, is more substantive: photon packets that strike the disk surface are subject to absorption and Compton recoil according to albedo and redistribution tables computed with PTRANSX output, using a procedure described in Section 2.4. An example cross section of an electron temperature map so computed is shown in Figure 3.

The end result is a complete description of the electron temperature and radiation field everywhere in the corona—including the flux irradiating the disk surface—that is consistent with the density and cooling structure of the GRMHD simulation as well as the temperature and ionization structure of the disk body.

2.3. Disk Reprocessing—PTRANSX

2.3.1. Defining the Problem

At each (r, ϕ), the region between $\Theta_{\rm top}$ and $\Theta_{\rm bot}$ constitutes a column of the disk body. Its vertical density and cooling rate profiles are known from the HARM3D simulation data, and the fluxes incident upon its upper and lower surfaces are computed by PANDURATA. In addition, some choice of elemental abundances is required. The chief assumption is that such a column can be treated as a finite, plane-parallel slab independent from its neighbors. The problem is to find, for each slab, a description of the radiation field and ionization balance at all vertical points that is energy-conserving, in photoionization equilibrium, consistent with the boundary conditions and structure provided, and which includes as many of the relevant physical processes as possible. Photoionization equilibrium is a reasonable assumption: for the densities and temperatures typical of the accretion disks we consider, the recombination timescale is short, $\sim 10^{-7}$ s (for highly ionized Fe), compared to the disk's dynamical timescale, $\sim 10^{-3}$ s. We accomplish this with the code PTRANSX, a version of which was introduced in Kinch et al. (2016); since then, PTRANSX has undergone major improvements and restructuring, particularly in its treatment of Compton scattering but in other areas as well. Its operation is described below.

2.3.2. The Transfer Solution

We solve the radiative transfer equation in plane-parallel geometry, including all relevant atomic processes and Compton scattering:

$$\mu \frac{dI_{\mu\varepsilon}}{dz} = -\alpha_\varepsilon I_{\mu\varepsilon} + j_\varepsilon - n_e \sigma_\varepsilon I_{\mu\varepsilon} + n_e \int_{-1}^{1} \int_{0}^{\infty} d\mu' d\varepsilon' \sigma_{\varepsilon'} R(\mu', \varepsilon'; \mu, \varepsilon) I_{\mu'\varepsilon'}. \quad (5)$$

We employ the Feautrier method (Mihalas 1978), which requires only that the redistribution function R has forward–backward symmetry, i.e., that $R(\mu', \varepsilon'; \mu, \varepsilon)$ is unchanged under

$$\Delta\theta = \cos^{-1}[\mu'\mu + \sqrt{(1-\mu'^2)(1-\mu^2)}] \rightarrow \Delta\theta \pm \pi. \quad (6)$$

R is a measure of the probability that photons with angle–energy (μ', ε') will scatter to angle–energy (μ, ε). With a specification of boundary conditions—$I_{\mu\varepsilon}$ *inward* at the upper and lower surfaces (the incident intensity from PANDURATA)—we solve a discretized version of the above transfer equation *directly* via a forward–backward recursive sweep (Mihalas 1985).

Our treatment of Compton scattering is expressed by our choice of R. Though we have gone to great lengths to describe Compton scattering as accurately as possible, we are required by the Feautrier method to make the following approximation:

$$R(\mu', \varepsilon'; \mu, \varepsilon) = \frac{1 + [\mu'\mu + \sqrt{(1-\mu'^2)(1-\mu^2)}]^2}{2} \mathcal{R}(\varepsilon', \varepsilon). \quad (7)$$

We replace the angular dependence of a more accurate redistribution function with the dipole phase function of Thomson scattering, which has the required forward–backward symmetry. The Klein–Nishina cross section *does not* have this symmetry—forward scattering is preferred to backward scattering, and significantly so at energies approaching and beyond $m_e c^2$. For the energies we are most concerned with ($\lesssim$50 keV), however, that preference is modest. $\mathcal{R}$ is the angle-average of the full Compton redistribution function (itself a function of the local electron temperature) computed directly with an independent Monte Carlo calculation using relativistic dynamics, the Klein–Nishina cross section, and the Maxwell–Jüttner velocity distribution. The same $\mathcal{R}$ is used in the ensemble scattering calculation of PANDURATA described above.

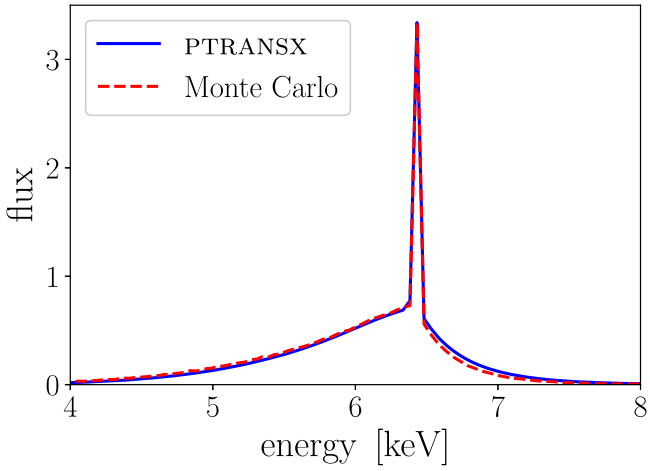

Figure 4. Comparison of example seed photon spectra (Fe Kα only) computed using both PTRANSX and an independent Monte Carlo code. The profiles are normalized to unity.

In order to demonstrate the correctness of our transfer solution in general, and of our treatment of Compton scattering in particular, we compare sample results between a PTRANSX solution and those from a completely independent Monte Carlo transfer code. A straightforward Monte Carlo implementation was supplied with the vertical structure for the density, temperature, emissivity, and absorption opacity from a PTRANSX slab—though, for ease of comparison, we consider here *only* the Fe Kα line emissivity. The independent Monte Carlo code does *not* use the angle-averaged Compton redistribution function described above; rather, it treats Compton scattering directly and with the appropriate angular dependence. Nevertheless, as Figure 4 indicates, the seed photon flux so computed agrees with the PTRANSX result exceedingly well. Note that the slight overprediction of upscattering relative to the Monte Carlo approach is due to the logarithmically spaced energy grid in PTRANSX. Even for a flat probability distribution, it is more likely for a photon to scatter to bin $i+1$ rather than bin $i-1$ if the bin width increases logarithmically. Increasing the number of grid points

alleviates the problem, though the discrepancy as it stands is well below the intrinsic error of any real X-ray detector.

2.3.3. Equilibrium-finding Procedure

We make use of subroutines of the photoionization code XSTAR (Kallman & Bautista 2001) in order to compute the local ionization balance—and consequent emissivity and absorption opacity—of gas at a fixed temperature and density, immersed in a known radiation field, in photoionization equilibrium.

At each (r, ϕ) sampled, the disk body is divided into some number of vertical cells (typically a few dozen, see Section 2.6 below). For the ith cell, the net energy balance y_i is defined as the difference between the net energy flux *out* of the cell and the total cooling within that cell. That is:

$$y_i = \int_0^\infty d\varepsilon (F_{i,\varepsilon}^{\rm top} - F_{i,\varepsilon}^{\rm bot}) - \mathcal{L}_i(\Delta z)_i. \quad (8)$$

The collection of these values for all cells in the given vertical column forms the vector $\boldsymbol{y}$; in total energy balance, $\boldsymbol{y} = 0$. The vector $\boldsymbol{y}$ is directly computable from a complete description of the radiation field, the result of solving the transfer equation. All heating and cooling processes that are represented in either the emissivity, the absorption opacity, or the redistribution function—that is, bremsstrahlung, all atomic processes (photoionization, recombination, and line emission) and (inverse) Compton scattering—have their effects on the energy balance included in the expression for $\boldsymbol{y}$. Knowledge of $I_{\mu\varepsilon}$ in each cell constitutes a full description of the radiation field; similar to $\boldsymbol{y}$, we call such a collection $\boldsymbol{I}$. Similarly, the collection of energy-dependent absorption opacities and emissivities, $(\alpha_\varepsilon, j_\varepsilon)$, in all cells is denoted $\boldsymbol{S}$. Finally, the cell-by-cell list of temperatures is $\boldsymbol{T}$.

The first step in the procedure is to zero out all emission and absorption and perform a transfer solution with only *Thomson* scattering. This yields a guess at the radiation field in each cell. For each cell independently, we supply to the relevant XSTAR subroutines the radiation field, density, temperature, and elemental abundances; XSTAR returns the photoionization equilibrium values for the ionization balance, emissivity (line and continuum), and absorption opacity. We do not use XSTAR's built-in transfer apparatus. As in Kinch et al. (2016), we ignore the resonant absorption of line photons on the basis of a post hoc analysis of their escape probabilities, computed by XSTAR—due to the extremely high local turbulent velocity of the disk gas, these are all very near to unity. For this first iteration, we use XSTAR's ability to find a local energy-conserving temperature while performing its photoionization equilibrium calculation. In addition to the heating and cooling rates that XSTAR considers (see Kallman & Bautista 2001 for details), we also supply the local cooling rate in the HARM3D simulation as an exogenous heating term. In subsequent iterations, we supply a local temperature according to the procedure described next. We now have a first guess at the absorption opacity, emissivity, and temperature in every cell. A second transfer solution performed with this $\boldsymbol{S}_0$ and $\boldsymbol{T}_0$ yields $\boldsymbol{I}_0$ and the corresponding $\boldsymbol{y}_0$.

We imagine our transfer/XSTAR scheme as a vector function: XSTAR requires $\boldsymbol{I}$ and $\boldsymbol{T}$ to determine the photoionization equilibrium $\boldsymbol{S}$, which via our transfer solution produces (a new) $\boldsymbol{I}$ and thus $\boldsymbol{y}$. That is, $F(\boldsymbol{I}, \boldsymbol{T}) = \boldsymbol{y}$. We seek a procedure by which, for a given $\boldsymbol{I}$, we can find the energy-conserving temperature structure $\boldsymbol{T}^*$, such that $F(\boldsymbol{I}, \boldsymbol{T}^*) = 0$.

To do so, we employ the multidimensional Newton–Raphson algorithm. Starting with $\boldsymbol{T}_0$ and a corresponding $\boldsymbol{y}_0$, we perform a finite difference estimation of a Jacobian of the form

$$J_{ij} = \frac{\partial y_i}{\partial T_j}. \quad (9)$$

The new guess at the energy-conserving temperature is

$$\boldsymbol{T} = \boldsymbol{T}_0 - J^{-1}\boldsymbol{y}_0. \quad (10)$$

From this new temperature structure we determine the new $\boldsymbol{S}$ with XSTAR, and with that perform a transfer solution to find the new $\boldsymbol{y}$. We repeat the procedure until all elements of $\boldsymbol{y}$ are sufficiently close to zero, i.e., for all i, the two terms on the right of Equation (8) differ by less than 1%. Thus we find $\boldsymbol{T}^*$. As a practical matter, it occasionally happens that elements in the new temperature vector are not reasonable (for example, negative temperatures); this typically occurs in situations where there are relatively sharp changes in density or cooling rate. In these cases, we require an additional step before the next iteration: these "problem" cells are isolated and their individual y roots found by varying only *their own* T using the secant or bisection methods; these new T values replace their nonsensical counterparts in $\boldsymbol{T}$, and iteration resumes. Following Nayakshin et al. (2000), we estimate that, for the disk temperatures we find in this paper, the maximum Thomson depth over which heat conduction dominates is $\sim 10^{-4}$; because this is much smaller than our cell sizes, it is safe to ignore heat conduction even for sharp cell-to-cell changes in the cooling rate.

It is important to stress that at no step of the procedure described above is the radiation field supplied to XSTAR altered. After the energy-conserving temperature structure $\boldsymbol{T}^*$ is found, and the absorption opacities and emissivities everywhere re-computed with it, one last transfer solution gives us a new, and by construction energy-conserving, radiation field. In fact, the full procedure can be thought of as a function that takes some radiation field $\boldsymbol{I}$ as input and returns a new radiation field $\boldsymbol{I}^*$—this new radiation field is energy-conserving, but the gas is in photoionization equilibrium with the *previous* radiation field $\boldsymbol{I}$. Naturally, then, we just repeat the entire process until $\boldsymbol{I}^* = \boldsymbol{I}$. Thus we accomplish our goal: we have a complete description of a radiation field for which energy is conserved everywhere and with which the gas is at all points in photoionization equilibrium.

In cases where the disk body is many Thomson depths in thickness, it becomes impractical (mainly due to memory constraints) to treat it as a single finite slab extending from one photosphere to the other. Rather, we are forced to set an interior boundary condition some number of Thomson depths inward from the photosphere (we choose 10, see Section 2.6 for details); the natural choice is to assume a blackbody flux *into* the computation volume from the otherwise excluded disk interior. This is similar to the interior boundary condition used by García & Kallman (2010) (i.e., the radiative diffusion approximation, Rybicki & Lightman 1986), but we do not set the disk temperature in advance according to Shakura & Sunyaev (1973). Like the temperature within the computation volume, this boundary temperature is not assumed a priori.

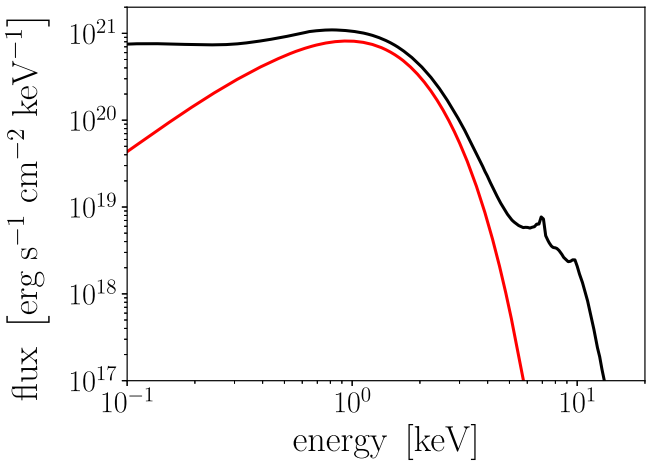

Figure 5. Comparison between the initial assumed seed photon spectrum (red curve) and the converged PTRANSX output spectrum (black curve), at $r = 10\,M$, $\phi = 0$, for $\dot{m} = 0.03$.

Rather, it is found in exactly the same way as part of the same formalism. At the inner boundary, we define an additional element of $\boldsymbol{T}$, $T^{\rm bound}$, and an additional element of $\boldsymbol{y}$,

$$y^{\rm bound} = \pm\int_0^\infty d\varepsilon F_\varepsilon^{\rm bound} - \frac{1}{2}(\text{total } \mathcal{L} \text{ in disk interior}), \quad (11)$$

with the upper sign corresponding to the upper disk layers and the lower sign to the lower disk. Defined so, $y^{\rm bound} = 0$ (at both upper and lower interior boundaries) indicates that the net flux into the computation volume is equal to the total cooling rate *excluded* by the computation volume—like all elements of $\boldsymbol{y} = 0$, it is a statement of energy conservation. With some reasonable starting guess for $T^{\rm bound}$ (e.g., $2\sigma T^4 =$ total $\mathcal{L}$ in disk interior), our multidimensional Newton–Raphson method will find the energy-conserving interior boundary temperature as part of its overall solution.

2.4. The Reprocessed Spectrum

Next we compute from the output of PTRANSX a *new* seed photon spectrum at all (r, ϕ) points on the disk surface. This is done simply by performing one last radiative transfer solution but with the incident intensity set to zero—the seed photon spectrum includes only those photons *emitted* by the gas in the disk, not those that are reflected by it (see Section 2.5 below). In Figure 5 we compare the initially assumed hardened blackbody seed photon spectrum at one point on the disk surface to the PTRANSX output seed photon spectrum. The PTRANSX spectrum is broader, higher at all energies, and has a prominent H-like Fe Kα emission feature at 7 keV and a small K-edge absorption dip near 9 keV. Though the PTRANSX spectrum is slightly harder, there is still virtually zero emission above 10 keV. This pattern generally holds (though with a variable dominant Fe ionization state) for all accretion rates we consider, and at all radii except for where the disk's total Thomson thickness is $\lesssim$1—there the PTRANSX seed photon spectrum is just optically thin free–free emission. For the example point shown, the integrated power of the PTRANSX seed photons is ∼50% greater than that of the hardened blackbody—this is because the PTRANSX seed photons must carry additional energy from Compton and photoionization heating of the disk due to its irradiation by the corona. PANDURATA then ray-traces these new seed photon packets from the disk surface with a limb-darkened angular distribution consistent with PTRANSX's transfer solution. In the initial PANDURATA solution, however, photon packets that strike the disk surface are simply reflected; for this second pass, the absorption and energy redistribution of photons impinging the disk surface is informed by the PTRANSX output.

When a photon packet with energy-dependent intensity I_ε intersects the disk surface, it is transformed into a reflected photon packet with intensity $I_\varepsilon^{\rm refl}$ according to

$$I_\varepsilon^{\rm refl}/\varepsilon = \int_0^\infty d\varepsilon' f(\varepsilon') G(\varepsilon, \varepsilon') I_{\varepsilon'}/\varepsilon'. \quad (12)$$

The albedo f is the fraction of incident photons at a specific energy that are reflected, i.e., not absorbed by the disk, regardless of their final, outgoing energy. G is a normalized description of the redistribution of photons from energy ε' at incidence to energy ε upon reflection. Both f and G are functions of position on the disk surface and of the photon packet's angle of incidence with respect to the local disk normal. These functions are tabulated using a separate, auxiliary Monte Carlo transport code. This additional code injects large numbers of photons at each energy and incident angle from the PTRANSX grids, for each point on the disk surface, using the PTRANSX output opacity and temperature structure; it records for each energy and angle the fraction that are reflected (the albedo f) and the energy distribution of the reflected photons (G in Equation (12)). The Compton scattering calculation therein is performed according to standard relativistic dynamics.

From this same Monte Carlo code, we have found that the distribution in angle of the outgoing photons is a very nearly linear function of μ with respect to the local disk normal (i.e., limb-darkening, but not exactly the expression for a pure scattering atmosphere of Chandrasekhar 1960); we therefore select the initial trajectory of the reflected photon packets according to this distribution. Small portions of the disk are only marginally optically thick, and in these regions there can be a significant *transmitted* fraction. Ideally, the transmitted fraction would spawn a new photon packet in addition to the reflected packet. This, however, is not computationally feasible at this time, so instead we include the transmitted fraction in the reflected photon packet. To the extent that the upper and lower halves of the computation volume are similar, this will ultimately produce the same result. For the cases considered in this paper, transmission through the disk is negligible for $\gtrsim$99% of the disk area. Figure 6 shows the energy-dependent albedo at several radii for $\dot{m} = 0.03$; Figure 7 shows the albedo at $r = 10\,M$ for the four sample accretion rates. The most dramatic feature in both is, not surprisingly, the highly ionized Fe K-edge at 8–10 keV: its depth increases at larger radii (as cooler, less ionized gas has a higher fraction of unstripped Fe atoms available for absorption) and with decreasing $\dot{m}$ (for the same reason; see Equations (1) and (2) and Figure 7).

2.5. PTRANSX and PANDURATA Communication

The revised seed photon packets are the reprocessed emission from the disk, including atomic emission features such as the Fe K lines. As PANDURATA transports them

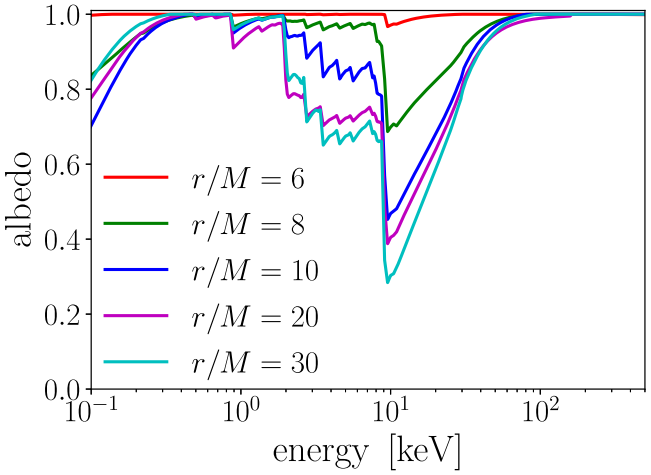

Figure 6. Energy-dependent albedo (including transmitted fraction) at the disk surface at several radii (all $\phi = 0$) for $\dot{m} = 0.03$.

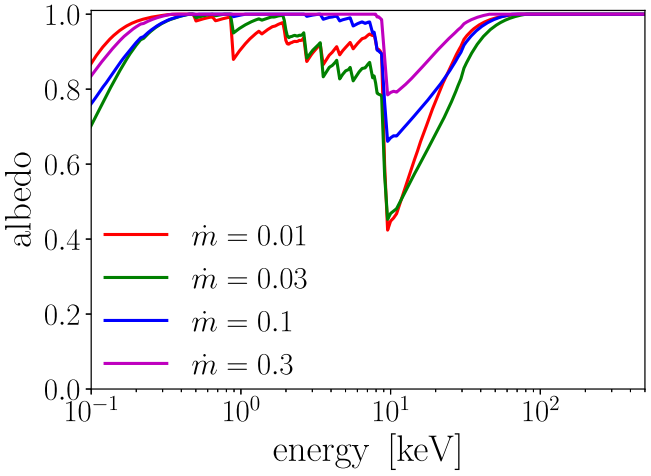

Figure 7. Energy-dependent albedo (including transmitted fraction) at the disk surface at $r = 10\,M$, $\phi = 0$, for several values of $\dot{m}$.

through the corona, they experience IC scattering in addition to all special and general relativistic effects. When they scatter off the disk surface, absorption features such as the Fe K-edge are imprinted. A different spectrum of seed photons than that originally assumed affects the efficiency of the IC cooling process in the corona—so we run PANDURATA again, with disk albedo and Compton recoil tables in hand from the last PTRANSX run, and thus determine a new coronal temperature map and radiation field. This PANDURATA run yields a new irradiating flux, and so we run PTRANSX again to obtain new seed photon spectra and albedo tables. The cycle repeats until the X-ray spectrum as seen by the distant observer changes by less than 1% from one pass to the next.

Our method separates *emission* from the disk and *absorption/reflection* by the disk into sequential steps: Fe Kα line photons, for example, are emitted as part of the seed photon flux for a point on the disk surface consistent with the incident flux at that point found from the previous PANDURATA run; likewise, photon packets that strike the disk as they are ray-traced through the corona are subject to absorption according to the disk's opacity found in the preceding PTRANSX run. The same is true for the overall energy balance—the power in the seed photon flux accounts for both the disk's internal dissipation *and* Compton and photoionization heating of the disk's gas consistent with the incident flux from, again, the previous PANDURATA run. By cycling between the two codes until the global radiation field (including the disk incident flux) no longer changes, we ensure *global* energy balance and the self-consistency of our non-simultaneous treatment of absorption and emission.

The final run of PANDURATA yields the desired product: a spectrum as seen by a distant observer—a *prediction*, arrived at through consideration of the relevant physical principles, for what we expect to observe from an accreting stellar-mass black hole, specifying only the *physical* parameters of mass, spin, accretion rate, and elemental abundances.

2.6. Numerical Specifics

In describing our technique, we have been intentionally vague concerning any numerical values. While our general approach is applicable to a large volume of the stellar-mass (and even AGN) black hole parameter space, the specific resolutions, samplings, etc., that we use in a real calculation must be tailored to the kind of problem we want to solve—we must balance the desired accuracy and completeness of our prediction with realistic computational constraints. In practice, this requires numerical experimentation: resolutions and samplings start off coarse and are repeatedly refined until the final results (presented in the next section) no longer appear to change.

Here we consider four cases, at $\dot{M}/\dot{M}_{\rm Edd} = \dot{m} = 0.01$, 0.03, 0.1, and 0.3; per Equations (1) and (2) in Section 2.1, the choice of $\dot{m}$ translates HARM3D data snapshots into physical (cgs) values for the density and cooling rate. For each, the central black hole mass is $10\,M_\odot$, the spin is zero, and the abundances are solar (values of Grevesse et al. 1996). When running PANDURATA, the coronal volume is divided into ~24,000 sectors (the exact number varies with the location of the disk photosphere, and therefore decreases with increasing accretion rate) of $\Delta\theta = \pi/36$ and $\Delta\phi = \pi/32$ radians each, with a logarithmically increasing radial extent such that $\Delta r/r = 0.062$, starting at the event horizon, $r = 2\,M$. For the snapshots we used, smaller sectors than these do not result in an appreciable change to the shape or strength of the X-ray flux incident upon the disk surface or seen by the distant observer—doubling the number of sectors results in less than a 1% change to the observable spectrum in the range 1–30 keV. The majority of these sectors lie wholly in the corona, but those that overlap with the disk have only their coronal part included in PANDURATA's calculation. When running PTRANSX, we sample ~300–500 (r, ϕ) points per case (with more for case with the higher accretion rate because the inner edge of the disk photosphere extends further inward). These are chosen uniformly in azimuth (at 8ϕ angles) and logarithmically in r such that $\Delta r/r = 0.062$. For each (r, ϕ) slab, the vertical cells are spaced semi-logarithmically in Thomson depth τ, such that the increase in Thomson depth *into* the disk over one cell is $\Delta\tau/\tau = 0.25$, but with $\Delta\tau$ limited to a maximum of 0.4; the grid is laid out so that the cells follow this semi-logarithmic spacing *into* the disk starting from both upper and lower photospheres, meeting at the midplane. Slabs with a total Thomson depth of 20 or greater are cleaved into upper and lower computation volumes as described above, with the interior boundary always placed at a Thomson depth of 10 as measured from the relevant photosphere. The number of vertical cells used varies from six at the extreme inner edge of

the disk to 50 at its thickest extent before the approximation just described is employed; the "cleaved" slabs are separated into upper and lower volumes of 27 cells each. Neither finer spacing in optical depth nor a deeper interior boundary results in an appreciable change to the output seed photon spectrum.

The angle with respect to $\hat{z}$, the cosine of which is μ in our transfer equation, is discretized such that 16 bins uniformly spaced in μ cover the range -1 to 1. This is more than sufficient to capture the angular dependence of the radiation field, which is nearly isotropic for most of the disk body. Our energy grid is more complex. For the purposes of determining the temperature structure and photoionization balance via a multidimensional Newton–Raphson scheme, we span the range from 1 eV to 10^8 eV with a coarse 161-point grid whose energy resolution is $\Delta\varepsilon/\varepsilon = 0.122$. The computational cost of the transfer solution scales poorly with the number of energy bins (cubically) and our multidimensional Newton–Raphson scheme requires the transfer solution to be performed *many* times—typically 20–80 iterations per slab, depending on its thickness and ionization parameter. Yet because the bulk of the power in the radiation field is in the continuum, and the broadband continuum can be well represented on such a coarse grid, increasing the energy resolution further results in little change in the equilibrium temperature structure. Thus the approach we take is to use a coarse grid to find the equilibrium temperature structure, then re-bin to a finer 801-point grid ($\Delta\varepsilon/\varepsilon = 0.0233$) on which we perform *one* last transfer solution at a resolution high enough so that line features are clearly distinguishable; we use this same procedure (and identical energy grids) with PANDURATA as well.

We take 1% as sufficient for all convergence tests. That is, "energy conservation" (for both PTRANSX and PANDURATA) means (energy in) = (energy out) is satisfied in all cells/sectors (and globally) to within at most 1%. The majority are better converged by the time this is achieved—typically, ∼90% of cells/sectors conserve energy within 0.1%. For determining whether $\boldsymbol{I}^* = \boldsymbol{I}$, we compute the first several energy moments of the mean photon intensity in the range 1–30 keV (the region of the outgoing spectrum we are most concerned with) at each cell; when the greatest fractional difference between any of these values and its counterpart in the previous iteration has dropped below 1%, we consider the radiation field to have converged. Finally, the cycling between PTRANSX and PANDURATA ceases when the greatest difference in the spectrum as seen by a distant observer (at any inclination or energy in the range 1–30 keV) from one iteration to the next differs by, again, at most 1%—for the cases we discuss in this paper, this takes between 5 and 10 passes.

3. Results

3.1. Continuum

The key results of our calculation are predicted X-ray spectra as seen by a distant observer. Figure 8 shows the broadband spectral luminosity for the four accretion rates we consider. Because the dynamical timescale for stellar-mass black holes is many times smaller than the integration time for any reasonable observation of them, we present all distant observer spectra as azimuthally averaged. Particularly for the cases with $\dot{m} \geqslant 0.03$, these broadband spectra reproduce the forms inferred by phenomenological fitting of real black hole X-ray binary data in the steep power-law state (Remillard & McClintock 2006):

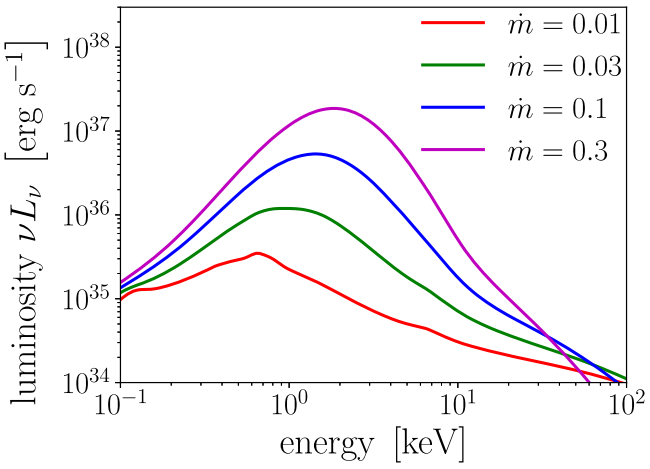

Figure 8. Spectral luminosity at four accretion rates, each with a central black hole mass of 10 $M_\odot$ and $a = 0$. Note that the spectra soften as the accretion rate increases. Although not easily visible in this representation, the equivalent width of the Fe Kα feature diminishes with increasing $\dot{m}$.

there is a quasi-thermal bump at 1–3 keV that is extended to high energy as a steep power law that hardens slightly above ∼10 keV. The photon index Γ computed over the 5–30 keV band ranges from 2.7 for $\dot{m} = 0.01$ to 4.5 for $\dot{m} = 0.3$, comparable to the values observed for black hole binaries in the thermal and steep power-law states, $\Gamma = 2.1$–4.8 (McClintock & Remillard 2006). The thermal bump is not too surprising, and its existence and temperature follow from having a dense, optically thick disk body with a sub-Eddington accretion rate around a $\sim 10\,M_\odot$ black hole. On the other hand, the prediction of a steep power-law component, due to the IC upscattering of thermal seed photons in a hot corona, represents a triumph for the theory of MHD accretion disks. No coronal emission at all is predicted by the models of Shakura & Sunyaev (1973) and Novikov & Thorne (1973). We find here that a purely physical calculation, starting from a GRMHD simulation of black hole accretion, gives rise *naturally* to the approximate spectral shape observed for black hole binaries in the steep power-law state, with no phenomenological descriptions of the accretion geometry (of disk or corona) or parameter-tweaking required. This result was first shown by Schnittman et al. (2013), also using PANDURATA analysis of HARM3D simulations. With our more careful treatment of the seed photon spectrum and the inclusion of disk absorption enabled by coupling to PTRANSX, we predict slightly softer spectra than those reported in Schnittman et al. (2013). Curiously, using simulations nominally similar to ours (GRMHD thin-disk simulations with prescribed cooling functions) and a post-processing procedure calculating the Comptonization of initially thermal photons, Narayan et al. (2016) were unable to find any high-energy extension of the thermal component.

The spectra in Figure 8 have power-law tails that extend to very high energies. In Figure 9, we show (for $\dot{m} = 0.03$) the distribution of IC power generation in the corona as a function of the electron temperature. The photons that make up the observable spectrum were upscattered by electrons with a wide range of temperature, 1–1000 keV, but the distribution is distinctly bimodal with peaks at 10–30 keV and 400–800 keV. While the majority of the cooling is due to electrons with temperatures less than 100 keV, approximately 20% of the coronal power is radiated from electrons with temperatures in excess of 400 keV. The other three cases of accretion rate also

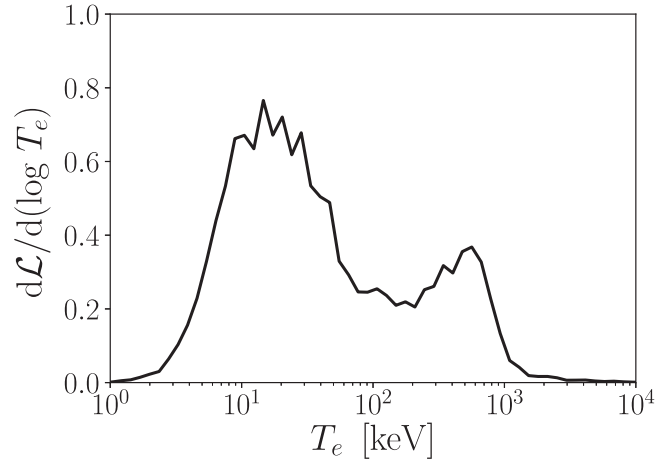

Figure 9. Distribution of IC power as a function of electron temperature in the corona for $\dot{m} = 0.03$. The data are normalized to the total cooling in the corona. Note both the broad range in temperature and the distinctly bimodal shape of the distribution.

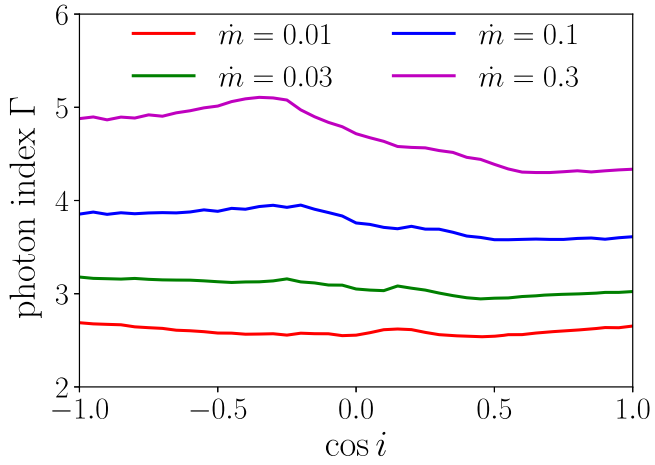

Figure 10. Photon index of the predicted spectrum, measured in the range 5–30 keV, as a function of observer angle at four accretion rates.

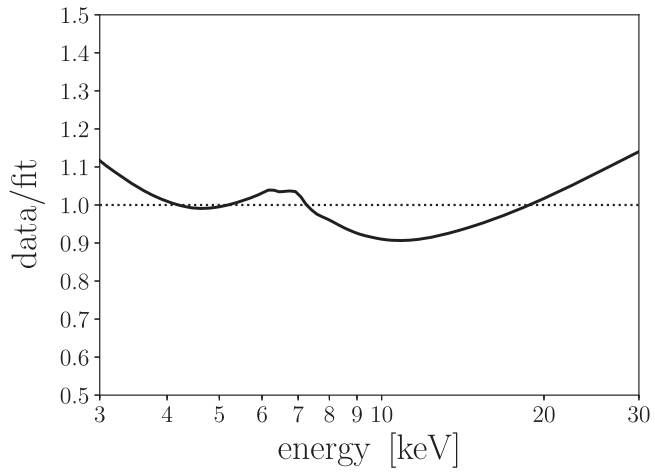

Figure 11. Predicted spectrum divided by a power-law fit to the range shown for $\dot{m} = 0.01$ and $i = 25°$.

have similarly bimodal distributions. It is evident from this figure that a single-temperature Comptonization model of the corona cannot adequately describe our results. Similarly, a single Compton y-parameter does not adequately describe coronal scattering. A thermal 1 keV seed photon that escapes to the distant observer will typically undergo 3–7 scatters. If it scatters through electrons at $T_e = 20$ keV (roughly the location of the first peak in Figure 9), then $y \approx 0.7$; however, if it scatters through electrons at $T_e = 500$ keV (roughly the location of the second, smaller peak), $y \approx 70$. From Figure 3 (which, while shown for $\dot{m} = 0.01$, is qualitatively similar to the $\dot{m} = 0.03$ case), we see that there is a clear dependence of the coronal temperature on polar angle—hotter regions are those more inclined relative to the midplane. We see from this figure as well that a single-temperature description is unsatisfactory: a thermal disk photon will likely pass through several layers of gas with very different temperatures, scattering in any or all of them.

We do not see clear evidence for Compton bumps in the spectra of Figure 8. These bumps can be seen most clearly when the disk is absorptive across the soft X-ray band and up through the Fe K-edge, and when the underlying continuum is relatively hard so that there are plentiful photons above ∼10 keV to scatter. Here neither is the case. The result is that the numerous photons with energy below the onset of absorption at the Fe K-edge can be upscattered and smooth over the feature. We expect the reflection hump to be visible when we scale these same simulations to AGN masses and temperatures. At that scale, we expect elements other than Fe to produce important spectral features as well.

Figure 10 shows the photon index measured in the range 5–30 keV as a function of observer angle for each accretion rate. The power-law slope varies irregularly, but only slightly, with inclination; its range increases with increasing $\dot{m}$. The extent of the top–bottom asymmetry seen in Figure 10 (and also in Figure 13 below, showing Fe Kα equivalent width in the same fashion) provides a rough indication of the "cosmic variance" expected for these simulations.

It is important to emphasize again (see Section 2.1) that each case of accretion rate we consider is the same underlying simulation snapshot with the density and cooling rate scaled. The simulation used here is most physically realistic for $\dot{m} = 0.1$–0.3; it is not surprising that our results for this range of accretion rates most closely resemble observations. HARM3D is unable to relate disk vertical structure to accretion rate because it has an ad hoc procedure for radiative cooling and does not include radiation forces at all. However, techniques for coupling radiation transport to MHD are rapidly improving (Jiang et al. 2014; Sądowski et al. 2014). In the future, it will be possible to reapply our method to data from codes of that variety in order to work with a more realistic connection between accretion rate and disk structure.

3.2. Fe Kα

To illustrate our predictions of the Fe Kα line profile, we adopt a procedure mimicking a common approach to presenting observational data: we divide the data by a simple prescription for the continuum—in this case, a power-law fit to the region 3–30 keV. Figure 11 shows this procedure applied to the $\dot{m} = 0.01$ case at an observer inclination of 25°, for which the fitted power law has photon index $\Gamma = 2.7$. We reproduce the features expected: a relativistically broadened Kα emission line near 6.4 keV and a K-edge absorption trough centered roughly at 10 keV. However, the contrast of both features relative to the power-law fit is quite small, only 5%–10%. In addition, above 15–20 keV there is a slight

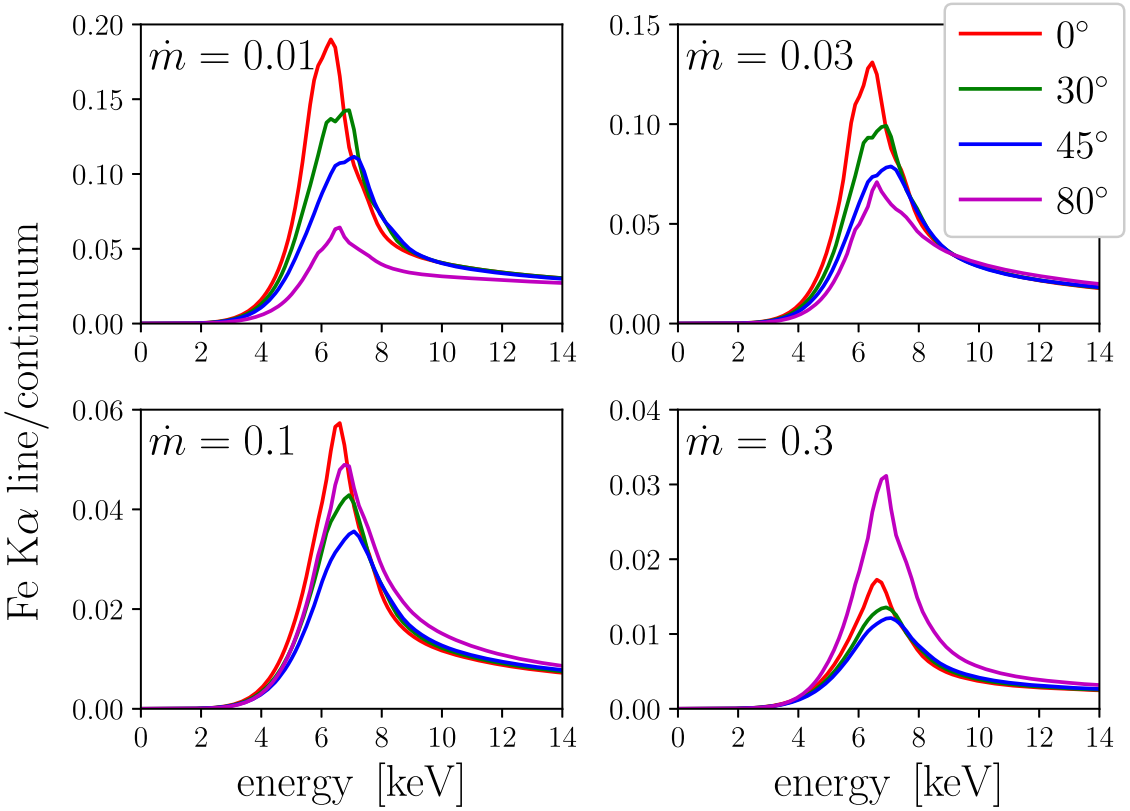

Figure 12. Photons whose origin is an Fe Kα transition, as a fraction of all continuum photons, once they have reached the distant observer; for four accretion rates at several inclinations each. The "shelf" at high energy is due to IC upscattering in the corona taken in ratio to a steeply declining continuum. Note the difference in scale for each subplot.

hardening of the continuum relative to $\Gamma = 2.7$, the single value that best fits the 3–30 keV continuum slope.

In Figure 12, we show those photons that *originate* from the Kα transition as a fraction of continuum photons at the same energy, as seen by a distant observer, at several inclinations for each accretion rate. It is important to note that while this representation emulates model-fitting procedures, these are *not* themselves model fits divided by the total flux. We produce these plots by keeping track of Kα photons as they diffuse from their point of creation to the disk surface and are then ray-traced to infinity, with no continuum-fitting needed. We calculate the equivalent width (EW) directly as well, as shown in Figure 13 as a function of observer angle.

The Fe Kα line profiles strongly resemble actual spectral data in the sense that they are fairly broad and their EWs are in the range often measured ($\sim$100 eV, see below). On the other hand, they also differ in some respects. In particular, the "shelf" at high energies in Figure 12 is due to the upscattering of Fe Kα photons as they traverse the corona. With the notable exception of composite models such as those presented in Steiner et al. (2017), this "shelf" line flux is typically not accounted for by continuum-fitting models; rather, these upscattered Fe Kα photons would appear as continuum photons, potentially introducing a systematic bias in a continuum-subtraction procedure as the level of the perceived continuum at energies above the Kα feature is artificially elevated by a few per cent. This coronal upscattering effect, in addition to our use of a non-spinning black hole, also explains the blueward asymmetry of the line profiles in Figure 12. The simple data/fit representation of Figure 11 is the closest approximation to actual observed line profiles: compared to,

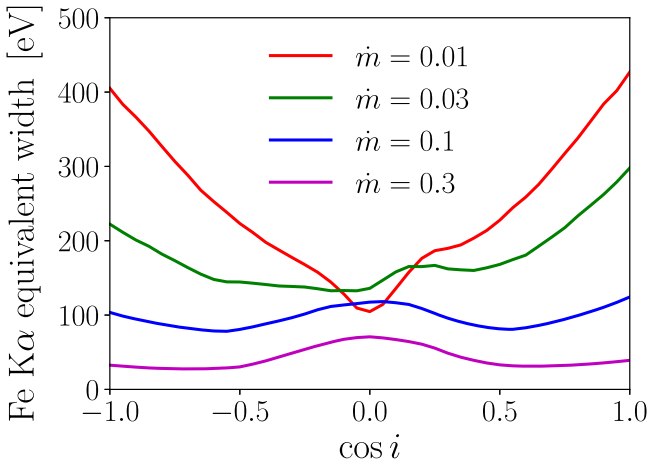

Figure 13. Fe Kα equivalent width, as a function of observer angle, for each of the four accretion rates we consider. The EW is computed in the range 2–8 keV.

e.g., Cyg X-1 line profiles (Reynolds & Nowak 2003; Walton et al. 2016), ours are nearly as broad, though with a slightly less extended red wing.

The Fe Kα feature becomes relatively weaker with respect to the continuum as the accretion rate increases—the peak contrast drops by a factor of six from $\dot{m} = 0.01$ to $\dot{m} = 0.3$ in Figure 12. This is consistent with the increasing ionization parameter (Tarter et al. 1969): $\log \xi$ varies with position (decreasing with radius), but is in the range 2.5–3.0 for $\dot{m} = 0.01$, increasing to 3.0–3.5, 3.5–4.0, and 4.0–4.5, for $\dot{m} = 0.03$, 0.1, and 0.3, respectively.

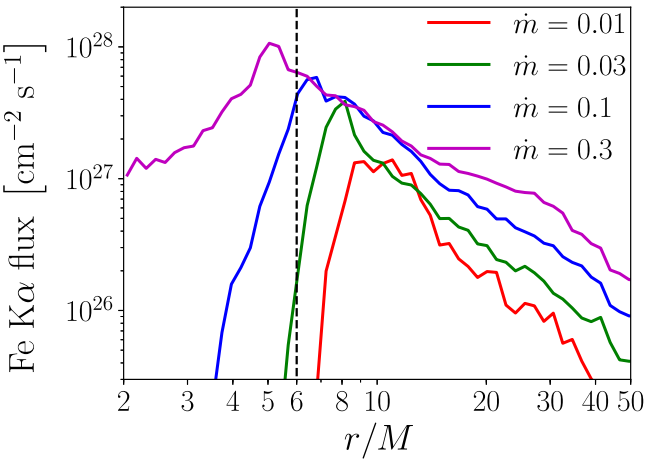

Figure 14. Fe Kα surface brightness, averaged over azimuth and both disk surfaces, for four accretion rates. Note the location of the peaks with respect to the ISCO at $r = 6\,M$.

For the lower accretion rates, the line contrast drops monotonically with increasing inclination. However, as $\dot{m}$ increases, the near edge-on ($i = 80°$) view becomes stronger relative to the other viewing angles, and is twice the strength of the other three sample inclinations for $\dot{m} = 0.3$. The overall line flux for the near edge-on viewing angles is less than for the face-on inclinations for all accretion rates, but the strength of the line *relative* to the continuum is greater when viewed nearly edge-on because the disk itself obscures emission from outer radii (where the line flux is weak compared to the continuum) while, due to lensing, light from the inner disk region (where the line flux is strong compared to the continuum) still reaches the distant observer; this is only the case when there *is* any Kα emission at the innermost radii, i.e., for higher accretion rates (see Figure 14).

The line strengths we find are, overall, comparable to those typically observed, with EWs in the range 40–400 eV. Compare, for example, to the analysis of Cyg X-1 soft state spectra by Walton et al. (2016); they report an Fe Kα EW = 300–330 eV (see also the discussion in Reynolds & Nowak 2003). It is a well-known phenomenon that modeling–fitting of black hole X-ray spectra often results in inferred Fe abundances that are several to many times the solar value (García et al. 2018; Tomsick et al. 2018). Because we have not yet fit real data with our simulated spectra, nor do we analyze our theoretical spectra with the same techniques used by observers, it is not possible to make *direct* comparisons between our line strengths and those reported in, e.g., Walton et al. (2016). Nevertheless, that we use *only* solar Fe abundances yet still find strong Kα lines is encouraging. Our approach is fundamentally different from those often employed when fitting actual spectra, so it is difficult to pinpoint a single reason why we do not need supersolar Fe abundances to achieve strong lines. A major contributor, however, is likely our naturally extended corona (see Figure 3), which allows for Fe Kα production over a larger fraction of the disk surface. Figure 14 shows the radial dependence of the Fe Kα surface brightness for each accretion rate, averaged over azimuth and the two surfaces of the disk. Like most phenomenological models, the variation with radius roughly follows a power law—our power laws, however, are approximately $\propto r^{-2}$ (steepening slightly with decreasing accretion rate), a shallower profile than is typically assumed from lamppost geometries ($\propto r^{-3}$ or steeper, Wilkins & Fabian 2012; Dauser et al. 2013). Azimuthal variations superimposed on these radial gradients can be sizable: the relative standard deviation of the ϕ-variation of the Kα emission decreases with increasing $\dot{m}$, from $\sim$50% for $\dot{m} = 0.01$ to $\lesssim$10% for $\dot{m} = 0.3$.

The radius of peak Kα surface brightness moves inward with increasing accretion rate, from $\sim 10\,M$ for $\dot{m} = 0.01$ to $\sim 5\,M$ for $\dot{m} = 0.3$. The Kα surface brightness at radii interior to the peak increases with accretion rate as well; for $\dot{m} = 0.3$, there is significant Kα production even just outside the event horizon. While the strength of the Fe line relative to the continuum diminishes with increasing accretion rate (Figure 13), the *number* of Kα photons increases as the accretion rate, and therefore the total luminosity, increases. Though we do find that the peak in Kα surface brightness occurs somewhat *near* the ISCO, we do not find a sharp cutoff exactly *at* the ISCO. This should not be too surprising. If gas flows into the black hole, there must be some gas between the event horizon and the ISCO that might produce Fe Kα emission; conversely, if the accretion rate is low and the illuminating flux particularly powerful, there could be no available unstripped Fe to emit photons except for well outside the ISCO. In general, the interior Kα emission cutoff cannot just be a function of the spin alone—it depends also on the disk's surface density and ionization state, which themselves depend on the accretion rate, as we demonstrate in Figure 14 (see also the discussions in Reynolds & Begelman 1997; Krolik & Hawley 2002; Beckwith et al. 2008).

We have considered only the $a = 0$ case here, but the variation of the peak surface brightness with accretion rate for *non*-spinning black holes has important implications for spin-measuring techniques that rely on identifying the interior cutoff of Kα production as the ISCO. In future work, we will apply our method to simulations with $a > 0$ as well. We have already done this for the continuum flux method for spin measurement (Schnittman et al. 2016); it will be very interesting to see how the Fe line technique compares to the continuum method.

4. Conclusion

The most important result is simply that our machinery can manufacture *forward* predictions of the entire X-ray spectrum radiated by an accreting stellar-mass black hole—line *and* continuum features together—in a self-consistent, energy-conserving fashion, directly from the output of high-resolution 3D GRMHD simulations. It is worth repeating that we have required *no* assumptions about the accretion flow geometry at any point—no lamppost coronae, or disk inner edges fixed exactly at the ISCO. And yet, by specifying *only* the physically meaningful parameters of mass, spin, accretion rate, and elemental abundances, and then applying to the simulation data the relevant physical principles and carrying out detailed radiative transfer and photoionization calculations, we are able to produce spectra similar in shape and principal features to those actually observed from stellar-mass black holes. It also bears emphasizing that we employed standard techniques without preference for a desired outcome—a Monte Carlo radiation transport code that treats only Compton scattering is the natural choice in the hot, optically thin corona, while a plane-parallel Feautrier method that treats Compton scattering, free–free, and all atomic processes, is the natural choice in the cooler, optically thick disk. That the output of these methods when applied to simulation data resembles so well the familiar

X-ray spectra of their real counterparts encourages us to attempt to understand these objects from the standpoint of direct application of well-understood physical principles, as opposed to phenomenological modeling. In this vein, we ultimately plan to use this method for the production of grids of spectra—allowing observers to attempt to fit real spectral data with an XSPEC package that requires only a relatively small set of physical parameters. Moreover, as simulation codes are improved (in particular with regard to the equation of state), our methods can be readily employed upon the data they produce.

We thank Jon Miller, Andy Fabian, Javier García, and Chris Done for helpful conversations. This work was partially supported by NASA ATP grant NNX14AB43G and NSF grants AST-1516299 and AST-1715032. J.D.S. was partially supported by NASA grant ATP13-0077. We are also grateful for generous allocations of computer time through the XSEDE program of NSF (TG-MCA95C003) and at the Maryland Advanced Research Computing Center (MARCC).

Software: XSTAR (Kallman & Bautista 2001), HARM3D (Noble et al. 2009), PANDURATA (Schnittman & Krolik 2013), PTRANSX (Kinch et al. 2016).

ORCID iDs

Brooks E. Kinch https://orcid.org/0000-0002-8676-425X
Jeremy D. Schnittman https://orcid.org/0000-0002-2942-8399
Timothy R. Kallman https://orcid.org/0000-0002-5779-6906
Julian H. Krolik https://orcid.org/0000-0002-2995-7717

References

Beckwith, K., Hawley, J. F., & Krolik, J. H. 2008, MNRAS, 390, 21
Chandrasekhar, S. 1960, Radiative Transfer (New York: Dover)
Dauser, T., Garcia, J., Wilms, J., et al. 2013, MNRAS, 430, 1694
García, J., Dauser, T., Lohfink, A., et al. 2014, ApJ, 782, 76
García, J., Dauser, T., Reynolds, C. S., et al. 2013, ApJ, 768, 146
García, J., & Kallman, T. R. 2010, ApJ, 718, 695
García, J., Kallman, T. R., & Mushotzky, R. F. 2011, ApJ, 731, 131
García, J. A., Kallman, T. R., Bautista, M., et al. 2018, in ASP Conf. Ser. 515, Workshop on Astrophysical Opacities, ed. C. Mendoza, S. Turck-Chiéze, & J. Colgan (San Francisco, CA: ASP), 282
Grevesse, N., Noels, A., & Sauval, A. J. 1996, in ASP Conf. Ser. 99, Cosmic Abundances, ed. S. S. Holt & G. Sonneborn (San Francisco, CA: ASP), 117
Haardt, F., & Maraschi, L. 1991, ApJL, 380, L51
Hawley, J. F., Richers, S. A., Guan, X., & Krolik, J. H. 2013, ApJ, 772, 102
Jiang, Y.-F., Stone, J. M., & Davis, S. W. 2014, ApJS, 213, 7
Kallman, T., & Bautista, M. 2001, ApJS, 133, 221
Kinch, B. E., Schnittman, J. D., Kallman, T. R., & Krolik, J. H. 2016, ApJ, 826, 52
Krolik, J. H., & Hawley, J. F. 2002, ApJ, 573, 754
Liang, E. P. T. 1979, ApJL, 231, L111
McClintock, J. E., & Remillard, R. A. 2006, in Black Hole Binaries, ed. W. H. G. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 157
Mihalas, D. 1978, Stellar Atmospheres (2nd ed.; San Francisco, CA: Freeman)
Mihalas, D. 1985, JCoPh, 57, 1
Miller, M. C., & Miller, J. M. 2015, PhR, 548, 1
Mitsuda, K., Inoue, H., Koyama, K., et al. 1984, PASJ, 36, 741
Narayan, R., Zhu, Y., Psaltis, D., & Sądowski, A. 2016, MNRAS, 457, 608
Nayakshin, S., Kazanas, D., & Kallman, T. R. 2000, ApJ, 537, 833
Noble, S. C., Krolik, J. H., & Hawley, J. F. 2009, ApJ, 692, 411
Noble, S. C., Krolik, J. H., Schnittman, J. D., & Hawley, J. F. 2011, ApJ, 743, 115
Novikov, I. D., & Thorne, K. S. 1973, in Black Holes (Les Astres Occlus), ed. C. Dewitt & B. S. Dewitt (New York: Gordon and Breach), 343
Remillard, R. A., & McClintock, J. E. 2006, ARA&A, 44, 49
Reynolds, C. S. 2013, CQGra, 30, 244004
Reynolds, C. S., & Begelman, M. C. 1997, ApJ, 488, 109
Reynolds, C. S., & Nowak, M. A. 2003, PhR, 377, 389
Ross, R. R., & Fabian, A. C. 2005, MNRAS, 358, 211
Rybicki, G. B., & Lightman, A. P. 1986, Radiative Processes in Astrophysics (Weinheim: Wiley)
Sądowski, A., Narayan, R., McKinney, J. C., & Tchekhovskoy, A. 2014, MNRAS, 439, 503
Schnittman, J. D., & Krolik, J. H. 2013, ApJ, 777, 11
Schnittman, J. D., Krolik, J. H., & Noble, S. C. 2013, ApJ, 769, 156
Schnittman, J. D., Krolik, J. H., & Noble, S. C. 2016, ApJ, 819, 48
Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 24, 337
Steiner, J. F., García, J. A., Eikmann, W., et al. 2017, ApJ, 836, 119
Tanaka, Y., Nandra, K., Fabian, A. C., et al. 1995, Natur, 375, 659
Tarter, C. B., Tucker, W. H., & Salpeter, E. E. 1969, ApJ, 156, 943
Tomsick, J. A., Parker, M. L., García, J. A., et al. 2018, ApJ, 855, 3
Walton, D. J., Tomsick, J. A., Madsen, K. K., et al. 2016, ApJ, 826, 87
Wilkins, D. R., & Fabian, A. C. 2012, MNRAS, 424, 1284