Durham E-Theses

Accretion and Ejection around Astrophysical Black Holes

GARDNER, EMMA,LOUISE

How to cite:

Use policy

Accretion and Ejection around Astrophysical Black Holes

Emma Gardner

A thesis submitted for the degree of

Doctor of Philosophy

Department of Physics

University of Durham

June 2015

Accretion and Ejection around Astrophysical Black Holes

Emma Gardner

Submitted for the degree of Doctor of Philosophy

June 2015

Abstract

Astrophysical black holes should be simple objects with only two parameters: mass and spin. As material accretes onto the black hole this adds two further parameters: accretion rate and, since accretion generally occurs through a preferential plane, the inclination at which we view the system. Inclination becomes particularly important when a fraction of the inflowing material is not accreted but is instead ejected from the system in powerful, highly collimated and sometimes highly relativistic jets. It is these luminous accretion flows and jets that allow accreting black holes to be detected across the entire range of the electromagnetic spectrum from radio up to gamma-ray energies. The emission from the accretion flow and jet should be completely determined by the four fundamental parameters of mass, spin, accretion rate and inclination. Variations in these four parameters should be all that is required to explain the enormous variety of spectra from accreting black holes. In this thesis I present five papers studying emission from black holes of all size-scales and in all accretion regimes, including inclination effects and investigating the possible effect of the most difficult parameter to measure — black hole spin. Black holes do not exist in isolation. Stellar mass black holes are fed by their companion stars and supermassive black holes by gas from their host galaxies. Not only does the galaxy fuel the growth of the supermassive black hole but equally the outflows that result from black hole accretion affect the growth of the galaxy, heating gas and suppressing star formation. This adds extra importance to understanding black hole accretion and the interplay between accretion and ejection. I find a scenario where low spin black holes are limited to feeding back via winds and moderately relativistic jets, while only the highest spin black holes are capable of producing the most powerful, highly relativistic jets, may be consistent with current observations.

Declaration

Portions of this work have appeared in the following papers:

- Gardner E., Done C., 2013, MNRAS, 434, 3454

- Gardner E., Done C., 2014, MNRAS, 438, 779

- Gardner E., Done C., 2014, MNRAS, 442, 2456

- Gardner E., Done C., 2015, MNRAS, 448, 2245

Acknowledgements

Contents

List of Figures

List of Tables

Chapter 1

Introduction: Fundamental Properties of Black Holes

A black hole is formed when material collapses to such high densities that not even light can escape its gravitational potential well. The strength of a gravitational field can be characterised by the velocity a test mass would require to escape to infinity. The closer to the black hole the test mass is launched, the stronger the gravitational attraction it experiences and the larger the velocity required to escape to infinity. The radius around a black hole at which the escape velocity equals the speed of light is called the event horizon. Escaping from within the horizon requires an escape speed greater than the speed of light. Since nothing can travel faster than the speed of light this means the region of space-time inside the horizon is cut off from the rest of the Universe. The radius of the event horizon (R_H) depends on both the mass and the spin of the black hole. For a non-spinning (Schwarzschild) black hole, $R_H = 2GM/c^2 = 2\,R_g$, where R_g is the gravitational radius (used hereafter to write distances in dimensionless units as $r = R/R_g$). However if the black hole is spinning it will drag space-time around with it, reducing the radius of the event horizon. For a spinning black hole, $r_H = 1 + \sqrt{1 - a^2}$ (Fig.1.1, top panel), where $a = J/McR_g$ is the dimensionless black hole spin parameter ($-1 \leq a \leq 1$) and J is the angular momentum of the black hole (Kato et al., 1998). A black hole therefore has only two fundamental properties: its mass and its spin.

By its very nature a black hole cannot be seen directly, but its presence can be

Figure 1.1: Radius of the horizon and innermost stable circular orbit and black hole efficiency as a function of black hole spin.

inferred through its effects on its surroundings. The extreme gravity around a black hole strongly distorts space-time, producing regions with high curvature. Light rays from a distant source passing close to the black hole but not close enough to be captured follow this curved space-time so that their path appears bent around the black hole. In this way black holes become gravitational lenses (see, for example, the review by Bozza, 2010).

Any material entering the black hole's sphere of influence may be accreted onto the black hole. In order to accrete, the material must lose potential energy and angular momentum, and some of this energy is lost via radiation. The nearer to the black hole the accreting material, the deeper the potential well and the higher the energy of the emitted radiation. As a result, the emission from such accretion flows peaks in the X-rays, with luminosities as high as $10^{43}\,\mathrm{erg\,s^{-1}}$ for a typical supermassive black hole, making accreting black holes some of the brightest objects

in the Universe. The black hole itself may be invisible, but the luminous accretion flows around them provide a means of studying them and their environment. Material accretes in a series of concentric circular orbits with decreasing radius, until it reaches the innermost stable circular orbit (R_{isco}), after which general relativity predicts there can be no more stable orbits and material goes into free fall towards the horizon and the black hole. Just like R_H, R_{isco} depends on the mass and spin of the black hole, varying from $1\,R_g$ for $a = 1$ (maximal prograde spin, i.e. accretion flow and black hole spin perfectly aligned) to $9\,R_g$ for a maximal retrograde black hole with $a = -1$ (Fig.1.1, middle panel). The total amount of energy the accretion flow can radiate depends on R_{isco}. Hence the efficiency of the black hole at converting accreting matter to radiation depends on its spin (Fig.1.1, bottom panel). Not only does the accreted material add to the mass of the black hole but it also carries its own angular momentum, which can in turn act to increase or decrease the spin of the black hole.

Black hole accretion rates are measured in terms of the ratio of the accretion flow luminosity with respect to the Eddington luminosity (L_{Edd}), where the Eddington luminosity is the luminosity at which radiation pressure pushing outwards balances the inward gravitational force. L_{Edd} therefore represents a theoretical limit on the accretion rate. If the accretion rate exceeds the Eddington limit then radiation pressure from the accreting material is greater than the gravitational attraction of the black hole and pushes material away, preventing the accretion rate rising any further. The Eddington limit depends on the mass of the black hole as:

$$L_{Edd} = \frac{4\pi G M m_p c}{\sigma_T} \sim 1.3 \times 10^{38} M/M_\odot \qquad (1.0.1)$$

where σ_T is the Thompson cross section for electron scattering and $M_\odot$ is a solar mass. The Eddington limit assumes spherical accretion, so in practice the Eddington limit can be overcome if the radiation from the accreting flow is directed away from the in-falling material, reducing the radiation pressure on it. In this way super-Eddington accretion flows can and probably do occur.

Therefore, in addition to a black hole's two fundamental properties of mass and spin, a black hole may have two further important parameters: the mass accretion

rate onto it and, since accretion often occurs in a preferential plane, the inclination of the accreting system with respect to our line of sight. This fourth parameter of inclination becomes especially important when considering material that is not accreted by the black hole, but is instead ejected from the system in powerful, highly collimated and sometimes highly relativistic jets perpendicular to the accretion flow. Such outflows are often found to accompany accretion but the conditions required to produce and power them are still not fully understood and still debated.

In this thesis I present five papers focussing on accretion and ejection around both stellar mass and supermassive black holes and what the emission from these systems might tell us about the black holes at their centres and their four physical parameters of mass, spin, accretion rate and inclination.

Chapter 2

Introduction: Stellar Mass Black Holes

Stellar mass black holes are formed from the collapse of massive stars. During most of a star's life, nuclear reactions in the core provide enough outward radiation pressure to balance the inward force of gravity. When the star runs out of fuel there is nothing to oppose the star's self gravity and it collapses.

If the core of the collapsing star is less than $\sim 1.4\,M_\odot$ (where $M_\odot$ is a solar mass), electron degeneracy pressure will eventually halt the collapse and the star will become a white dwarf (Chandrasekhar, 1931). If the star is more massive than this, electron degeneracy pressure will not be enough to halt the in-falling material. When the the outer layers hit the hard stellar core a shockwave is produced which propagates outwards causing the star to explode in a supernova. If the supernova remnant is less than $\sim 2-3 M_\odot$ neutron degeneracy pressure will be strong enough to prevent further collapse, and the remnant becomes a neutron star. If the stellar remnant is larger than $\sim 2-3 M_\odot$ then there is nothing left to prevent collapse and the result is a black hole (Chitre & Hartle, 1976; Kalogera & Baym, 1996).

Stars are rarely formed in isolation and many exist in binary systems, where two stars are gravitationally bound and orbit one another. If one of the stars goes supernova, leaving a black hole, material can be accreted onto the black hole from the companion. As the material is accreted it shines brightly in the X-rays, allowing these systems to be detected, and as a result they are termed black hole binaries

Figure 2.1: High/soft (red) and low/hard (blue) spectral states of Cyg X-1, from Gierliński et al. (1999).

(BHBs). The mass of the black hole can be estimated by measuring the orbital motion of the companion star and estimating the companion star's mass from its optical spectrum (e.g. Orosz et al., 2002 and references therein). The X-ray spectra of BHBs are not constant but vary with time. They can be classified in terms of spectral states, determined by the mass accretion rate onto the black hole.

2.1 Spectral States

The two canonical spectral states of BHBs are the low/hard state (LHS) and the high/soft state (HSS). The LHS is seen at low luminosities and corresponds to low mass accretion rates. As accretion rate increases, luminosity increases and the source may make a transition to the HSS. The accretion rate at which the transition occurs varies not only between sources, but also within the same source, and depends on whether the transition is from LHS to HSS or from HSS to LHS. The transition from LHS to HSS is generally at a higher luminosity ($\sim 0.10.2 \, L/L_{Edd}$, Done et al., 2007) while the transition from HSS back to LHS occurs at lower luminosity ($\sim 0.02 \, L/L_{Edd}$, Maccarone & Coppi, 2003).

Fig.2.1 shows an example of a LHS and a HSS (in blue and red, respectively) of the BHB Cyg X-1. The HSS (Fig.2.1, red points) is characterized by a strong

Figure 2.2: Hardness-intensity diagram showing spectra colour-coded by spectral state (HIMS, hard intermediate state; SIMS, soft intermediate state) from three outbursts of GX 339–4, from Plant et al. (2014).

thermal component in the soft X-rays peaking at $\sim 1\,\mathrm{keV}$, which dominates the total luminosity, and a weak non-thermal tail to high energies. The emission in the LHS (Fig.2.1, blue points) is dominated by a power law like tail peaking in the hard X-rays at $\sim 100\,\mathrm{keV}$. A lower temperature thermal component may be observed but the X-ray band pass is now dominated by the hard X-ray tail.

The fact the LHS is dominated by hard emission and occurs at low luminosities and the HSS is dominated by soft emission and occurs at high luminosities can be used to construct a hardness-intensity diagram, where a measure of the 'hardness' of the spectrum is plotted against its luminosity. Fig.2.2 shows a hardness-intensity diagram for the BHB GX 339–4. Three outbursts of the the source are shown, with each point representing an individual observation coloured according to its

spectral state. Prior to the outburst the black hole is in quiescence. As accretion rate increases it progresses up the right hand side of the diagram in the LHS until it reaches the top where it crosses to the left hand side and switches to the HSS. The luminosity then gradually drops again until the source moves back along the bottom horizontal track back to the LHS and finally drops down into quiescence. The diagram clearly demonstrates how the luminosity of the transition is variable; while the source generally always returns to the LHS along the bottom horizontal track, the luminosity at which it switches from LHS to HSS varies such that it may make the transition before reaching the maximum luminosity of the LHS branch by taking the middle track rather than the upper (Plant et al., 2014). This hysteresis which causes systems to switch from LHS to HSS at a higher luminosity and switch back at a lower luminosity has led to Fig.2.2 being referred to as a 'turtle head' diagram.

The HSS and LHS are not the only states BHBs show. Intermediate states are seen when sources make the fast transition from LHS to HSS and back again (green triangles, Fig.2.2). Some BHBs also show a very high state (VHS) at luminosities much higher than a typical HSS (Miyamoto et al., 1991). The VHS is characterised by high luminosity in both the soft thermal component and the hard power law and is often accompanied by strong jet emission.

A model proposed to explain these observed spectral changes with accretion rate is the truncated disc model (Done et al., 2007), illustrated in Fig.2.3. In this model the soft thermal component is black body (BB) emission from an accretion disc. In the HSS the disc extends down to the last stable orbit around the black hole. The non-thermal tail in the HSS is then produced by Compton up-scattering disc photons off hot electrons in an optically thin corona above the disc.

As accretion rate drops and the source enters the LHS, the disc recedes, evaporating from the inside out to be replaced by a hot inner flow (Meyer & Meyer-Hofmeister, 1994; Mayer & Pringle, 2007). In order to form an accretion disc the material must cool by radiating away its potential energy. The electrons emit the radiation but it is the more massive protons that carry most of the potential energy. In order for the material to cool there must be efficient transfer of energy

Figure 2.3: Schematic showing the proposed geometry of the accretion flow corresponding to the low/hard state, very high state and two high/soft state spectra (thermal dominant and ultra soft) according to the truncated disc model. In the high/soft state, the flow consists of an outer cool, geometrically thin, optically thick accretion disc, which extends down to the last stable orbit. As accretion rate drops in the low/hard state, the disc, which produces the soft thermal emission, is truncated at progressively larger radii. Within the truncation radius the flow is hot, geometrically thick and optically thin and produces the hard Compton emission. From Done et al. (2007).

between protons and electrons (Stepney, 1983). When accretion rate drops, the density of the accreting material drops and there are fewer collisions between protons and electrons. As a result the material can no longer cool into a geometrically thin, optically thick disc, but remains hot, geometrically thick and optically thin. A fraction of the accretion energy is retained by the protons and advected with the flow, leading to these hot flows being termed Advection Dominated Accretion Flows (ADAFs, Narayan & Yi, 1995). The remains of the outer accretion disc provide a low temperature soft BB component, but the X-rays are now dominated by emission from the hot flow, where hot electrons Compton up-scatter seed photons from the disc and cyclo-synchrotron emission produced from magnetic fields within the hot flow, to provide the hard power law tail. The brightest LHS spectra have an optical depth (τ) inferred from their spectra of $2 - 3$, corresponding to an accretion rate of $\dot{m} = L/L_{Edd} \sim 0.1$ (Ibragimov et al., 2005, Malzac & Belmont, 2009), implying this is the critical density at which energy transfer between ions and electrons becomes inefficient and the hottest material nearest to the black hole can no longer cool and condense into a disc.

In the context of the truncated disc model, the VHS should represent the highest accretion rates, where there is so much inflowing gas that the disc is throwing off material, potentially disrupting its structure, and making both a strong disc and corona, as well as expelling excess material via the jet.

Modelling the spectra of BHBs therefore involves modelling both the emission from an optically thick accretion disc and from hot electrons in the inner hot flow.

2.1.1   Disc Emission

The radiation from a thin accretion disc can be estimated by considering the gravitational potential energy liberated at each radius (R):

$$\frac{dE}{dt} = \frac{GM\dot{M}}{R^2}dR \tag{2.1.1}$$

where M is the mass of the black hole and $\dot{M}$ is the accretion rate. Following the virial theorem, half of this can be radiated, giving:

$$dL = \frac{GM\dot{M}}{2R^2}dR \tag{2.1.2}$$

If the radiation thermalises then setting this equal to $dL = dA\sigma T^4$, where σ is the Stefan-Boltzmann constant, gives:

$$T^4 = \frac{GM\dot{M}}{8\pi\sigma R^3}dR \tag{2.1.3}$$

Each radius emits as a black body with temperature increasing with decreasing radius. The peak of the disc spectrum is therefore determined by the inner radius of the disc. In the HSS this will be at the last stable orbit around the black hole (R_{isco}). In the LHS this will be the radius at which the disc is truncated (R_{trunc}) and replaced by the hot inner flow, as in Fig.2.4.

This is the standard Shakura & Sunyaev (1973) disc solution. If the disc extends to down to R_{isco} then a stress free inner boundary condition must be included. This accounts for the fact there is no viscosity inside R_{isco}. Small relativistic corrections to the emissivity can also be included following Novikov & Thorne (1973). A colour

Figure 2.4: Model disc spectrum around a $10\,M_\odot$ black hole with $R_{trunc} = 20\,R_g$. Each annulus of the disc radiates as a black body with temperature increasing as radius decreases. The peak disc luminosity is therefore determined by the inner radius of the disc.

temperature correction can also be included, which accounts for the fact the disc emission may not completely thermalise at every radius. This is because the absorption opacity depends on temperature, density and the energy of the radiation, all of which change with radius, while the electron scattering opacity is constant. For hot inner radii, the absorption opacity can be much less than the scattering opacity, causing a higher effective temperature than otherwise predicted, so the resulting modified BB spectrum will be hotter (Shimura & Takahara, 1995; Done et al., 2012).

Nonetheless, assuming all the energy (more or less) thermalises, so that the energy released at each radius is radiated locally, remains a good approximation. This requires the disc is optically thick and results in a geometrically thin disc ($H/R \sim 0.01$, where H is the vertical scale height of the disc). As material accretes inwards, angular momentum is transported outwards by viscosity. The material is likely to be threaded with magnetic fields and differential rotation causes field lines connecting material at different radii to become stretched and tangled. Simulations show this magneto-rotational instability (MRI; Balbus & Hawley, 1991) can provide

the stresses required to transport angular momentum outwards.

2.1.2 Hot Flow Emission

In the hot flow within a truncated disc (as in the corona above a HSS disc) energy transfer between ions and electrons is inefficient, due to the lower density of the material. As a result the hot flow forms a two temperature plasma. The ion temperature will be set by the virial temperature, such that:

$$kT_{ion} = \frac{GMm_p}{R} = \frac{m_p c^2}{r} \tag{2.1.4}$$

where r is the radius in units of gravitational radii ($R_g = GM/c^2$) and n is the density of accreting material, related to the optical depth (τ) by $n = \tau/(\sigma_T R_{trunc})$. This gives ion temperatures of 10^{12} K, whilst the radiating electrons are much cooler, remaining at 10^9 K (Narayan & Yi, 1995).

Just like the disc, the hot flow also contains tangled magnetic fields. An estimate of the magnetic field strength can be made by assuming the magnetic energy density (U_B) saturates at 10% of the gas pressure:

$$U_B = \frac{B^2}{8\pi} = 0.1 n k T_{ion} \tag{2.1.5}$$

The radiative inefficiency of the hot flow means that it will only radiate a fraction, f, of the gravitational energy radiated by a disc spanning the same radii:

$$L_{hot} = f L_{disc}(R_{trunc} < R < R_{isco}) \tag{2.1.6}$$

The hot flow therefore consists of tangled magnetic fields and hot electrons with an input power of L_{hot}: it will produce Comptonised emission from up-scattering disc photons, and will also produce cyclo-synchrotron radiation through the interaction of electrons with the magnetic field.

2.1.2.1 Thermal Comptonisation

The energy of hot flow electrons is much greater than that of any photons from the disc. A disc photon passing through the hot flow which happens to collide with an electron will therefore gain energy and be Compton up-scattered. The energy of the photon after the collision (E_2) depends on its original energy (E_1), the Lorentz factor of the electron ($\gamma = (1 - \beta^2)^{-1/2}$) and the angles between the electron and original photon ($\theta_{e,1}$), electron and final photon ($\theta_{e,2}$) and the original and final photons ($\theta_{1,2}$) (Rybicki & Lightman, 1979):

$$E_2 = \frac{E_1(1 - \beta \cos \theta_{e,1})}{1 - \beta \cos \theta_{e,2} + (E_1/\gamma)(1 - \cos \theta_{1,2})} \tag{2.1.7}$$

Averaging over angles, assuming an isotropic distribution of original seed photons and $\theta_e \gg 1$, gives an average energy gain by each photon of $(1 + 4\Theta)$, where $\Theta = kT_e/m_e c^2$ is the dimensionless electron temperature. The up-scattered photons can go on to collide again with more electrons, gaining $(1 + 4\Theta)$ each time until they reach 3Θ. They cannot be up-scattered further than this since the photons cannot gain more energy than the electrons began with.

However the photons will only be up-scattered by the electrons if they collide. The probability of a photon colliding with an electron depends on the optical depth as $e^{-\tau}$, so only a fraction of the photons will be boosted on to higher energies each time. This succession of scattering orders, with a constant fraction, τ, of the photons being boosted on by a constant amount, $(1 + 4\Theta)$, each time gives the power law seen in LHS spectra, with the roll-over of the Compton spectrum determined by the electron temperature. This process is shown schematically in Fig.2.5a.

This has led to LHS spectra often being modelled by a power law with an exponential cut-off at the electron temperature. However, a simple power law is not always a good approximation to the true Comptonised spectrum. The actual cut-off is steeper than an exponential, and there will be a bump in the spectrum at the seed photon energy, due to the fraction of seed photons that escape without scattering, and a downturn below the seed photon energy which must also be modelled (Kubota & Done, 2004). Moreover, at very low mass accretion rates, such as those seen in

Figure 2.5: Schematic showing formation of Compton power law from successive electron scatterings. Red line shows the seed photon population, a fraction τ of these are up-scattered by energetic electrons to form the first scattering order (blue), a fraction τ of these photons are then up-scattered again, with each scattering boosting the photons by $1 + 4\Theta$. In panel a). a high optical depth produces a smooth power law spectrum, while the lower optical depth in b). results in a bumpy spectrum. From Done (2010).

dimmer LHS, the optical depth is expected to be $\ll 1$. At such low optical depths, the probability of scattering is less and the individual scattering orders become separated, giving a spectrum that is bumpy (Fig.2.5b). Even far from the seed photon energy and high energy cut-off, a power law is no longer a good description of the spectrum.

A more accurate model of Comptonisation can be produced using the code EQ-PAIR (Coppi, 1999). This calculates the emission from a homogeneous sphere, given inputs of the total heating power to the electrons (L_{hot}), the optical depth and size of the region, and the power and typical energy of the seed photons for Compton cooling (L_{seed} and kT_{seed}). EQPAIR derives the electron temperature self-consistently from the cooling rates and can accommodate the low optical depths expected at low mass accretion rates. It also provides a more intuitive way to understand Comptonisation: rather than considering the spectrum in terms of optical depth and electron temperature, a better framework is in terms of optical depth and the ratio L_{hot}/L_{seed}. More seed photons ($L_{hot}/L_{seed} \sim 1$) means more Compton cooling, a lower electron temperature and a softer Compton spectrum. Fewer seed photons ($L_{hot}/L_{seed} > 1$),

Figure 2.6: Example Compton spectra of Cyg X-1 showing spectral softening and decreasing electron temperature from red to blue as the disc moves inwards and seed photon flux increases, from Ibragimov et al. (2005).

i.e. if the disc has receded, means the electrons remain hotter and the spectrum is harder. The photon index (Γ) is often used as a measure of the hardness of an X-ray spectrum, where Γ is defined as $\nu L(\nu) \propto \nu^{-\Gamma+2}$ and $\Gamma > 2$ is defined as soft and $\Gamma < 2$ as hard. Fig.2.6 shows three Compton spectra from Cyg X-1, showing how the spectrum softens and electron temperature (i.e. Compton roll-over) decreases as the seed photon flux increases, due to the disc moving inwards and increasing the amount of cooling.

In practice, approximating the hard emission as a single Compton spectrum can be overly simplistic. Regions of the hot flow nearer the disc will receive a greater flux of seed photons so will emit a softer Compton spectrum than central regions, which intercept fewer seed photons. In other words, Comptonisation is probably inhomogeneous. This is most clearly the case in BHBs approaching the transition from LHS to HSS, for which the hard Compton tail can be better modelled as two

Figure 2.7: Sequence of low/hard state spectra from Cyg X-1, from Yamada et al. (2013). Spectra are numbered according to the order in which they were observed, but presented in order of decreasing luminosity. Grey line shows spectrum from observation 3 for comparison. Fitting the brightest states requires both a hard and a soft power law, suggesting Comptonisation is inhomogeneous.

power laws — a soft power law, which increases in luminosity as the disc undercuts the hot flow and increases the flux of seed photons in the outer regions, and a hard power law from the remaining central regions. Fig.2.7 shows a sequence of spectra from Cyg X-1 for which this is clearly the case in the brightest spectra (Yamada et al., 2013).

2.1.2.2   Cyclo-Synchrotron

In the presence of a magnetic field, electrons will be accelerated and if an electron accelerates it will radiate. If the electrons are travelling at non-relativistic velocities then the emission is termed cyclo-synchrotron radiation. The process is essentially the same as Comptonisation, except that the seed photon is now a virtual magnetic

Figure 2.8: Cyclo-synchrotron spectrum for a hot flow of radius $20\,R_g$. Dotted line shows unabsorbed spectrum, solid line shows spectrum taking into account self-absorption.

field photon, which interacts with a hot electron to produce a 'real' photon.

Cyclo-synchrotron emission forms a steep power law. However, most of the emission is self-absorbed, so the spectrum actually peaks at the self-absorption frequency (ν_{csa}). Fig.2.8 shows an example of a self-absorbed cyclo-synchrotron spectrum from a hot flow with a radius of $20\,R_g$. This internally generated cyclo-synchrotron emission provides a second source of seed photons for Comptonisation in the hot flow. It is particularly important in the inner regions of the flow, which are shielded from disc flux when $\tau \gtrsim 1$, and also when the disc truncation radius is large and consequently the number of disc seed photons intercepted by the hot flow is low.

Self-absorption greatly reduces the amount of cyclo-synchrotron emission that is produced. So much so that a very high electron temperature — roughly twice as high as is observed — is required to produce enough seed photons to make a LHS spectrum using a population of thermal electrons (Yuan et al., 2007). However, if the electron population has a non-thermal tail, cyclo-synchrotron emission becomes much more efficient and can reproduce observed spectra with electron temperatures in agreement with the observations (Wardziński & Zdziarski, 2001; Veledina et al.,
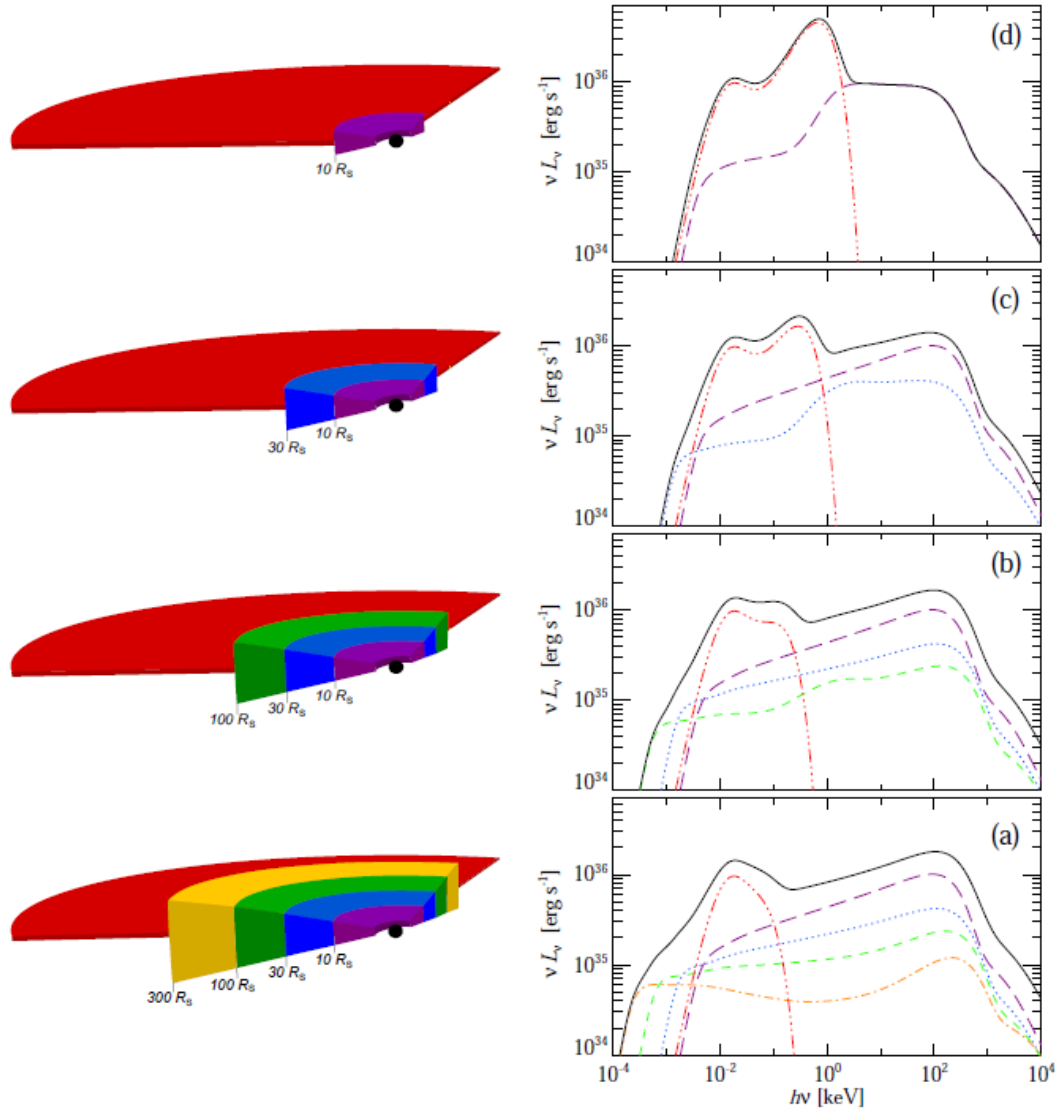
Figure 2.9: Schematic of the evolution of the hot flow spectrum with hot flow size during a transition from high/soft state to low/hard state. Red line shows contribution of the outer disc (including irradiation), while other lines show the cyclo-synchrotron-self-Compton emission from each region of the hot flow, with the zone closest to the disc also including Comptonisation of disc seed photons, from Poutanen & Veledina (2014).

2011; Veledina et al., 2013). The existence of an MeV power law tail in the LHS of Cyg X-1 (Zdziarski et al., 2012) and a weak power law tail in the HSS (Grove et al., 1998; Gierliński et al., 1999; Zdziarski et al., 2001; McConnell et al., 2002) supports the theory that non-thermal electrons are present. In which case the hot flow should contain a hybrid distribution of electrons, with a main population of thermal electrons heated by Coulomb collisions with protons, plus a proportion of non-thermal electrons accelerated by shocks or magnetic reconnection (Fragile & Blaes, 2008; Das et al., 2009; Henisey et al., 2012; Ding et al., 2010; Riquelme et al., 2012; Hoshino, 2013). Fig.2.9 shows how the spectrum from such a hybrid electron distribution changes as the the disc recedes and the size of the hot flow increases.

2.1.3 Reflection

A fraction of the hard coronal flux will illuminate the outer disc. These photons can scatter off electrons in the disc and be reflected. This adds a third spectral component to the intrinsic disc and coronal emission. The amount of reflection depends on the relative importance of electron scattering versus photoelectric absorption.

For neutral material, photoelectric absorption dominates at low energies so there is very little reflected flux below $10\,\text{keV}$. The cross section for photoelectric absorption decreases as the energy of the incident photons increases. Consequently the reflected flux increases and reaches a peak between $20 - 50\,\text{keV}$, called the reflection hump (George & Fabian, 1991). Above this the reflected flux decreases again due to Compton down-scattering; the highest energy photons scatter inelastically, losing some of their energy to the electrons and so the reflected photon has lower energy.

Fig.2.10 shows a reflection spectrum from neutral material, together with the illuminating power law. In addition to the reflection hump, the reflection spectrum also shows emission lines. The incident electrons that are not reflected are photo-electrically absorbed, exciting the ions in the disc, which then decay and re-emit the energy as emission lines. The most strongest line is generally $K\alpha$, where the excited electron drops from shell $n = 2$ to $n = 1$. Since photoelectric absorption dominates at low energies this is where the emission lines occur. Most are not visible in the total spectrum since they are heavily diluted by the illuminating continuum. In

Figure 2.10: Reflection spectrum from neutral material with solar abundance and column $N_H = 1.5 \times 10^{24}\,\mathrm{cm}^{-2}$ (blue) together with the original illuminating continuum (grey), from Done (2010).

addition, low atomic number elements also have a high probability of Auger ionization de-excitation, whereby the excited electron is ejected rather than decaying to a lower energy level. Iron has a high atomic number, so fewer Auger losses, and a high atomic number also means its lines are emitted at higher energies where the reflection spectrum is less diluted. The iron $K\alpha$ line at ~ 6.4keV is therefore the most prominent emission line in neutral reflection spectra (e.g. Suzuki et al., 1984).

If the material is not neutral but ionised, the shape of the reflection spectrum will change. This is because the amount of photoelectric absorption at low energies depends on the ionisation state of the reflecting material. As the ionisation state increases, the cross section for photoelectric absorption decreases, so more of the low energy flux will be reflected. Fig.2.11 shows how the reflection spectrum changes for increasing ionisation parameter, ζ. The absorption edges and emission lines also broaden as ζ increases as there are more free electrons so more free-electron scattering, which alters the photon energies and causes broadening.

The emission lines from BHB reflection spectra are not just broadened by electron

Figure 2.11: Ionised reflection spectra from material with increasing ionisation parameter (ζ), assuming solar abundance and column $N_H = 1.5 \times 10^{24}\, \mathrm{cm}^{-2}$, with original illuminating continuum shown in grey, from Done (2010).

scattering, but also by relativistic effects (Fabian et al., 1989). The accretion disc is rapidly rotating in a strong gravitational field. Photons from the approaching side of the disc are blue shifted and Doppler boosted while those from the receding side are red shifted and deboosted. Both these effects are a function of inclination angle of the disc with respect to our line of sight, while the additional effects of time dilation and gravitational redshift cause line broadening in systems at any inclination. Fig.2.12 shows how each of these effects contributes to the final broadened line. The smaller the radius from which the line is emitted, the faster the material is travelling and the broader the line. Hence line broadening and line shifts can be used to estimate the radius of the inner edge of the accretion disc and from this estimate the spin of the black hole if the disc extends down to the innermost stable orbit (e.g. Reis et al., 2009).

The reflection contribution to the spectrum will be at its strongest when the disc subtends the largest possible angle with respect to the illuminating continuum, allowing it to intercept the maximum amount of flux, i.e. when the disc inner radius

Newtonian

Special relativity
Transverse Doppler shift
Beaming

General relativity
Gravitational redshift

Line profile

0.5   1   1.5

ν_{obs}/ν_{em}

Figure 2.12: Schematic illustrating how Doppler and relativistic effects contribute to produce a broadened emission line profile, from Fabian et al. (2000).

is small (Gierliński et al., 1999). This is also when the iron line will be broadest as it will be dominated by reflection at the smallest radii. Fig.2.13 shows a sequence of spectra demonstrating the strengthening of the reflection component and broadening of the iron line as luminosity increases and the disc inner radius decreases.

2.1.4   Reprocessing

The fraction of hard coronal flux that intercepts the disc and is not reflected or emitted as lines will be reprocessed. The photon energy goes into heating the disc which then re-radiates this as quasi-thermal emission that adds to the observed disc spectrum (Malzac et al., 2005). If the disc is energetically dominant and the hard

Figure 2.13: Sequence of low/hard state spectra ordered by luminosity showing the changing strength of the reflection spectrum; brighter spectra have smaller disc truncation radii so stronger reflection. The bottom panel shows the ratio of the spectrum to an absorbed power law, clearly showing stronger deviations caused by reflection at soft energies and at the iron line in the brighter spectra, from Kolehmainen et al. (2014).

Figure 2.14: Model showing increasing inner disc reprocessing with increasing ratio of Compton/disc flux. Grey arrow shows the direction of increasing L_c/L_d and dashed line shows disc with no irradiation, from Gierliński et al. (2008).

power law tail is weak then this additional reprocessed flux will only increase the intrinsic disc emission by a very small amount. However if the hard power law carries a large fraction of the bolometric luminosity, e.g. in the LHS when the disc is truncated, then the reprocessed emission can contribute a significant fraction of the observed disc flux.

Fig.2.14 shows how the disc spectrum changes as the fraction of energy in the hard power law increases, increasing the amount of reprocessing on the inner edge of the disc. The illumination heats the inner edge of the disc, causing it to appear brighter and more luminous, hence it is important to account for reprocessing when fitting spectra to derive inner disc radii (Gierliński et al., 2008).

The hot inner disc will in turn illuminate the outer disc. For large disc radii, which are cool and radiate much less intrinsic flux, this reprocessed emission can

Figure 2.15: Sequence of observations of the BHB XTE J1817—330 as it makes the transition from high/soft state (observations 1, 8 and 13) to low/hard state (observations 20 and 22), showing how the inner and outer disc reprocessing increases as the fraction of flux in the power law increases. Blue lines show hard coronal power law, red lines intrinsic disc emission, from Gierliński et al. (2009).

dominate the intrinsic disc flux, giving the disc spectrum a distinctive shoulder in the optical region of the spectrum (Gierliński et al., 2009). Fig.2.15 shows how both the inner and outer disc reprocessing increases as the BHB XTE J1817—330 makes the transition to the LHS.

2.2 Variability

BHBs do not just show long term variability on day–month timescales in terms of spectral state changes. Within a given spectral state they also show much shorter timescale variability, with luminosity fluctuations occurring on ms–100 s timescales. Spectral state change are due to the entire structure of the accretion flow altering in response to global changes in mass accretion rate. Short timescale fluctuations

are probably also caused by mass accretion rate fluctuations, but on a much smaller scale. The structure of the accretion flow cannot change on $100\,$s timescales, but the MRI — the mechanism which transports angular momentum outwards and therefore regulates the flow of material inwards — is stochastic, causing the mass accretion rate between adjacent radii to constantly vary. As a result the X-ray emission from BHBs shows aperiodic variability over a range of timescales — their accretion flows 'flicker'.

The mass accretion rate fluctuations at a given radius (r) will occur around the viscous frequency (f_{visc}) for that radius:

$$f_{visc} = \alpha(H/R)^2 f_{dyn} = \frac{\alpha\,(h/r)^2\,c}{2\pi R_g(r^{1.5} + a)} \tag{2.2.8}$$

where f_{dyn} is the dynamical frequency, α is the viscosity parameter, H is the vertical scale-height at that radius and a is the spin of the black hole (Shakura & Sunyaev, 1973). It is immediately clear that the scale-height of the accretion flow is important. The smaller (h/r), i.e. the thinner the accretion flow, the smaller f_{visc}, so the longer the characteristic timescale of fluctuations from a given radius. (h/r) is typically 0.01 for an optically thick, geometrically thin accretion disc, but can be as large as 0.2 for an optically thin, geometrically thick hot flow. This means a soft BHB accretion disc typically varies on timescales $\gg 1000\,$s while the hard hot flow emission varies on timescales of the order $\sim 1\,$s. Consequently for a typical $1\,$ks observation, the soft emission should appear much less variable than the hard emission (Churazov et al., 2001). Fig.2.16 shows that this is indeed the case. At low accretion rates the source is in the LHS, the count rate is low and the emission is dominated by the variable power law from the hot flow, so that while the source brightens along the 'hard line' it shows consistently high fractional variability $(\sim 30 - 40\%)$. As soon as the source transitions to the HSS at high count rates, the X-ray band becomes dominated by the stable disc and the variability drops significantly to $< 5\%$ (Plant et al., 2014).

Figure 2.16: Root mean square variability versus count rate for multiple observations of GX339–4, from Plant et al. (2014).

2.2.1 RMS–Flux Relation

A useful way to quantify variability is to simply calculate the root mean square variability (RMS) of a light curve. RMS is defined as:

$$\sigma = \sqrt{\frac{\sum_{i=1}^{N}(x_i - I)}{N - 1}} \tag{2.2.9}$$

where x_i is the flux at time t_i, I is the mean flux and N is the total number of timesteps in the light curve (e.g. Nowak et al., 1999). The fractional variability of the light curve is then $F_{var} = \sigma/I$.

BHBs obey a linear RMS–flux relation (Fig.2.17). If a light curve is split into equal length segments and the RMS and mean flux of each segment is calculated, the size of the RMS fluctuations and the mean flux are linearly related so that σ/I is a constant.

A linear RMS–flux relation cannot be produced by summing independent events but requires a multiplicative process (Uttley & McHardy, 2001). This is easily

Figure 2.17: RMS–flux relation for Cyg X-1, from Uttley & McHardy (2001).

provided in the context of an accretion flow, because low frequency mass accretion rate fluctuations produced at larger radii propagate inwards. The fluctuations at smaller radii therefore consist of the long timescale fluctuations propagating inwards from outer radii modulated by the faster fluctuations generated at smaller radii (Lyubarskii, 1997).

2.2.2 Power Spectra

Another way to quantify variability is to calculate the power spectrum of the light curve. This is the modulus squared of the Fourier transform of the source's light curve and shows how much power there is at different frequencies. Each radius has a different viscous frequency (with viscous frequency increasing with radius for a constant scaleheight). The power spectrum of the fluctuations generated at each radius is a Lorentzian centred around the viscous frequency for that radius. Since the total light curve consists of emission from all radii, its power spectrum will be a sum of Lorentzians generated at all relevant radii.

The hard band light curves of LHS BHBs therefore show 'flat topped noise'

$(P(f) \propto f^{-1}$, i.e. $fP(f) \propto f^0)$ between a low frequency break, f_l (which is moveable), and a high frequency break, $f_h \sim 5\text{Hz}$ (Remillard & McClintock, 2006). Below the low frequency break the power scales as $P(f) \propto f^0$ and above the high frequency break it scales as $P(f) \propto f^{-2}$. In the context of the truncated disc model these two frequency breaks can be understood as representing the two boundaries of the hot flow. The low frequency break corresponds to the viscous frequency of the largest radius of the hot flow, i.e. where the hot flow meets the truncated disc. The high frequency break corresponds to the viscous frequency at the innermost stable circular orbit (Done et al., 2007).

The top panel in Fig.2.18 shows an example of a LHS power spectrum from XTE J1550–564. The following panels show how the power spectrum changes as the source makes a transition towards the HSS (Ingram & Done, 2012a). The high frequency break, corresponding to fluctuations generated at R_{isco}, does not change much. In contrast the low frequency break increases as the disc inner radius moves in, and the larger radii of the hot flow, which produce the lowest frequency fluctuations, are replaced by a stable (on these timescales) disc (Churazov et al., 2001).

The bottom panel in Fig.2.18 shows a strong peak at $\sim 4\text{Hz}$ in addition to the flat topped noise. A similar peak can be seen in the other power spectra, with the peak decreasing in strength as truncation radius increases. These are the signatures of quasi-periodic oscillations (QPOs; van der Klis, 1989). Although the majority of the variability is stochastic and aperiodic, the light curves of LHS BHBs can show additional quasi-periodic behaviour. Sinusoidal flux variations, which often slip out of phase, can be seen in their light curves and produce characteristic narrow peaks in their power spectra. One model to explain this phenomenon is that these signals arise due to frame dragging caused by the black hole spin axis being misaligned with the axis of rotation of the accretion flow (Ingram et al., 2009; Ingram & Done, 2011). This causes the hot inner flow to precess. Since the sound speed in the hot flow is faster than the precession speed, pressure waves allow the entire hot flow to precess as a solid body. This precession changes the orientation of the hot flow with respect to our line of sight producing periodic variations in flux. The characteristic frequency of these variations is related to the outer radius of the hot

Figure 2.18: Power spectra of XTE J1550–564 during a transition from the low/hard state (top panel) to the high/soft state (bottom panel), together with the truncation radii predicted from fitting a truncated disc model, from Ingram & Done (2012a).
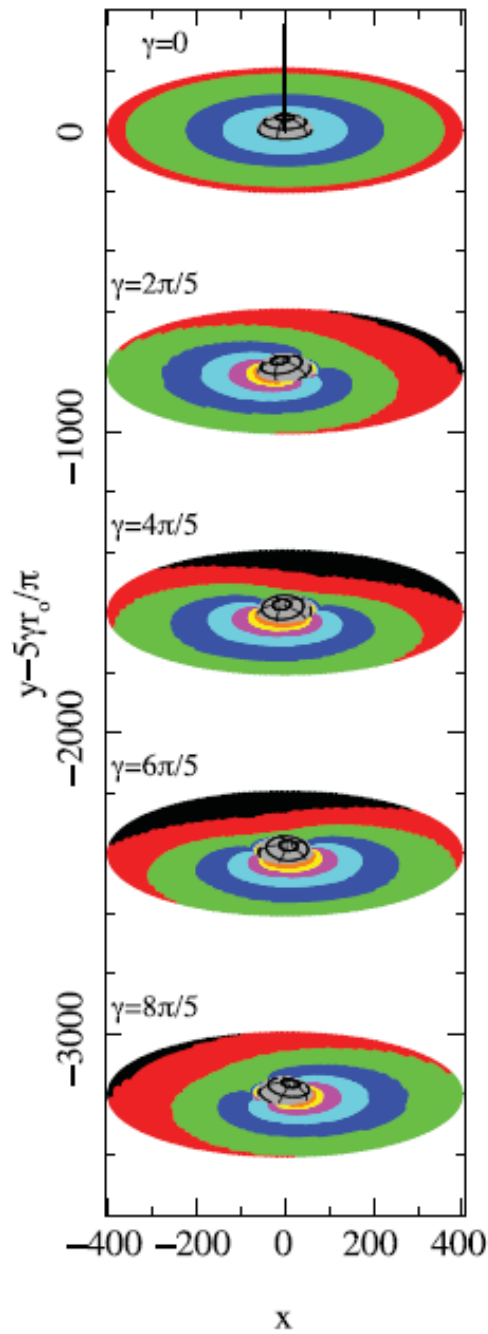
Figure 2.19: Changing hot flow orientation and disc illumination due to precession of the hot flow, caused by misalignment of the black hole spin axis and the axis of rotation of the accretion flow. Black line shows black hole spin axis, grey region inner hot flow and orange to black regions show brightest to dimmest regions of disc illumination, from Ingram & Done (2012b)

flow, hence the QPO frequency increases with the low frequency break as the disc moves inwards, as is seen in Fig.2.18. The sound speed is much slower in an optically thick, geometrically thin accretion disc, hence the disc warps rather than precessing and the QPO disappears when the inner hot flow is replaced by a stable disc in the HSS. Fig.2.19 illustrates this precession effect and also shows how the precessing flow alters the illumination pattern on the disc, which should cause the reflection spectrum to subtly change with QPO phase (Ingram & Done, 2012b).

Fig.2.20 (top panel) shows a collection of LHS power spectra in comparison with power spectra typical of other spectral states (Done & Gierliński, 2005). The bottom panel shows the corresponding energy spectra for reference. The extreme LHS in Fig.2.20c is an example of a spectrum where the disc truncation radius is extremely large, so that the hard power law produced by the hot flow accumulates fluctuations out to very large radii, giving a particularly low low frequency break. The LHS in Fig.2.20a and b have smaller truncation radii and so the flat topped fluctuations are limited to a narrower range of frequencies. In the HSS shown in Fig.2.20e, the disc extends down close to the last stable orbit so the corona is confined to only the smallest radii and the bulk of the power that remains is at high frequencies. The VHS in Fig.2.20f shows more power over a wider range of frequencies than the HSS due to the size of the corona increasing again as it extends out above the disc.

Almost all BHBs show HSS power spectra similar to Fig.2.20e. The exception is Cyg X-1, the HSS power spectrum of which is shown in Fig.2.20d, which instead shows a flat topped noise component with no low frequency break (Axelsson et al., 2005). Fig.2.21 shows how this flat topped component, not seen in other BHBs, increases in power as the source approaches the HSS until it dominates the power spectrum. Where this extra variability component comes from is not understood, however it implies that even in the HSS the variable coronal emission is still able to accumulate fluctuations from a wide range of radii. Perhaps the corona extends out above the disc in the HSS of Cyg X-1, rather than being confined to the innermost radii as in other BHBs. One potential reason for this may be that Cyg X-1 never fully makes the transition to the HSS.

Figure 2.20: Examples of typical power spectra and corresponding energy spectra during different spectral states, from Done & Gierliński (2005). Power spectra: black shows full $2 - 60$ keV power spectrum, blue shows power spectrum from just $13 - 25$ keV energy band. Energy spectra: red shows disc, blue shows Comptonisation, green reflection, solid magenta shows absorbed total and dashed magenta unabsorbed.

Figure 2.21: Power spectral evolution of Cyg X-1 during a transition from low/hard state to high/soft state, from Axelsson et al. (2005); Done et al. (2007).

2.2.3 Time Lags

It takes time for mass accretion rate fluctuations to propagate inwards. Since the highest energy emission is produced at the smallest radii this means higher energy bands should progressively lag softer energies.

Fig.2.22 shows the lag as a function of Fourier frequency for four increasingly hard energy bands with respect to a soft energy band ($0 - 3.9$ keV; Nowak et al., 1999). Slow frequency fluctuations ($\sim 0.1 - 1$ Hz) are generated at large hot flow radii and propagate inwards to the smaller radii producing the higher energy emission. Consequently the soft emission brightens first and then the higher energy emission follows when the accretion rate fluctuation has propagated down to the relevant radii. The lag of the softest hard band ($3.9 - 6$ keV) behind the soft band at 1 Hz is therefore shorter than that of the hardest hard band ($14.1 - 45$ keV) behind the same soft band (~ 0.004 s compared to 0.025 s) since the fluctuation must travel further to reach the smallest radii producing the hardest flux. That hard bands progressively lag soft bands supports the idea of inward propagation of fluctuations. For each lag-frequency spectrum, the fact the lags between hard and soft bands decrease as

Figure 2.22: Lag-frequency spectra between different energy band light curves of Cyg X-1 in the low/hard state, from Nowak et al. (1999).

Figure 2.23: Lag-energy spectra of GX 339–4, where the lag of each energy bin is calculated relative to the light curve over the total energy band such that energy bins with more positive lag values lag those with less positive values, from Uttley et al. (2011).

Fourier frequency increases indicates the low frequency fluctuations are originating in the soft component while high frequency fluctuations are produced by the hard emission regions, in agreement with the hard emission coming from smaller radii which can fluctuate faster.

However, not all of the soft emission leads the hard emission. Uttley et al. (2011) showed that the high frequency soft band fluctuations lag the hard band rather than leading it. This is due to thermal reprocessing: some of the hard coronal flux illuminates the inner parts of the disc, heating it up. The disc then re-radiates this energy as thermal emission at soft energies.

Fig.2.23 shows the lags as a function of energy for four different frequency ranges, measured during a LHS of GX 339–4 (Uttley et al., 2011). The lag of each energy bin is now measured relative to the total light curve summed over the whole energy range. Energy bins with more positive lag values lag behind those with less positive lag values. At low frequencies ($0.034 - 0.12 \, \mathrm{Hz}$), the soft energy bins ($< 1 \, \mathrm{keV}$) lead the hard energy bins, due to low frequency fluctuations being generated at large hot flow radii and in the disc, which together produce the softest emission. However at high frequencies ($2 - 8 \, \mathrm{Hz}$), all the fluctuations are generated at small radii, which produce the hardest emission, and this heats the soft disc causing it to re-radiate the energy as quickly varying soft emission, which lags behind the illuminating hard flux with a lag that corresponds to the light travel time from illuminating corona to soft disc.

2.3 Jets

BHBs do not only emit at high energies. Outer parts of the accretion disc emit at optical wavelengths and are frequently capable of outshining the companion star at high mass accretion rates. However still lower energy emission is detected from many BHBs — in the infra-red and radio regimes.

Even the outermost parts of a BHB disc — limited in size by the binary separation — are too hot to radiate at these wavelengths, although cyclo-synchrotron emission can make some contribution. However the bulk of this low energy emission

Figure 2.24: Each section of a conical jet produces self-absorbed synchrotron emission. Regions further out (red lines) are larger so have a lower self-absorption frequency. The sum of emission from all regions produces a flat spectrum in the radio ($L_\nu \propto \nu^0$ or $\nu L_\nu \propto \nu$ at $\sim 10^{10}$ Hz).

comes from material that is not accreted but is instead ejected from the system in the form of relativistic jets.

BHBs do not universally show jet emission. The presence or absence of a jet is related to the accretion state. In the LHS there is a steady jet and the radio and X-ray luminosities are correlated ($L_R \propto L_X^{0.7\pm0.1}$, Corbel et al., 2003). This is known as the radio–X-ray correlation. Measurements of the bulk Lorentz factor (BLF) of these jets find mildly relativistic values of $\Gamma \sim 1.4$ (Fender et al., 2004). In fact, for the radio–X-ray correlation to hold precludes these jets having highly relativistic BLFs ($\Gamma \sim 10$), since different orientations would result in different amounts of Doppler boosting, resulting in a large range in L_R for a given L_X, assuming the X-rays are produced by the accretion flow. The radio spectrum itself is flat and is produced by successive regions of self-absorbed synchrotron emission, as shown in Fig.2.24 (Blandford & Königl, 1979). Just as for the cyclo-synchrotron emission from the hot flow, this is produced by electrons spiralling around magnetic field lines, with the difference that the magnetic fields are now in the jet and the electrons now have

Figure 2.25: Radio image of the steady, milliarcsecond-scale, low/hard state jet of Cyg X-1 (left panel; Stirling et al., 2001) compared to a transient, arcsecond-scale, state transition jet from GRS 1915+105 (right panel; Mirabel & Rodríguez, 1994), from Gallo (2010)

relativistic energies. Magnetic reconnection and shocks within the jet are assumed to accelerate the electrons and their synchrotron emission is further boosted by the jet bulk motion.

As the system moves towards the HSS, the BLF of the jet is observed to increase to $\Gamma \sim 2$. This produces internal shocks within the jet, as faster material catches up with the previous slower ejecta, and optically thin flares are observed (Fender et al., 2004), which may be seen as individual knots moving away from the jet core (Mirabel & Rodríguez, 1994). Fig.2.25 shows an example of such a transient jet from GRS 1915+105, compared to the steady LHS jet of Cyg X-1.

When the source enters the HSS, the radio emission switches off. This is known as crossing the 'jet line'. Fig.2.26 shows a schematic of the track a BHB might take on a hardness-intensity diagram during a transition from the LHS to HSS and back again. The schematic also includes tracks taken by those systems that attempt to make transitions between the HSS and LHS, but fail. The jet line matches the transition from LHS to HSS.

Once in the HSS no jet is produced (or alternatively it is so weak that it is not detected) and there is no radio emission. Fig.2.27 shows the radio–X-ray correlation in GX 339–4 and the suppression of radio emission once the source crosses the jet line and reaches the HSS at high X-ray luminosities, together with the switching back on of the jet when the source drops back down into the LHS at lower luminosity (Corbel
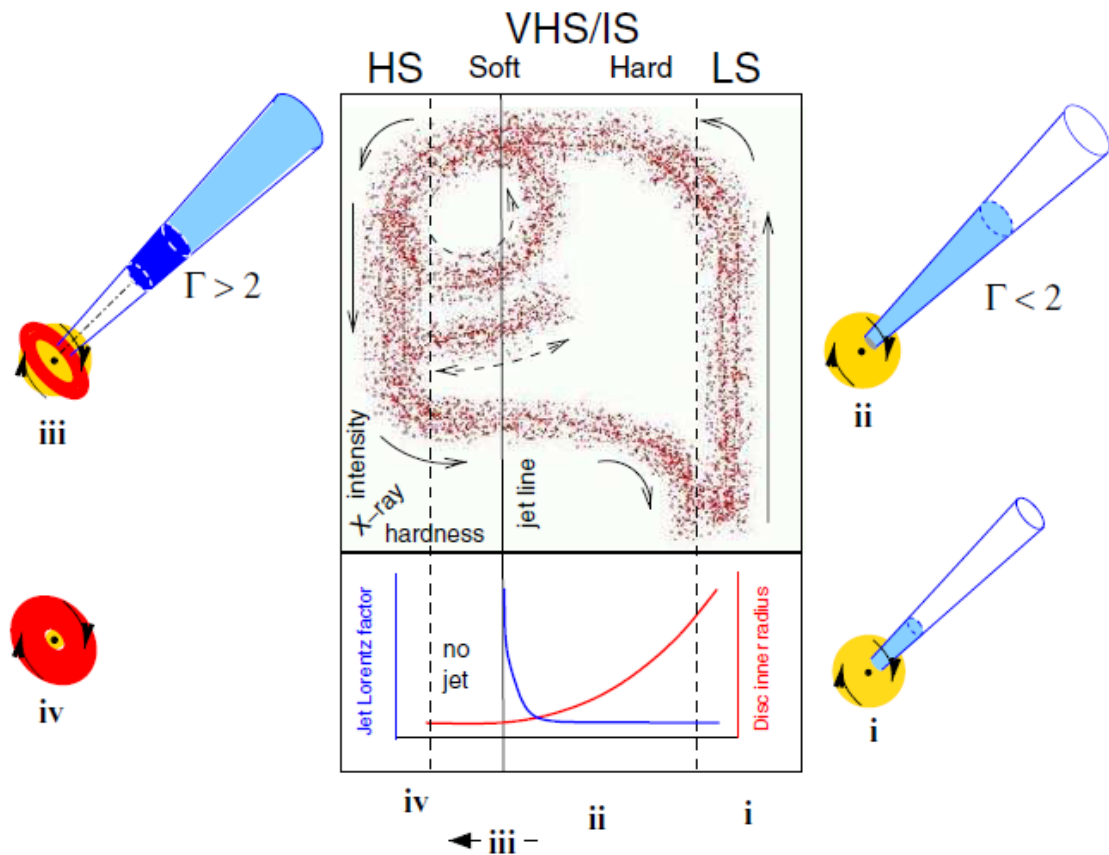
Figure 2.26: Schematic of the progress of a BHB on a hardness-intensity diagram during a transition from the low/hard to high/soft state and back again, where intensity increases vertically and hardness increases to the right, from Fender et al. (2004). The surrounding cartoons show the proposed behaviour of the jet during the transition: in the LHS there is a steady mildly relativistic jet (i and ii), as the system approaches the transition the jet BLF increases producing a travelling internal shock in the jet (iii) and in the HSS the jet switches off (iv) as the system crosses the jet line.
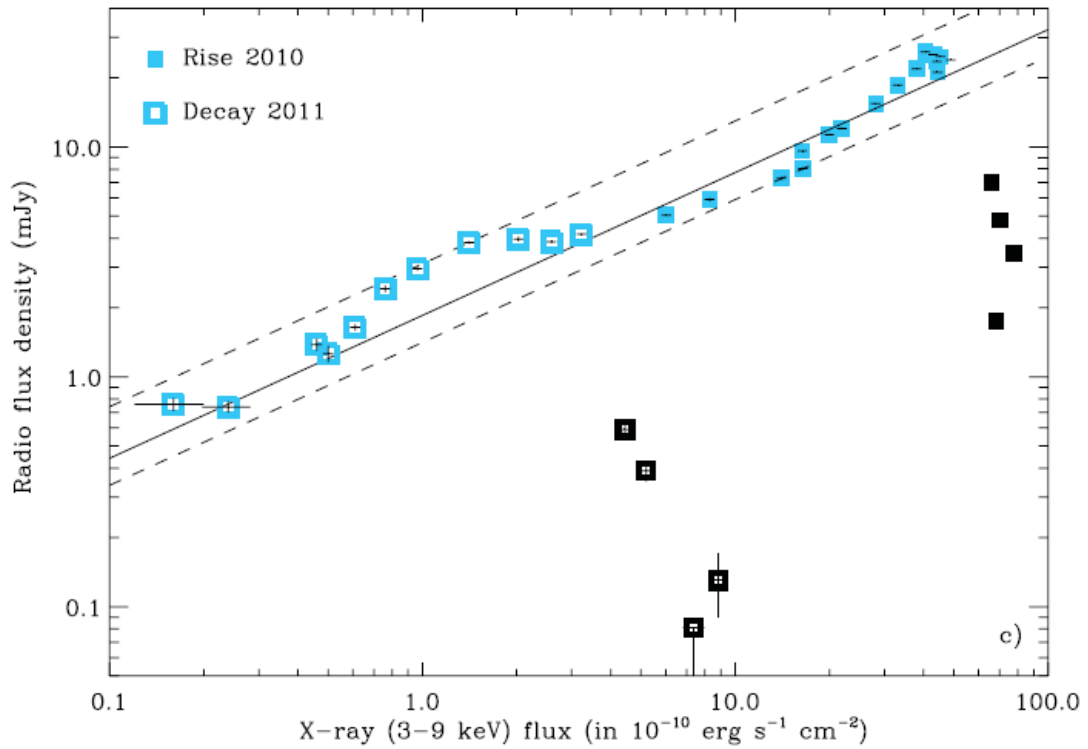
Figure 2.27: Radio–X-ray correlation in GX 339–4, showing the break in the correlation and suppression of radio emission at high X-ray luminosities when the source reaches the high/soft state and the switching back on of the jet when the source drops back down into the low/hard state at lower luminosity, from Corbel et al. (2013).

et al., 2013). This suppression of jet emission in the HSS suggests that whatever conditions are necessary for jet formation, they are present when the source is in the LHS and not when it switches to the HSS. The key difference between LHS and HSS accretion flows is that the LHS accretion flow has a much larger scale-height. Hydrodynamic simulations support the interpretation that a large scale-height flow is required to provide the large scale-height magnetic fields necessary for jet formation (Ohsuga & Mineshige, 2011).

Chapter 3

Introduction: Supermassive Black Holes

In contrast to stellar mass black holes, which are typically around $10\,M_\odot$, supermassive black holes have masses in the range of a million to a billion solar masses ($\sim 10^6 - 10^{10}\,M_\odot$). While stellar mass black holes are found dotted throughout galaxies, each galaxy generally has only one supermassive black hole at its centre. How these supermassize black holes formed and grew to such enormous sizes is still not well understood. Fig.3.1 shows several possible routes for supermassive black hole formation. One theory is that these black holes were formed from the collapse of the first generation of stars (Madau & Rees, 2001). A lack of heavy elements makes cooling inefficient causing gas clouds to collapse slowly and form massive stars $\gtrsim 250\,M_\odot$, which collapse to form supermassive black hole seeds of $\sim 100\,M_\odot$ (Bromm & Larson, 2004). Other theories suggest these first stars may merge to form a massive star which goes on to collapse into a black hole. These supermassive black hole seeds of $\sim 10^2 - 10^4\,M_\odot$ then grow through accretion to form the billion solar mass black holes in the present universe (Volonteri, 2012).

In the case of stellar mass black hole binaries (BHBs), the black hole accretes matter from its companion star, so this material has a well defined angular momentum. For a supermassive black hole, the gas supply for it to accrete comes from its host galaxy. As gas in the interstellar medium cools it forms stars. Any gas that does not go into star formation will sink towards the centre. If it loses enough grav-

Figure 3.1: Potential pathways for the formation of supermassive black hole seeds, from Volonteri (2012).

itational potential energy and angular momentum to come within the black hole's sphere of influence (i.e. where the gravitational field of the black hole dominates over that of the galaxy) it will be accreted.

A supermassive black hole may grow through either prolonged or chaotic accretion (Fanidakis et al., 2011; Fanidakis et al., 2012). In prolonged accretion the gas from a particular star formation episode is accreted in a single event. In chaotic accretion the available gas is not accreted in one go but split up into multiple smaller accretion events. If the total mass of in-falling gas is the same, the mass of the black hole will change by the same amount regardless of which accretion mode takes place. However the effect on the black hole's spin is very different if the accretion mode is prolonged compared to chaotic.

In the case of prolonged accretion, the gas is accreted in a single event and crucially from the same inflow direction. This results in net angular momentum transfer to the black hole. A prolonged accretion event will therefore quickly spin a supermassive black hole up to close to maximal spin (Volonteri et al., 2005; Volonteri et al., 2007). In contrast, in the case of chaotic accretion the gas is split into

multiple smaller accretion events, which enter from random directions with random angular momenta. This is possible because the scale-height of star formation is large compared to the black hole sphere of influence (King et al., 2008). As a result the black hole will be alternately spun up and spun down until the net angular momentum transfer is zero, resulting in a low spin black hole (Fanidakis et al., 2011; Fanidakis et al., 2012).

The energy released through accretion onto a central supermassive black hole can be so large that supermassive black hole accretion flows can outshine their host galaxies by several orders of magnitude. If the supermassive black hole at the centre of a galaxy is accreting — even at low accretion rates — then the galaxy is said to contain an Active Galactic Nucleus (AGN).

There is one final method by which supermassive black holes can grow, and that is through mergers. Galaxies grow through gas cooling onto them and through merging with other galaxies. When two galaxies merge so do their black holes. The black holes will sink to the centre of the merged galaxy. As they orbit one another they lose energy by emitting gravitational waves until eventually the two black holes coalesce (Merritt & Milosavljević, 2005). Just like prolonged accretion, this results in a final black hole that is highly spinning (Rezzolla et al., 2008).

That supermassive black holes grow in tandem with their host galaxies is evidenced by the $M - \sigma$ relation, shown in Fig.3.2. In other words, the stellar velocity dispersion (σ) of a galaxy bulge is found to correlate with the mass of the central black hole (Ferrarese & Merritt, 2000; Gebhardt et al., 2000). This means that the mass of the central black hole is closely linked to the size of its galaxy, with larger black holes residing in more massive galaxies. Clearly the more massive the galaxy the more gas and dust available to feed and grow the black hole. However the interaction is not just one way. Not only does the galaxy affect the growth of the black hole, but the growth of the black hole also affects the galaxy. The central black hole, with its sphere of influence of only a few parsecs, can affect the structure and growth of its host galaxy on megaparsec scales. This is because accretion is accompanied by ejection.

Radiation pressure can drive material off the accretion disc in the form of a wind.

Figure 3.2: $M - \sigma$ relation for supermassive black holes and their host galaxies, from McConnell & Ma (2013).

Such winds are wide angle and can reach velocities of $0.5\,c$ (Cappi, 2006; Tombesi et al., 2010). When the wind collides with material in the galaxy it dumps enormous amounts of energy into its surroundings, heating interstellar gas and suppressing star formation. Since only cool gas can sink to the centre and be accreted, this heating prevents further accretion. In this way the black hole 'feeds back' to its galaxy, tying together their growth (Silk & Rees, 1998). This is known as the radiative (or quasar) mode of feedback.

Jets provide an even more powerful method of feedback. Jets are highly collimated and can reach velocities $> 0.994\,c$. Consequently they can transport energy far beyond the black hole's host galaxy into the intracluster medium. There the energy may be dumped in giant radio lobes, reheating the intracluster medium and preventing the galaxy growing by cutting off its external gas supply. This is known as the kinetic (or jet) mode of feedback (Fabian, 2012).

Cosmological simulations have shown that such feedback mechanisms must be included in order to correctly reproduce the luminosity function of galaxies and its evolution with redshift (Bower et al., 2006). Hence understanding accretion around supermassive black holes and how they feed back and what conditions are required for them to feed back, as well as how much energy they return to their host galaxies and whether this energy comes from the in-falling material or the black hole itself, are of much wider importance.

3.1 Types of Active Galactic Nuclei

The ultra-violet (UV) and optical spectra of AGN almost always show strong emission lines and these have historically been used to classify AGN into different types. The larger masses of AGN compared to BHBs cause their disc emission to peak in the UV, rather than the X-rays. If the UV emission is strong it illuminates circumnuclear material, exciting the gas, which then remits the energy as optical and UV line emission. The strongest, most commonly observed lines are the Balmer series of Hydrogen, Hydrogen Lyα, MgII, CIV and OIII.

Fig.3.3 shows examples of optical spectra typical of seven main types of AGN.

Figure 3.3: Optical spectra from different types of active galactic nuclei, from http://www.astr.ua.edu/keel/agn/spectra.html.

Close inspection of the Seyfert 1, Broad Line Radio Galaxy (BLRG) and mean quasar spectra reveals that the emission lines have both a broad and narrow component, while the same lines in the Seyfert 2 and Narrow Line Radio Galaxy (NLRG) spectra show only a narrow component (e.g. Khachikian & Weedman, 1974). Low Ionisation Emission Region galaxies (LINERs) show predominantly low ionisation lines such as OI and NII (Heckman, 1980) and have much lower nuclear luminosities, while BL Lac spectra show either very weak lines or no lines at all (Padovani & Giommi, 1995).

Emission line widths are affected by the velocity of the material that emitted them. If the material is travelling more quickly, Doppler effects will cause the emission line to be broader, as photons from material travelling towards the observer will be blue shifted while photons emitted by material moving away will be red shifted. Gas closer to the central black hole will be orbiting faster, hence the broadest lines are emitted closest to the black hole, while narrower lines are emitted at much larger radii.

A given emission line will only be produced by gas close to the correct ionisation state for that transition. Ionisation state (ξ) depends on the density of the gas (n) and the illuminating flux, i.e. on the luminosity of the illuminating source (L) and the radius of the gas from it (R), as (Peterson, 1997):

$$\xi = \frac{L}{nR^2} \tag{3.1.1}$$

Hence, for a given luminosity, gas can be in the same ionisation state (and therefore emit the same line) at two different radii providing gas at a larger radius has a lower density. In this way, the same object can show both a broad emission line component produced closer to the black hole and a narrower component produced further out. Another consequence of the gas being denser at smaller radii and less dense at larger radii is that the broad lines are only due to permitted transitions, while the narrow lines include both permitted and forbidden transitions, since forbidden transitions are collisionally de-excited (and therefore suppressed) in the higher density central regions (Peterson, 1997).

Since $\xi \propto L$, this means the width of any given emission line will depend on the mass and accretion rate of the black hole, since $L \propto \dot{m}M$, where $\dot{m}$ is the Eddington scaled accretion rate ($\dot{m} = \dot{M}/\dot{M}_{Edd}$). The broad emission line component (for instance) will be produced by gas at some critical density and ionisation state. The radius at which this gas is located is then $R \propto (L/n_c\xi_c)^{1/2} \propto (\dot{m}M)^{1/2}$. In terms of gravitational radii (R_g), this becomes $r \propto R/R_g \propto (\dot{m}/M)^{1/2} \propto 1/v^2$, where v is the velocity of the gas assuming it's in a Keplerian orbit. The larger the black hole mass, the smaller the radius at which the line is produced, the faster the orbiting gas and hence the broader the line. Conversely, the higher the accretion rate, the larger the radius at which the line is produced, since the higher central luminosity over-ionises the gas out to a larger radius. Gas clouds at larger radii have a smaller orbital velocity and the line width is smaller. Broad lines can therefore be used in two ways to determine supermassive black hole masses: the time taken for the line to respond to an increase in the illuminating flux gives a measure of the light travel time to the broad line emitting gas which depends on the mass of the black hole — this is reverberation mapping (Peterson & Horne, 2004) — and the line widths

themselves can be used to estimate black hole mass (e.g. Corbett et al., 2003).

Variations in black hole mass and accretion rate can therefore explain why different classes of AGN show different broad line widths: Quasars and Seyfert 1s have similar accretion rates ($\dot{m} \sim 0.1$), but the black holes in quasars are typically an order of magnitude bigger ($M \sim 10^9$ as opposed to $10^8 \, M_\odot$), hence quasars show broader emission lines and have higher luminosity. A subset of Seyfert 1s, called Narrow Line Seyfert 1s (NLS1s), show both broad and narrow lines like all Seyfert 1s, however their broad lines are particularly narrow ($v < 2000 \, \text{cm s}^{-1}$, Osterbrock & Pogge, 1985). This is because NLS1s have smaller masses ($10^7 \, M_\odot$) and higher accretion rates ($\dot{m} \sim 1$) than ordinary Seyfert 1s. As a result they have the same central luminosity but narrower broad lines, since $v^2 \propto (M/\dot{m})^{1/2}$. LINERs show only low ionisation emission lines because their nuclear luminosities are too low to highly ionise the surrounding material — even though their masses are relatively high ($\sim 10^8 - 10^9 \, M_\odot$), this is not enough to compensate for their very low accretion rates ($\dot{m} < 0.01$) and hence radiatively inefficient flows.

However, differences in mass and accretion rate cannot explain why some AGN appear to lack broad emission lines altogether; Seyfert 2s have similar masses and accretion rates to Seyfert 1s, while NLRGs have similar masses and accretion rates to BLRGs. Yet Seyfert 2s and NLRGs show only narrow emission lines, with no broad component. A model developed to explain this is the unified model of AGN.

3.2   Unified Model of Active Galactic Nuclei

In theory, supermassive black holes (i.e. AGN) should simply be scaled up versions of stellar mass BHBs. Black holes and their accretion flows after all have only four parameters of mass, spin, accretion rate and inclination. However Fig.3.3 suggests there are many more types of AGN than there are variations in BHBs. This is partly due to the fact AGN span a much larger range in mass (~ 5 decades, while BHBs are all typically around $10 \, M_\odot$), and partly because supermassive black holes live in a much more complex environment, making inclination much more important.

Fig.3.4 shows the unified model of AGN, whereby the apparent presence or ab-

Figure 3.4: Unified model of active galactic nuclei, from Beckmann & Shrader (2012).

sence of broad emission lines can be explained by differences in viewing angle (Antonucci, 1993; Beckmann & Shrader, 2012). The maximum size of a BHB accretion disc is limited by the binary's orbital separation. In the case of AGN, the size of the accretion disc is limited by self-gravity. Beyond the self-gravity radius material is cool enough that it begins to clump and prevents an accretion disc forming. This limits the maximum radius of the accretion disc to $\sim 500 - 1000\, R_g$, depending on the black hole mass and accretion rate (Laor & Netzer, 1989). Far beyond this radius the temperature eventually falls below the dust sublimation temperature and according to the unified model a dusty torus forms (Antonucci & Miller, 1985). The torus has a large scale-height and is optically thick to X-rays. X-rays from the central regions are absorbed and re-emitted in the infra-red (IR). For large inclinations, where our line of sight intersects the torus, the central regions will be obscured. If the hydrogen column density $N_H > 1.2 \times 10^{24}\, \mathrm{cm}^{-2}$, such obscured AGN are termed 'Compton thick' (Maiolino et al., 1998).

Not only will the direct continuum emission from the accretion flow be obscured,

but so will the reprocessed line emission produced above the disc. The broad line emission is produced closer to the black hole than the narrow line emission, hence it is more likely to be obscured, while the narrow lines, which are produced further out, should always be visible.

For systems at low inclination $i \lesssim 30°$ (where $i = 0°$ is the axis of rotation of the accretion flow) we have a direct view down to the innermost parts of the accretion flow: the narrow line region (NLR), IR torus, broad line region (BLR) and direct accretion flow emission will all be observed. Consequently the AGN will be classed as a type I quasar or Seyfert 1 (depending on its luminosity) if it is radio quiet (i.e. has a weak radio jet), or as a BLRG or radio loud quasar if it has a strong radio jet.

As inclination increases, the torus gradually obscures the direct accretion flow and BLR emission. The scale-height of the torus is sufficiently large that, for $i \gtrsim 45°$, both the accretion flow and BLR are completely obscured and the only emission lines come from the NLR. As a result the object will be classed as a Seyfert 2, NLRG or type II quasar, depending on its luminosity and the strength of its jet. Differences in inclination can therefore explain the variation in emission line properties and the proportion of objects in each class.

For the case of radio loud AGN, low inclination has extra significance, since for $i \sim 0°$ our line of sight is orientated directly down the jet. The jet emission — which is already strong in a radio loud AGN — is maximally Doppler boosted due to its bulk motion and can dominate the spectrum of the AGN, in which case the AGN is classed as a blazar. In the case of flat spectrum radio quasars (FSRQs), broad and narrow emission lines are still seen, since FSRQs are high accretion rate systems with strong lines (Ghisellini et al., 1998; Ghisellini et al., 2010). In the case of BL Lacs, which are at much lower accretion rates, the jet emission swamps any weak emission lines from the accretion flow so they may appear lineless.

The unified model therefore makes it possible to separate out the effects of inclination from the more fundamental black hole parameters of mass, accretion rate and spin. By concentrating on unobscured objects — those which show both broad and narrow lines — it is possible to investigate whether the underlying AGN have spectral states like BHBs.

3.3 Spectral States

It is not possible to observe a single AGN progress through successive state changes as it is for BHBs. This is because all timescales scale with mass, so that 1 year for a $10^8\,M_\odot$ AGN corresponds to just $3\,\mathrm{s}$ for a $10\,M_\odot$ BHB. State changes in BHBs typically occur on a timescale of days. Correcting for mass and a smaller outer disc radius means the corresponding state changes in AGN should occur on megayear timescales, making state changes in individual objects impossible to observe. The only way to investigate the possible states of AGN accretion flows is therefore to examine the spectra of multiple objects with different masses and accretion rates, taking into account selection effects.

A further complication is that the accretion discs of AGN peak in the UV which is unobservable due to absorption in our own galaxy. This is because the disc temperature $T_d \propto L/A$ and $L \propto M$ while the emitting area $A \propto M^2$, because the inner disc radius scales with mass. This gives $T \propto 1/M$, shifting the peak luminosity of the disc from the X-rays for a BHB to the UV for an AGN. Nevertheless, by combining optical and X-ray data it is still possible to reconstruct the accretion flow spectra of AGN.

3.3.1 The Soft Excess

Fig.3.5 shows how the spectrum of an AGN changes with accretion rate. The three spectra are mean spectra obtained by binning Seyfert 1 spectra by accretion rate (Jin et al., 2012) and correcting for mass (Done et al., 2012), since low accretion rate Seyferts (i.e. BLS1s) generally have larger masses than higher accretion rate NLS1s. The three spectra correspond to a $10^8\,M_\odot$ AGN at $\dot{m} = 0.77$ (red), 0.25 (green) and 0.058 (blue).

Just like BHBs, AGN spectra can be fitted with a hard power law tail peaking at $\sim 100\,\mathrm{keV}$ and a black body (BB) disc component, peaking at 6×10^{-3}, 1.5×10^{-2} and $2.5 \times 10^{-2}\,\mathrm{keV}$ for the blue, green and red spectra, respectively. As accretion rate drops, the disc peak not only drops in luminosity but shifts to lower energy, just like in BHBs, suggesting a truncating disc. However it is clear that these two

Figure 3.5: AGN accretion flow spectra as a function of accretion rate for $\dot{m} = 0.77$ (red), 0.25 (green) and 0.058 (blue), for a $10^8\,M_\odot$ black hole. From Done et al. (2012).

components alone are not enough to fit the spectra. They cannot explain the excess of emission at soft X-ray energies ($\sim 0.01 - 1\,\text{keV}$), which is present in all three spectra but most prominent in the blue spectrum. This excess of emission above the extrapolation of the hard $2 - 10\,\text{keV}$ power law down to lower energies is referred to as the 'soft excess'.

The soft excess can be well fit by an additional Comptonisation component (Czerny et al., 2003; Gierliński & Done, 2004; Porquet et al., 2004). Although extreme relativistic reflection has been proposed as an alternative explanation, for the soft excess of NLS1s in particular (Crummy et al., 2006). BHBs do show evidence for an additional soft Comptonised component, but in bright intermediate states (Yamada et al., 2013), as the disc moves inwards with increasing accretion rate, dramatically increasing the seed photon flux as the system approaches the HSS. In contrast, in AGN, the soft Compton component is strongest at low accretion rates.

The disc, hard and soft Comptonised components can be understood in terms of a phenomenological model where the disc truncates as accretion rate drops and

the accretion energy below R_{trunc} is used to power the two Compton components (Done et al., 2012). Fits to multiple objects suggest the energy split between the two remains roughly constant, at 30% in the hard power law and 70% in the soft excess (Jin et al., 2012). The soft Comptonisation is optically thick and can be interpreted as coming from the inner regions of the disc below R_{trunc}, which for some reason do not completely thermalise so emit via optically thick Compton rather than standard BB emission. The seed photons for soft excess Comptonisation may well come from the disc. In contrast the seed photons for the optically thin, hard Comptonisation can be traced to the soft excess (Jin et al., 2013). This, combined with the fact the variability timescales increase from disc to soft excess to hard power law, support a picture where the outer disc is truncated and gives way to the optically thick region producing the soft excess emission, above which is an optically thin corona where soft excess seed photons are up-scattered into a hard power law. This scenario is shown in Fig.3.6.

There are two possible reasons AGN discs do not completely thermalise, while BHB discs appear to. One is that AGN discs are UV bright, due to being cooler, and UV line driving will drive material off the disc in the form of a wind, disrupting its structure regardless of whether the wind escapes or is failed (Risaliti & Elvis, 2010). Another possible reason is that AGN discs are more radiation pressure dominated than BHB discs, since the ratio of radiation pressure to gas pressure is $\propto (\dot{m}M)^{1/4}$ (Laor & Netzer, 1989). This leads to a larger scale-height and may affect the magneto-rotational instability (MRI), causing turbulent Comptonisation in the disc rather than thermalisation (Socrates et al., 2004).

By examining Fig.3.5 it is clear to see that BLS1s and their radio loud counterparts BLRGs should have accretion flow spectra similar to the blue spectrum. NLS1s should be similar to the red spectrum, but with the disc peak shifted to slightly higher energies due to their smaller mass. The standard quasar spectrum should be similar to the green line, but with the disc peak shifted to lower energies due to their large masses and comparison with the mean quasar spectrum of Elvis et al. (1994) shows that this is the case. Low accretion rate LINERs should be similar to the blue spectrum but with even larger disc truncation radii and lower lu-

Figure 3.6: Schematic showing AGN spectral components together with proposed accretion flow geometry, including emission from a truncated accretion disc (red), optically thick soft excess Comptonising region (green) and optically thin Comptonising corona (blue). Black line shows total spectrum, grey dashed line shows spectrum after including external absorption. From Collinson et al. (2015).

minosities. Following the unified model, the obscured and narrow line counterparts of all these objects should have the same underlying accretion flow spectra. The spectra we observe from them only appear different due to a combination of their intrinsic spectra and the limitations of our line of sight, through which absorption, scattering and reprocessing reduce and redistribute the emission we see.

3.3.2 Reflection

Reflection is seen in AGN just as it is in BHBs. Reflection from the torus produces a narrow neutral iron line and reflection continuum, while reflection from the accretion disc produces a broader iron line component (Reeves et al., 2001).

The broad component can generally be modelled by neutral to moderately ionised reflection from material within $50\,R_g$ with a standard r^{-3} illumination profile (Nandra et al., 2007). However, much more extreme parameters are needed if the reflection component is required to explain both the iron line and the soft excess emission. Such models are termed reflection dominated models.

3.3.2.1 Reflection Dominated Models

Reflection dominated models, or extreme relativistic reflection models, were developed to explain both the broad iron lines and strong soft excess emission of several high accretion rate NLS1s (MCG 6-30-15: Wilms et al., 2001, 1H0707–495: Fabian et al., 2002, e.t.c.). In these models the spectrum is decomposed into some BB disc emission at the very softest energies ($< 0.5\,\mathrm{keV}$), a hard coronal power law and a strong relativistic reflection component which provides both the iron line and all the excess emission at soft energies.

In order to produce enough reflected emission at soft energies requires extremely centrally concentrated illumination with r^{-5} or r^{-6}, often higher than solar abundances and reflection from $r < 3\,R_g$, in order to smear out the associated emission lines and produce the smooth soft spectrum that is observed (Fabian et al., 2004). Requiring reflection to occur from $r < 3\,R_g$ requires the disc to extend down to the last stable orbit and the central black hole to be highly spinning. Producing such centrally concentrated illumination requires the illuminating corona to be confined

Figure 3.7: Intensity-sorted X-ray spectra of the NLS1 MRK 766, showing systematic spectral changes with luminosity, from Miller et al. (2007).

to $\sim 2\,R_g$ on the spin axis above the black hole, leading to this scenario also being termed the 'lamppost model'.

Some support for the lamppost model comes from the fact that some NLS1s show spectra at low luminosities that look very like they are dominated by ionised reflection, while at high luminosities their spectra look much more power law like (Vaughan & Fabian, 2004). Fig.3.7 shows a sequence of such spectra from the NLS1 MRK 766. In the context of the lamppost model, this can be explained by the height of the illuminating corona changing. When the object is bright, the illuminating corona is higher above the black hole, more direct power law flux is seen and the reflected component is diluted. When the object is dim, the illuminating corona is closer to the black hole, more of the power law flux is bent back towards the disc due to gravitational light bending, less direct emission is observed and the spectrum is dominated by the reflected component (Miniutti & Fabian, 2004).

However, extreme relativistic reflection is not the only way to model these dim, apparently reflection dominated NLS1 spectra. A combination of moderate reflec-

tion, a continuum Compton component to produce the soft excess emission and clumpy absorption can also fit such spectra (Pounds et al., 2004). Some extreme NLS1s (e.g. MRK 766) do show absorption features around 6.7 keV, supporting the idea that absorption may play a role (Miller et al., 2007).

3.3.3 Reprocessing

As in BHBs, any flux incident on the accretion disc that is not reflected will be reprocessed. Again this heating effect on the inner edge of the disc must be taken into account when fitting disc spectra for inner disc radii (Done et al., 2013). However, since most AGN discs peak in the unobservable UV and only the smallest, highest mass accretion rate AGN show disc emission in the soft X-rays, this is not often attempted.

The presence of reprocessing at larger radii in the optical regions of the disc is evidenced by the fact the optical light curves of many AGN appear to lag the illuminating hard X-rays on short timescales (e.g. Arévalo et al., 2008). However, reprocessing should never dominate the outer disc emission in AGN as it does in BHBs. This is because reprocessing only dominates at the very largest radii $\sim 1000 - 2000\,R_g$, where the illuminating flux is greater than that produced by the disc (van Paradijs & McClintock, 1995). BHB discs can easily extend to these radii, depending on the binary separation. In contrast, AGN discs are much smaller, since they are truncated at the self-gravity radius, which limits their outer disc radius to $\sim 800\,R_g$. Consequently the intrinsic disc flux should be brighter than the illuminating flux at all radii in an AGN disc.

3.3.4 Comparison with BHB Spectral States

It is useful to note that the accretion flow structure of AGN at a given accretion rate does not necessarily match that predicted by scaling up a BHB. LINERs, with their discs truncated at large radii and their jets, correspond well to LHS BHBs. However BLS1s, BLRGs and quasars should correspond to the disc dominated HSS. Yet they still show significant disc truncation and a large soft excess despite accretion rates

comparable to disc dominated HSS BHBs, as evidenced by a hard coronal power law (e.g. ARK 120: Matt et al., 2014; Haardt & Maraschi, 1991) and disc peaks at low energies (e.g. MRK 509: Mehdipour et al., 2011). In addition, radio loud quasars show powerful jets, while the HSS of BHBs is exclusively radio quiet. Finally the Eddington rate NLS1s should correspond to the VHS of BHBs, where the jet switches back on and the hard power law starts to carry more of the energy budget again. Yet NLS1s are resolutely disc dominated, with a weak power law tail and the majority are radio quiet. Perhaps mass loss in UV winds and a tendency towards larger scale-height discs mean a higher accretion rate is required in an AGN to obtain the same geometry as a BHB, i.e. thermalised disc dominated states are harder to achieve so require higher mass inflow densities and therefore higher accretion rates.

3.4 Winds from Active Galactic Nuclei

Although the unified model can explain the presence/absence of broad lines in different objects through the visibility/obscuration of different regions of circum-nuclear gas, it doesn't explain how the gas producing the emission got there to begin with. Disc winds provide an obvious mechanism for supplying circum-nuclear material.

There are four main mechanisms by which material can be driven off the accretion disc — in a thermal wind, a dust driven wind, UV line driving and through radiation pressure. The presence or absence of each type of wind in a particular AGN depends on the state of its accretion flow and the spectral shape of its emission.

The central X-ray source illuminates the surrounding accretion disc and torus, heating the material to the Compton temperature ($\sim 10^7$ K). At large radii, material heated to the Compton temperature will be unbound since the thermal velocity of the material becomes greater than the escape velocity. This material is then driven from the system in the form of a thermal wind (Begelman et al., 1983). Thermal winds only exist at large radii so are associated with the torus and BLR.

The inner edge of the torus is set by the dust sublimation radius. Inside this radius the illuminating X-ray flux is too high, heating the material above the dust sublimation temperature so dust can no longer exist. However, material inside the

accretion disc is shielded from the central X-ray source by the rest of the disc, so dust can exist (Czerny & Hryniewicz, 2011). Radiation pressure from photons emitted by the accretion disc itself pushes on this dust lifting it up out of the disc. Once out of the disc it is illuminated by the central X-rays, the dust sublimates and either the material falls back down (failed wind) or if it has gained sufficient height its rotational velocity will be larger than the Keplerian velocity for its new larger radius and it will escape from the system. Czerny et al. (2015) propose that this dust driven wind region could provide the source of the broad line emission and this scenario is shown in Fig.3.8. The inner edge of this dust driving region is set by the radius at which the temperature within the disc reaches the dust sublimation temperature and the disc no longer contains dust grains.

AGN with strong UV disc emission can also power UV line driven disc winds. This requires moderately ionised gas, allowing for multiple UV line transitions (Castor et al., 1975). As the gas absorbs UV photons it also absorbs their momentum. This accelerates the gas to a higher velocity. Its UV line transitions are then Doppler shifted allowing it to absorb more UV photons of higher energy, accelerating it again. The process repeats, potentially accelerating the gas up to velocities of $> 0.1\,\mathrm{c}$ (Proga & Kallman, 2004). The range of radii over which UV line driving occurs is restricted to the range of disc radii emitting in the UV (Shlosman et al., 1985; Risaliti & Elvis, 2010). UV line driven disc winds are also restricted to X-ray weak sources, since if the X-ray illumination is strong it over-ionises any material that lifts off the disc so it no longer interacts with UV photons, effectively switching off UV line driving. If the wind self-shields, this can help (Murray & Chiang, 1996; Proga & Kallman, 2002, but see also Higginbottom et al., 2014), but primarily a UV strong, X-ray weak accretion flow spectrum is required for UV line driving to be effective.

The final mechanism for driving material off the disc is radiation pressure in the form of an Eddington wind. In the highest accretion rate sources, the luminosity of the central regions of the accretion flow is sufficiently high that radiation pressure alone will blow material out of the disc in the form of an Eddington wind (Shakura & Sunyaev, 1973; King & Pounds, 2003). Since the luminosity is highest in the innermost regions, Eddington winds are generally confined to small radii.

Figure 3.8: Schematic showing the formation of the torus and broad line region. The torus exists beyond the dust sublimation radius. Below the dust sublimation radius dust can only exist in the disc, where it is shielded from the central X-ray source. Radiation pressure lifts this dust up out of the disc, where irradiation increases, the dust evaporates and the remaining material falls back down. The whole dust driving region is then the source of the broad lines, with the inner edge set by the radius at which dust within the disc sublimates. From Czerny et al. (2015).

Fig.3.9 summarises which of these different winds may be present in AGN of different masses and accretion rates.

High mass, high accretion rate AGN should show all four types of wind — thermal, dust driven, UV line driven and Eddington — since their disc spectra peak in the UV. Low mass, high accretion rate AGN should lack a UV line driven disc wind since their lower masses means hotter accretion discs peaking in the extreme UV (EUV) and soft X-rays rather than the UV (Hagino et al., 2015).

Moderate accretion rate AGN will lack a radiation pressure Eddington wind due to their lower accretion rates and lack a UV line driven wind since their spectra are UV weaker and X-ray stronger, over-ionising any UV wind that tries to form.

The lowest accretion rate AGN are so under-luminous they may be unable to drive a thermal wind, as the heating rate is less than the cooling rate of the rising and expanding material (Begelman et al., 1983), and they may even lack a torus due to the low densities of inflowing material (Gonzalez-Martin et al., 2015). Their spectra are UV weak and X-ray strong, due to the disc being replaced by a hot, inefficient, ADAF-like flow, such that even dust driving may be inefficient.

One explanation for the source of the narrow lines is that they are produced at the back of the torus, since within the torus, dust will absorb the line emission so that narrow line emission will only escape from the back, where the dust grains are clumping and the density is lower (Netzer & Laor, 1993; Netzer, 2015). This spatial separation of the broad and narrow line regions by the dusty absorbing torus then helps to explain why there are distinct broad and narrow line components rather than a continuum of line widths. However observations of extended narrow line regions with lower gas velocities, suggest some narrow line emission must also be produced by gas at much larger radii above the accretion flow, as in the traditional unified model scenario (Unger et al., 1987). This emission arises from illuminated interstellar gas in the host galaxy itself, rather than gas expelled from the accretion flow.

Figure 3.9: Schematic showing accretion flow geometry and spectral shape for AGN of different masses and accretion rates and the potential winds that may result.

Figure 3.10: Comparison of two AGN power spectra with power spectra of Cyg X-1 in the low/hard and high/soft states, from McHardy et al. (2004).

3.5   Variability

AGN show stochastic X-ray variability just like BHBs. They also follow a linear RMS–flux relation implying, just like BHBs, this variability is due to accumulation of mass accretion rate fluctuations generated at the viscous frequency at each radius (Uttley & McHardy, 2001). However since f_{visc} scales inversely with mass, the timescale of fluctuations is correspondingly longer; while BHBs show variability on millisecond timescales, the fastest variability in AGN is of the order of minutes.

3.5.1   Power Spectra

Due to their longer timescales, observations spanning years/decades are required to measure AGN power spectra. Fig.3.10 shows the power spectra of two AGN (NGC 3516 and NGC 4051) compared to the power spectra of the BHB Cyg X-1 in its hard and soft state.

The black hole in NGC 4051 is roughly $10^6\,M_\odot$, while Cyg X-1 is $\sim 10\,M_\odot$.

Scaling the 10 Hz high frequency break of the HSS Cyg X-1 power spectrum down by the five orders of magnitude mass difference implies the high frequency break of NGC 4051 should be at 10^{-4} Hz. Fig.3.10 shows that the high frequency break of NGC 4051 is in fact an order of magnitude higher than expected ($\sim 10^{-3}$ Hz). McHardy et al. (2006) showed this can be explained if the high frequency breaks of AGN scale with accretion rate as well as mass, so that $f_h \propto 1/T_h \propto M^{-1.12}\dot{m}^{0.98}$. This suggests there is some fundamental difference in the variability properties of AGN compared with BHBs, since most HSS BHBs show a high frequency break at 10 Hz regardless of accretion rate, expect for some states of GRS 1915+105 (at Eddington) which actually break at a lower frequency rather than higher (Zdziarski et al., 2005).

Another difference is that AGN do not seem to show low frequency breaks in their power spectra (one exception is the NLS1 galaxy ARK 564 for which a low frequency break has been measured at 10^{-6} Hz, McHardy et al., 2007). For very large mass, low accretion rate AGN, measuring the low frequency break would require a prohibitively long light curve. However for many lower mass AGN, where a low frequency break could be observed, it is not; instead below the high frequency break the power spectrum shows flat topped noise without any detection of a low frequency break (Uttley & McHardy, 2005). This is similar in shape to the HSS power spectrum of Cyg X-1 (which is itself not typical of other HSS BHB power spectra, Done & Gierliński, 2005). The absence of a low frequency break suggests the hard X-ray emission is responding to fluctuations accumulated from all radii. In the HSS of Cyg X-1 this can be explained by the hard fluctuations coming from a corona extending above the disc, in contrast to LHS BHBs where the strongest fluctuations are produced in the hot flow and so restricted to the range of hot flow radii. Perhaps AGN lacking low frequency breaks have similarly extended coronae.

Not only does the high frequency break change with mass and accretion rate scaling from BHBs to AGN, but so can the total power. While many AGN show flat topped hard band noise at $fP(f) \sim 0.01$ similar to BHBs (e.g. NGC 3516: Uttley & McHardy, 2005, PG1244+026: Jin et al., 2013), some show much more power, with $fP(f)$ nearing 0.1 (e.g. NGC 4051: Vaughan et al., 2011, 1H0707–495: Zoghbi

et al., 2011). Winds and obscuration provide an additional source of variability in AGN. Material thrown off the disc in a wind and even BLR clouds may interrupt the line of sight causing absorption events which alter the spectrum and introduce an extra source of variability. Such events must be excluded when investigating the intrinsic variability of the accretion flow. Since they generally occur in high inclination systems, this effect can be minimised by studying broad line systems at low inclination.

3.5.2 Time Lags

Just like BHBs, AGN can show time lags between hard and soft X-ray bands. As in BHBs, positive lags indicate propagation while negative lags (soft lagging hard) indicate reverberation. Fig.3.11 shows lag-frequency spectra for a sample of AGN, some of which show both positive and negative lags.

The discovery of short $\sim 30\,\text{s}$ high frequency soft lags in several NLS1s (e.g. Fabian et al., 2009; Emmanoulopoulos et al., 2011) lent support to the suggestion that the excess of soft emission in these objects was the result of relativistic reflection of the hard coronal power law off the accretion disc. Such reflection dominated models require reflection from disc radii extending down to $\sim 1\,\text{Rg}$ around a highly spinning central black hole. At first glance, the very short $30\,\text{s}$ reverberation lag is in good agreement with the idea the distance from hard illuminating source to reflecting disc is $< 5\,R_g$, since this roughly matches the light travel time assuming a $10^6\,M_\odot$ black hole. However it is important to note that the lag measured between two bands is strongly affected by dilution (Uttley et al., 2014; Kara et al., 2013). If the hard band contained only coronal power law emission and the soft only reflected emission, then the lag between hard and soft bands would be the lag between the coronal power law and its reflection. However, in practice, the hard and soft band will contain a contribution from both components and this reduces the measured lag from its 'true' value. Therefore if the lag between hard and soft bands is $\sim 30\,\text{s}$ and the hard contains primarily power law emission while the soft mostly reflected emission, then it is more accurate to say that the lag between the coronal power law and its reflection is at least $30\,\text{s}$ or longer. This is even more true if the soft
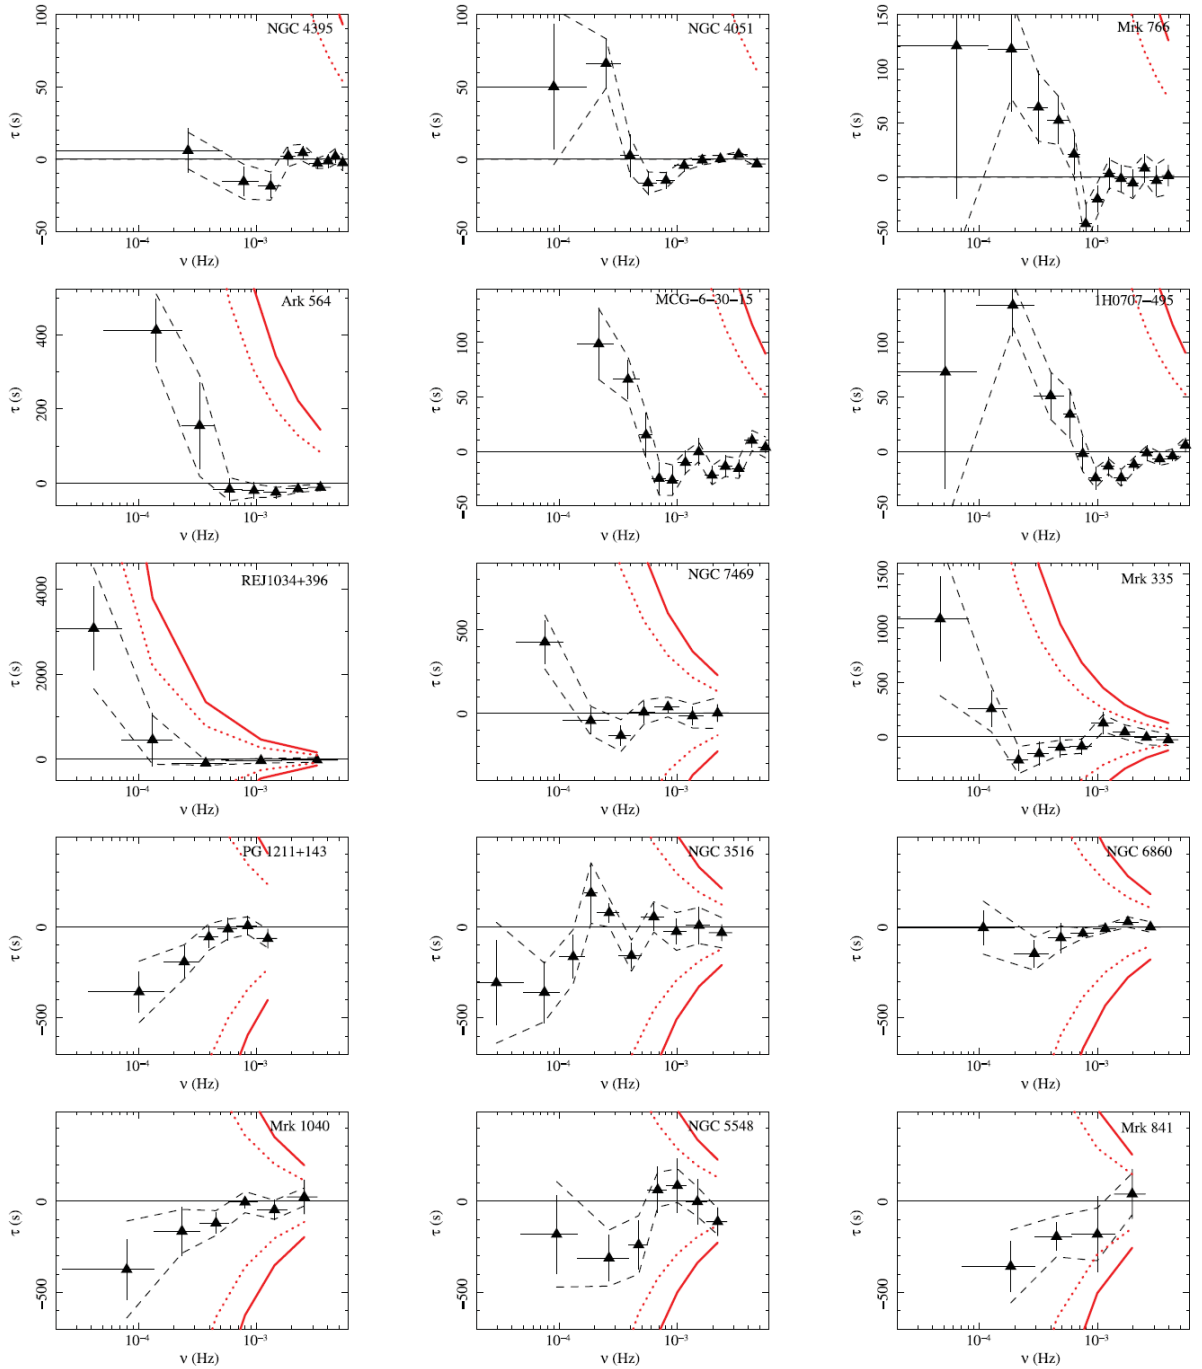
Figure 3.11: Lag-frequency spectra for a sample of AGN, where positive lags indicate the hard X-ray band lagging the soft X-ray band, from De Marco et al. (2013). Dashed black lines bound the 1σ confidence levels for the observed lags. Red solid lines show the range of allowed lag values at each frequency ($\tau = -\nu/2, +\nu/2$). Red dotted lines show the standard deviation of a uniform distribution defined on the same interval of lag time permitted values — Poisson noise errors become important when the spread on the measured lags (defined by the 1σ confidence levels) approaches this line.

Figure 3.12: Frequency and amplitude of the negative reverberation lag as a function of black hole mass for a sample of AGN, from De Marco et al. (2013). Coloured lines (from bottom to top) in the right hand panel show the light crossing time at $1\,R_g$, $2\,R_g$ and $6\,R_g$ as a function of black hole mass.

band contains an additional component which leads the hard power law, in which case these positive leads must be subtracted before estimating the true value of the negative lag.

De Marco et al. (2013) showed that both the frequency of the negative reverberation lag and its amplitude correlate with the mass of the central black hole (Fig.3.12). All size-scales scale with black hole mass, so light travel times should be longer in larger mass AGN (assuming reflection occurs at the same gravitational radius) and fluctuations occur at lower frequencies, so this is expected. The dashed lines in the right panel of Fig.3.12 show the light crossing time as a function of mass at different radii. However it is not appropriate to simply interpret the lag amplitude as a light travel time for any given AGN mass, since it is a combination of the actual light travel time to the reverberating component with that component's fractional contribution to the X-ray bands used to calculate the lag, plus any positive lags between the two bands at those frequencies. The fact that the lag amplitudes show a different dependence on mass to the light travel time at a constant radius, suggests either reverberation does not always occur at the same gravitational radius, or the fractional contribution of the reverberating component to the chosen X-ray bands is changing, or the strength of the positive lags is changing, or all three. Similarly the frequency that the minimum of the negative reverberation lag occurs at will depend on how much it is diluted by positive propagation lags, i.e. how strong the

positive lags are and how much of the soft band flux comes from the reverberating component and how much comes from soft continuum components.

3.6   Jets

AGN are historically divided into two classes: radio loud and radio quiet, where the radio loudness parameter (R) is defined as the ratio of the radio luminosity at $5\,\mathrm{GHz}$ to the optical luminosity at $4400\,\mathrm{\mathring{A}}$ and $R = L_R/L_{opt} > 10$ is classed as radio loud (Kellermann et al., 1989). Since the radio emission comes from the jet and the optical from the accretion flow, radio loudness is often used as a proxy for the strength of the jet. However, the amount of optical emission depends on the accretion state of the flow, which in turn depends on accretion rate, so radio loudness is affected by accretion rate. Low luminosity sources such as LINERs, with radiatively inefficient accretion flows and steady jets, will have a higher ratio of radio–optical luminosity, and so be more radio loud than, for example, a Seyfert with a bright radiatively efficient UV-optical accretion disc. Fig.3.13 shows that radio loudness does indeed decrease with increasing accretion rate. Similarly BHBs show a radio loud LHS at low accretion rates, with strong steady jet emission and a radiatively inefficient flow, and switch to a radio quiet HSS at high accretion rates, when the accretion flow emission is bright and efficient and the jet emission quenched. However it is important to note that, while high accretion rate AGN are less radio loud due to their more efficient accretion flow emission, they do not seem to show the additional jet quenching that HSS BHBs do. High accretion rate AGN are capable of sustaining steady radio jets while HSS BHBs do not. This may be additional support for larger scale-height flows in AGN, if a large scale-height flow is required to support a jet (Ohsuga & Mineshige, 2011).

Changing accretion state can explain the general trend of decreasing radio loudness with increasing accretion rate. However the top panel of Fig.3.13 shows that, despite separating AGN by accretion rate, there is still a large scatter in radio loudness (~ 5 orders of magnitude) for any given accretion rate. Some of this can be explained by differences in mass and some by differences in environment.
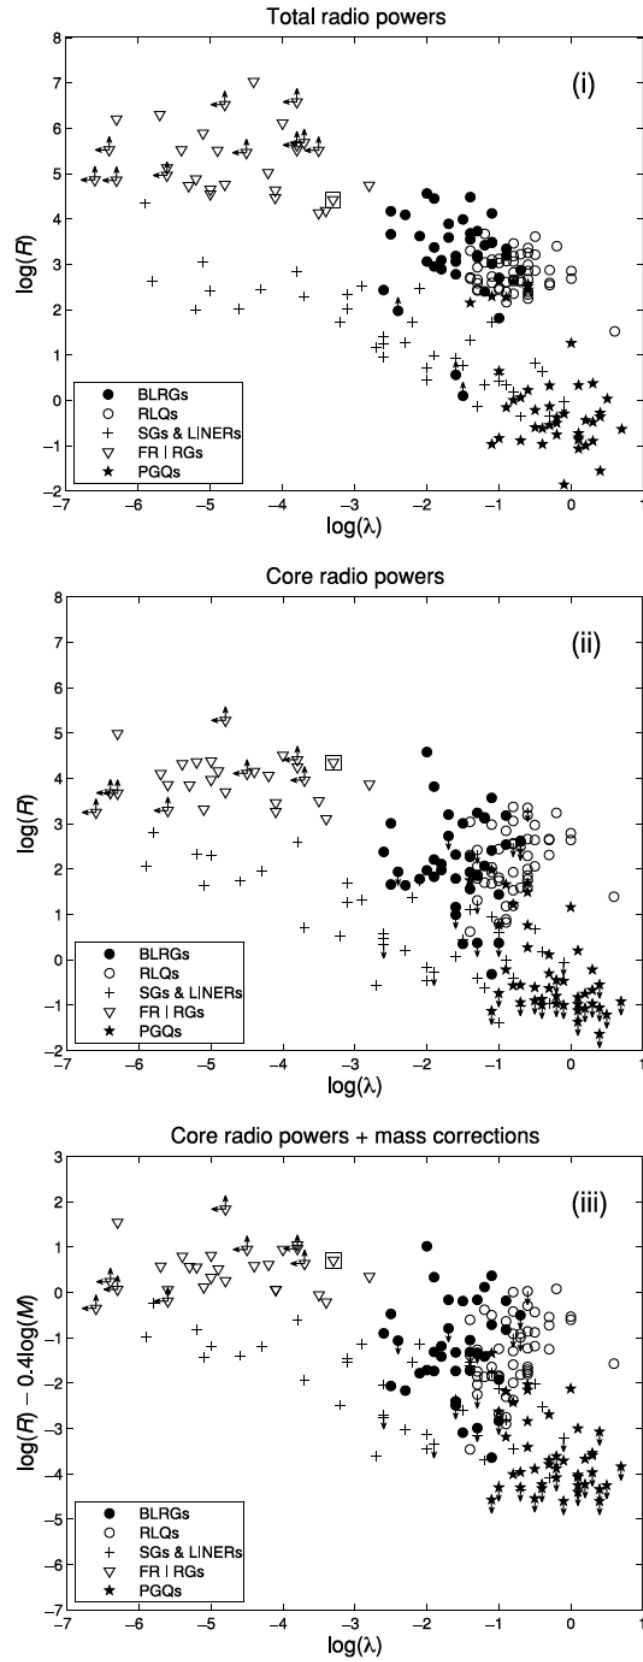
Figure 3.13: Radio loudness (where $R = L_{5\,\text{GHz}}/L_{4400\,\text{Å}}$) as a function of Eddington ratio calculated using i). total radio luminosity, ii). only core radio luminosity and iii). core radio luminosity plus a correction for AGN mass, from Broderick & Fender (2011).
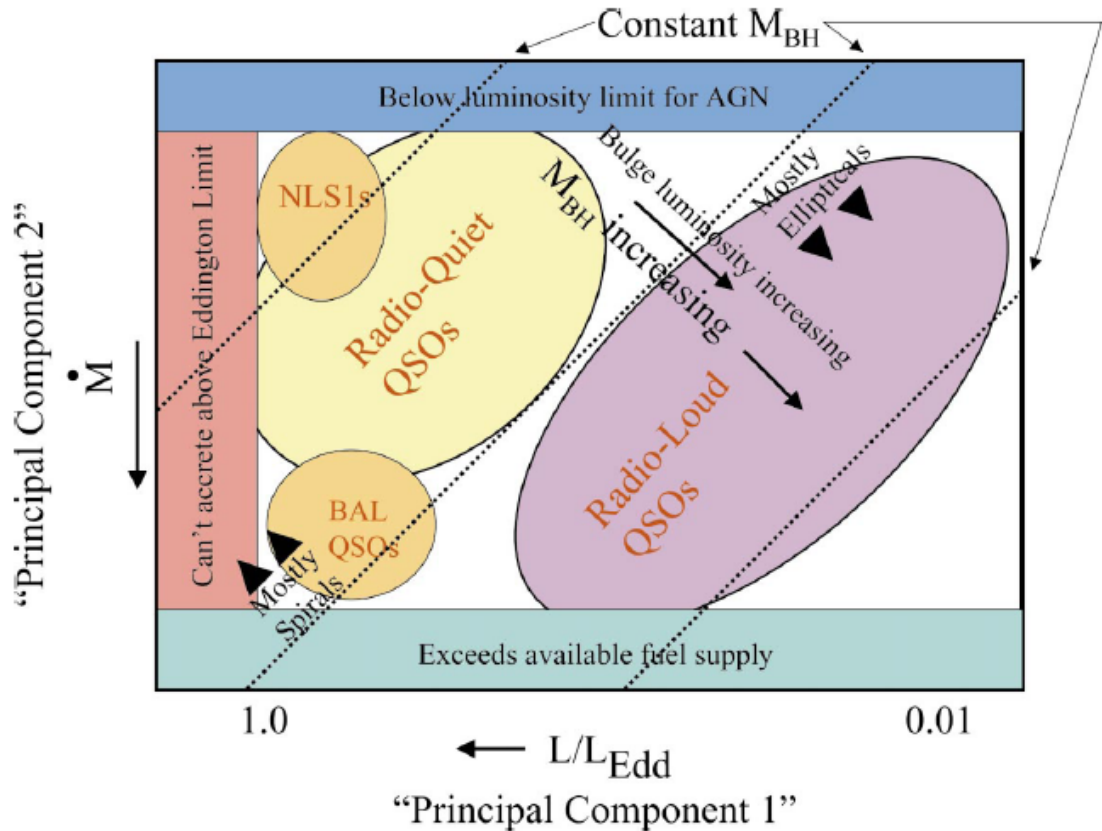
Figure 3.14: Different AGN groupings as a function of luminosity and accretion rate, from Boroson (2002).

More massive AGN have larger jets which are less self-absorbed so more radiatively efficient at a given frequency for a given accretion rate. A mass correction term must therefore be applied when comparing radio emission from different AGN (Heinz & Sunyaev, 2003). In particular, Seyferts can be two orders of magnitude smaller than radio loud quasars, which have typical masses of $10^8 - 10^9 M_\odot$, so their radio luminosity should be roughly three orders of magnitude lower ($L_R \propto L_X^{0.6} M^{0.8} \propto \dot{m}^{0.6} M^{1.4}$, Broderick & Fender, 2011).

Fig.3.14 shows how radio loud and radio quiet AGN are grouped by luminosity and mass and related to host galaxy size and type. Larger mass black holes tend to reside in more massive elliptical galaxies in denser environments. The emerging jet ploughs into the surrounding gas and dust often producing extended radio emission or evacuating bright radio lobes in the intracluster medium. In contrast, small black holes generally reside in small spiral galaxies which occur in less dense environments,

where there is much less surrounding material for the jet to interact with (Donoso et al., 2010; Hardcastle & Krause, 2013).

One way to reduce the effects of environment is to calculate radio loudness using the core radio emission and exclude extended emission produced by the jet interacting with the external environment. The bottom panel of Fig.3.13 shows that, by correcting for mass differences and using only core radio emission, the scatter can be reduced to ~ 3 orders of magnitude.

Nevertheless, at low accretion rates, Fanaroff-Riley Type I (FRI) radio galaxies remain systematically more radio loud than LINERs. While at high accretion rates, BLRGs and radio loud quasars remain systematically more radio loud than Seyferts and traditionally radio quiet Palomar-Green (PG) quasars. Having accounted for differences in mass, accretion rate and environment, leaves only one remaining parameter — spin. One explanation may be that jets are predominantly accretion powered, but if the black hole is highly spinning the jet may gain extra power by extracting some of the rotational energy of the black hole through the Blandford-Znajek mechanism (Blandford & Znajek, 1977). A difference in black hole spins from $a = 0.2 - 1$ could then explain the spread in radio loudness for a given accretion rate (Heinz & Sunyaev, 2003), since spin powered jets are expected to scale as a^2. An alternative explanation is that jet production relies not on black hole spin but on accumulation of magnetic flux from the galaxy, so that jet strength, and therefore radio loudness, depends on the amount of available flux and the ability of the black hole to accrete it (Sikora & Begelman, 2013), however such models have limited predictive power.

3.6.1 Fundamental Plane

LHS BHBs follow the radio–X-ray correlation, where the radio emission from the steady jet is correlated with the X-ray emission (generally assumed to be) from the accretion flow. With the addition of a mass correction term, this correlation can be extended to include AGN. This is known as the fundamental plane of black hole activity (Merloni et al., 2003), shown in Fig.3.15.

Figure 3.15: The fundamental plane of black hole activity, from Merloni et al. (2003).

Merloni et al. (2003) showed that a sample of LINERs, Seyferts and PG Quasars all lie on the fundamental plane. These are the AGN that lie on the lower half of the trend shown in Fig.3.13. In other words, the radio emission in these AGN is as expected from scaling up the steady jets of LHS BHBs, suggesting the mechanism of jet production in these objects is scale invariant. If spin is the source of the remaining scatter in Fig.3.13, it also suggests that the sub-sample of LINERs, Seyferts and PG quasars must have similar spin to the BHBs, since they follow the same correlation. If their spins differed they would be offset and require an additional correction term, besides just mass.

Fig.3.15 shows that the radio loud quasar 3C 273 is clearly offset from the fundamental plane. If the BHBs, LINERs, Seyferts and radio quiet PG quasars all have low to moderate spins and accretion powered jets, this is consistent with 3C 273 (and the other radio loud quasars, BLRGs and FRI objects on the upper half of the trend in Fig.3.13) potentially having higher spins and an additional spin contribution to
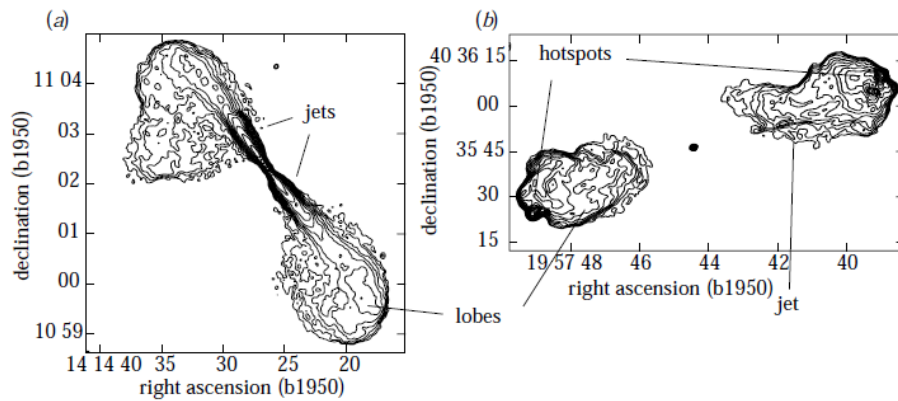
Figure 3.16: a). Fanaroff-Riley Type I radio galaxy 3C 296 (Leahy & Perley, 1991) and b). Fanaroff-Riley Type II radio galaxy Cygnus A (Perley et al., 1984), from Hardcastle (2005).

their jet powers.

3.6.2 Highly Relativistic

The FRI sources, BLRGs and radio loud quasars that lie in the upper half of the trend shown in Fig.3.13 represent some of the most powerful astrophysical jets. The most massive radio loud AGN, which live in the densest environments, are separated into two types according to their extended radio emission — Fanaroff-Riley Type I (FRI) and Fanaroff-Riley Type II (FRII) (Fanaroff & Riley, 1974), shown in Fig.3.16. FRI sources show fluffy radio jets that are brightest in the centre and fade away at large radii, in the form of wide angle plumes. The jets of FRII sources are much narrower and fainter and end in bright hot spots. For some reason FRI jets are inefficient, losing a large fraction of their energy as they travel, and so they fizzle out in plumes, while FRII jets radiate very little energy and remain well collimated, so that when they collide with external gas clouds all the energy is dumped in a bow shock and radiated away in giant radio lobes. One explanation for this may be a difference in jet opening angle. Simulations suggest jets with wider opening angles suffer more from entrainment and deceleration due to radiative loses (Krause et al., 2012). In which case, FRI sources should have wider angle jets. In addition FRI jets originate from low luminosity, low accretion rate systems ($\dot{m} < 0.01$), while FRII

jets originate from high luminosity, high accretion rate radio galaxies/radio loud quasars (see Fig.3.13). The higher power, and potentially greater mass loading, in higher accretion rate FRIIs may help them power through surrounding material and remain collimated. While the presence of an accretion disc down to smaller radii may also help to confine the jet and produce a smaller opening angle in FRIIs, compared to the large scale-height, inefficient accretion flows in FRIs (where the disc is truncated at very large radii or absent), if jet production is restricted to regions of the accretion flow with large scale-height.

Only sources viewed from the side ($i \sim 90°$) will be classed as FRI/FRII sources. As viewing angle decreases the receding jet becomes increasingly deboosted until only radio emission from the approaching jet is visible. If the jet is viewed head on ($i \sim 0°$) the source is classed as a blazar and Doppler boosted emission from the approaching jet dominates the spectrum of the AGN.

Blazars can be classed as either BL Lacs or FSRQs, with BL Lacs being FRIs viewed face on and FSRQs FRIIs viewed face on (following the unified model in Fig.3.4). Population studies taking into account beaming effects confirm that the observed numbers of BL Lacs/FSRQs are consistent with FRIs/FRIIs, respectively, being the parent populations (Urry & Padovani, 1995).

Fig.3.17 (left panel) shows the mean broadband spectrum from a sample of BL Lacs compared to the mean spectrum from a sample of FSRQs. In each case the plot models only the emission from the base of the jet, which dominates the spectrum, and neglects regions further out which sum to give the flat radio spectrum. The different spectral shapes of BL Lacs and FSRQs are a direct consequence of their different accretion regimes (Fig.3.17, right panel). Both BL Lacs and FSRQs show a hump at low energies due to synchrotron radiation from relativistic electrons in the jet, presumably accelerated by shocks at the base of the jet. FSRQs are in the radiatively efficient accretion regime ($\dot{m} > 0.01$) so they show a second smaller peak at $\sim 10^{15}$ Hz, which is BB emission from the UV bright accretion disc. BL Lacs have lower accretion rates ($\dot{m} < 0.01$) and radiatively inefficient flows so show no UV accretion disc and the jet emission drowns out any emission from the low luminosity flow. The higher energy hump ($\sim 10^{24}$ Hz) in the mean BL Lac spectrum is due to
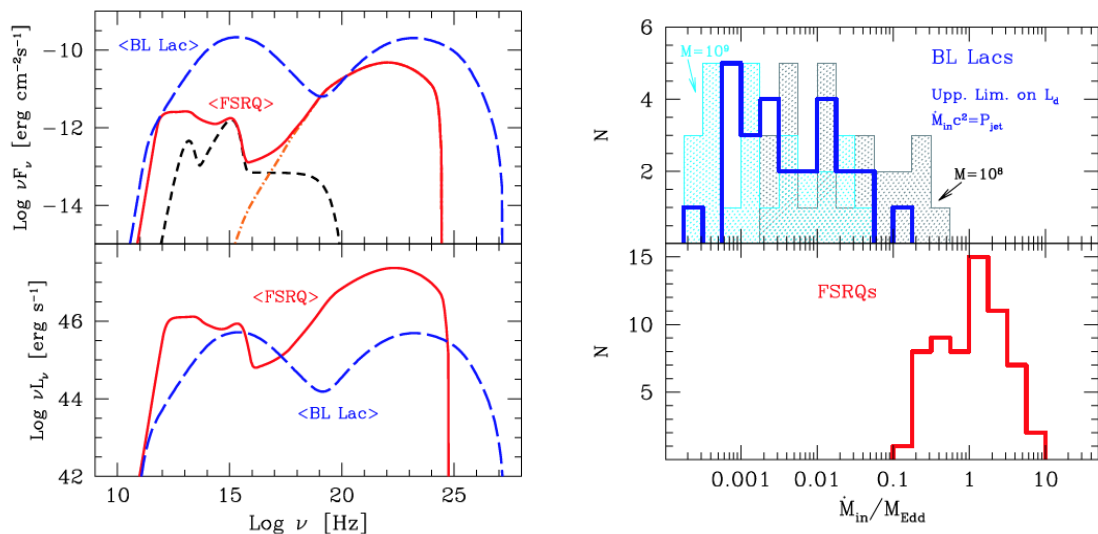
Figure 3.17: Mean BL Lac and FSRQ spectra (left panel) together with the number of BL Lacs and FSRQs as a function of accretion rate (right panel), from Ghisellini et al. (2010).

Compton up-scattering of the synchrotron radiation produced by the electrons and magnetic fields in the jet. The entire BL Lac spectrum is therefore composed of synchrotron and synchrotron self-Compton (SSC) emission.

The highest energy hump ($\sim 10^{22}$ Hz) in the FSRQ spectrum likewise contains a contribution from SSC emission. However the higher accretion rate FSRQs have an additional source of seed photons for Compton up-scattering in the jet. The UV bright accretion disc provides some seed photons, although these are behind the jet so strongly deboosted in the jet frame. Importantly the UV disc illuminates the BLR which (in the context of the traditional unified model) is *in front* of the jet emission region. The jet electrons are travelling towards this reprocessed line emission at relativistic velocities due to the bulk motion of the jet. Consequently they see this external source of seed photons strongly Doppler boosted so that external-Compton (EC) emission dominates over SSC. The higher jet power and greater seed photon flux means the Compton humps of FSRQs are more luminous than those of BL Lacs, and more seed photons means more cooling, so they also peak at slightly lower energies. The combination of Compton up-scattering and Doppler boosting makes BL Lacs and FSRQs some of the brightest objects detected at gamma-ray

energies (Chadwick et al., 1999; Nolan et al., 2012).

Spectral fitting of BL Lacs and FSRQs allows not only jet powers and magnetic fields to be estimated but also the BLF of the jet. Ghisellini et al. (2010) typically find $\Gamma \sim 13$ is required to fit FSRQ spectra and $\Gamma \sim 15$ for BL Lacs. These are much higher BLFs than the mildly relativistic $\Gamma \sim 1.4 - 2$ more commonly assumed for BHBs. This suggests radio louder AGN (BL Lacs/FRIs/FSRQs/FRIIs and potentially BLRGs/NLRGs) may have much more relativistic jets than BHBs and, by extension, the LINERs, Seyferts and radio quiet quasars that lie on the same fundamental plane. Again, one possibility is that spin may be increasing the BLF in these objects.

However, not quite all Seyferts are radio and gamma-ray quiet. So far 7 radio loud NLS1s have been detected with gamma-ray emission (Foschini et al., 2015). Since NLS1s are high accretion rate systems ($\dot{m} \sim 1$) with $M \sim 10^7 - 10^8 \, M_\odot$, they should be scaled down versions of the higher mass FSRQs. Just like FSRQs, their broadband spectra show a low energy synchrotron hump, high energy EC dominated hump and require $\Gamma \sim 13$ to fit (Abdo et al., 2009), so they do appear like mini-FSRQs (Fig.3.18). The discovery of these gamma-ray loud NLS1s (γNLS1s) provides a new opportunity to test existing blazar jet models and scalings over an extended mass range.

Figure 3.18: Spectrum of the Gamma-Ray Loud Narrow Line Seyfert 1 PMN J0948+0022, from Foschini et al. (2012).

Chapter 4

Jets and the Accretion Flow in Low Luminosity Black Hole Binaries

4.1 Introduction

In this chapter we focus on accretion and ejection around black hole binaries in the low/hard state.

The low/hard state (LHS) of black hole binaries (BHBs) is typically seen at mass accretion rates below a few per cent of the Eddington limit. It is characterised by a hard X-ray spectrum, rising in νf_ν to a peak at a few hundred keV, in sharp contrast to the typical temperature of a few hundred eV expected from an optically thick, geometrically thin accretion disc. These hard X-rays are also strongly variable on short (sub second) timescales, again, in sharp contrast to the long (few hour) viscous timescale expected from even the innermost radii of a thin disc. These properties are more typical of the alternative set of solutions of the accretion flow equations, where the flow is geometrically thick, and optically thin. The most well known of these alternative solutions are the Advection Dominated Accretion Flow (ADAF) models (Narayan & Yi, 1995), but these are only an analytic approximation to what is almost certainly a more complex solution, as the flow must be threaded by magnetic fields. Differential rotation shears the field azimuthally, while buoyancy

Figure 4.1: Low/hard state X-ray photon index as a function of luminosity in Eddington units, from Sobolewska et al. (2011)

lifts it vertically, and the combination sets up a turbulent magnetic dynamo which acts to transport angular momentum outwards so material can fall inwards. Close to the horizon, this turbulent field can also produce a jet, as observed in this state (Fender et al., 2004).

These hot flow solutions are only possible at low mass accretion rates, collapsing to the standard disc solutions when the flow becomes optically thick. The transition is complex, but the data can be largely fit into a picture where the thin disc progressively replaces the hot flow down to smaller radii as the mass accretion rate increases. Such truncated disc models predict that the contribution from the thin disc becomes stronger with increasing mass accretion rate, increasing the seed photons for Compton cooling of the hot flow, so the hard X-ray spectrum becomes softer, as observed.

Conversely, as accretion rate drops, the X-ray spectrum should harden. In fact, the data show that the X-ray spectral index first becomes harder, but then softens again below an (Eddington scaled) luminosity of $L/L_{Edd} \sim 10^{-2}$ (e.g. Corbel et al., 2008; Russell et al., 2010; Sobolewska et al., 2011, hereafter S11, see Fig.4.1). The

truncated disc/hot inner flow model gives a possible explanation for this behaviour of the X-ray spectral index. The disc recedes as mass accretion rate drops, which leads to a decrease in the seed photon luminosity intercepted by the hot flow so its Compton spectrum hardens. However, there is another source of seed photons, from cyclo-synchrotron emission generated within the flow itself by the hot electrons spiralling in the turbulent magnetic field. This source of seed photons increases as the mass accretion rate drops, as the drop in density means that the emission is much less self-absorbed and this more than compensates for the drop in emissivity. Thus the flow should make a transition from hardening due to Compton scattering on the receding disc, to softening due to Compton scattering of cyclo-synchrotron photons within the flow (S11).

However there is an alternative possibility for the change in X-ray spectral index, where this marks instead the change from a flow dominated X-ray spectrum to a jet dominated X-ray spectrum (Russell et al 2010; S11). Such a transition to a jet dominated flow (JDAF) was suggested by Yuan & Cui (2005) and there are coupled ADAF-jet models in the literature where the X-rays are produced by synchrotron jet emission at low mass accretion rates (e.g. Yuan & Cui, 2005; Yu et al., 2011).

Here we make a simple model of the accretion flow (truncated disc and hot, radiatively inefficient inner flow) to gain a quantitative understanding of how the X-ray spectrum evolves with accretion rate in terms of the contribution of disc and cyclo-synchrotron seed photons for flow Comptonisation. We then couple this to a standard conical jet model (Blandford & Königl, 1979; Merloni et al., 2003) to assess the relative contribution of the flow and the jet to the X-ray emission, using independent constraints from the observed radio–X-ray correlation (Hannikainen et al., 1998; Corbel et al., 2003; Gallo et al., 2003; Corbel et al., 2013).

4.2 The Fiducial Truncated Disc/Hot Inner Flow Model

In all the following we use dimensionless radii $r = R/R_g$, where $R_g = GM/c^2$, and mass accretion rates $\dot{m} = \dot{M}/\dot{M}_{Edd}$, where the Eddington limit $L_{Edd} = \eta \dot{M}_{Edd} c^2$

and $\eta = 0.057$ for a Schwarzschild black hole with innermost stable circular orbit $r_{isco} = 6$. We plot models for a $10\,M_\odot$ black hole.

Our main aim is to explore the origin of the X-ray flux in the LHS, firstly whether the truncated disc/hot inner flow model can produce the observed change in behaviour of spectral index with $\dot{m}$, and then to see whether this can also be produced by jet models. Previous models which included both truncated disc seed photons and internally generated cyclo-synchrotron seed photons (Narayan et al., 1997; Esin et al., 1997) did not explicitly explore this, and are also based on a pure ADAF model for the accretion flow. Such pure ADAF models are too hot and optically thin (maximum optical depth $\tau \propto \dot{m} < 1$) to match the observed hard X-ray emission (Yuan & Zdziarski, 2004). Allowing advection to be negative (heating the flow) as well as positive (assumed in the ADAF solution) takes the (luminous hot accretion flow: LHAF) models closer to the data, but there is still a clear mismatch, with data extending up to an optical depth of $\tau \sim 2$ (Yuan & Zdziarski, 2004). This probably reflects the fact that all such analytic models are only an approximation to a more complex reality, with magnetic fields threading the flow. Hence rather than build a full ADAF/LHAF model, which is known not to match the data, we instead take the key aspects of these models (radiatively inefficient flow, i.e. luminosity $L \propto \dot{m}^2$, which exists only up to a maximum mass accretion rate, $\dot{m}_c$) and set the parameters of this flow from the data, i.e. we take $\dot{m}_c = 0.1$ and $\tau = \tau_{max}(\dot{m}/\dot{m}_c)$ with $\tau_{max} = 2$ (e.g. Ibragimov et al., 2005; Torii et al., 2011; Yamada et al., 2013).

In the truncated disc/hot inner flow geometry, this radiatively inefficient flow exists inside a Shakura-Sunyaev disc truncated at radius $r_t \geq r_{isco}$. Evaporation of the cool disc by thermal conduction from a hot corona is known to produce this geometry at low mass accretion rates (Liu et al., 2002; Mayer & Pringle, 2007), where it typically gives $r_t \propto \dot{m}^{-1/2}$ below the critical mass accretion rate, $\dot{m}_c$, at which the hot flow collapses (e.g. Czerny et al., 2004). Evaporation models show that the disc is still substantially truncated at this critical mass accretion rate, but the value of this minimum truncation radius is $\sim 40\,R_g$ ($20\,R_{sch}$: Czerny et al., 2004). However, the evaporation rates assume the hot flow is an ADAF, whereas our flow is denser and cooler. The conductive flux depends more strongly on density, so we expect

stronger evaporation. This combined with weaker constraints on the observed disc radius in the LHS (Yamada et al., 2013) motivates us to choose $r_t = 20(\dot{m}/\dot{m}_c)^{-1/2}$.

We assume a standard Novikov-Thorne emissivity for a disc from $r_{out} = 10^5$ to r_t, and assume that all this energy thermalises, giving L_{disc}. The remaining energy from the Novikov-Thorne emissivity from r_t to r_{isco} is available to power the hot flow, $L_{hot,power}$, but this is radiatively inefficient so we take the actual radiated power to be $L_{hot} = (\dot{m}/\dot{m}_c)L_{hot,power}$, i.e. assume that the flow is as efficient as a thin disc at $\dot{m}_c$.

The hot flow radiates L_{hot} via Comptonisation (which depends on seed photon luminosity from both the disc and cyclo-synchrotron photons generated by the electrons interacting with the magnetic field in the hot flow) and Bremsstrahlung (which depends on density). We assume that the hot flow is a homogeneous sphere. The obvious radius of this sphere is r_t, but the emission should be centrally concentrated, so instead we assume that all the energy is dissipated in a region $r_h = 20$. At any radius r in the disc, we calculate the fraction of photons illuminating the hot flow, so the seed photon luminosity, $L_{seed,disc}$ is given by this integrated over all the disc from r_{out} to r_t. The density of the flow is then $n \sim \tau/(\sigma_T r_h R_g)$.

Radiatively inefficient flows are also generically two temperature, with ion temperature set by the virial temperature $kT_{ion} \approx m_p c^2/r$, while the electron temperature is set by the balance of heating and cooling. We assume that the flow is homogeneous within r_h so $kT_{ion} \sim m_p c^2/r_h$. Simulations show that the energy density in the tangled magnetic field saturates to $\sim 10\%$ of the gas pressure, so $U_B = B^2/(8\pi) = 0.1nkT_{ion}$. The cyclo-synchrotron emission from the hot flow then extends as an approximate steep power law from $v_B = eB/(2\pi m_e c) = 2.6 \times 10^6 B$. However, the majority of this emission is self-absorbed, so the emission peaks instead at the self-absorption frequency $\nu_{csa} = \frac{3}{2}\nu_B \theta_e^2 x_m$ where the electron temperature $\theta_e = kT_e/m_e c^2$ (found iteratively, see below) and x_m typically has values of a few hundred to a few thousand (full details are given in the Appendix). The luminosity is then $L_{seed,cyclo} \propto n\nu_{csa}^2 V$, where $V = \frac{4}{3}\pi r_h^3 R_g^3$ is the volume of the hot flow.

The total seed photon luminosity $L_{seed} = L_{seed,disc} + L_{seed,cyclo}$. We take the seed photon energy (ν_{seed}) as the weighted mean of the inner disc temperature and the

cyclo-synchrotron self-absorption frequency. The electron temperature can then be derived self-consistently from balancing heating (L_{hot}) and cooling (determined by L_{seed}, but also including Bremsstrahlung) rates using the publicly available EQPAIR code. This calculates the electron temperature and resulting emission spectrum from a homogeneous sphere, given inputs of the heating power to the electrons (L_{hot}), the optical depth and size of the region (τ and r_h), and the power and typical energy of the seed photons (L_{seed} and ν_{seed}) for Compton cooling (Coppi, 1999). The resulting spectrum incorporates both Bremsstrahlung and Compton components and does not assume that the Compton emission can be approximated as a power law. This is increasingly important as the flow density drops, as each successive Compton scattering order is separated by a factor $1/\tau$, making the spectrum increasingly bumpy as the mass accretion rate decreases.

To summarise: our accretion flow model consists of a truncated disc where the truncation radius increases with decreasing $\dot{m}$, and a radiatively inefficient inner hot flow powered by the remaining gravitational energy that is not dissipated in the truncated disc. We allow the optical depth of this hot flow to decrease with $\dot{m}$, and use both the external disc photons intercepting the hot flow and internal cyclo-synchrotron photons generated within the hot flow as seed photons for Comptonisation.

4.2.1 Spectral Changes with Accretion Rate

Fig.4.2 shows a sequence of model spectra for $\dot{m} = \dot{m}_c = 10^{-1}$ to $\dot{m} = 4 \times 10^{-3}$ (i.e. $r_t = 20$ to 100). Solid lines show the total emission, long dashed, short dashed and dotted lines show the individual components of cyclo-synchrotron, truncated disc and Comptonisation, respectively.

The proportion of luminosity in the disc compared to the hot flow, $L_{disc}/L_{hot} \approx (\dot{m}/\dot{m}_c)^{-1}(\frac{r_t}{r_{isco}} - 1)^{-1}$. Since we also know how r_t depends on $\dot{m}$ we can simplify this further to $\approx 0.3(\dot{m}/\dot{m}_c)^{-1/2} \propto r_t$. Thus decreasing $\dot{m}$ by a factor of 25 increases the disc truncation radius by a factor of 5 and increases L_{disc}/L_{hot} by a factor 5. This is not a large factor, but is evident in Fig.4.2 by comparing the ratio between the peak νf_ν flux of the disc and Comptonised emission for the highest and lowest

Figure 4.2: Model SEDs, with truncated disc (short dashed line), hot flow cyclo-synchrotron emission (long dashed line) and Comptonisation of both disc and cyclo-synchrotron seed photons (dotted line) for increasing truncation radius: $20\,R_g$ (black), $35\,R_g$ (blue), $50\,R_g$ (green), $70\,R_g$ (red) and $100\,R_g$ (magenta). Solid line shows sum of all three components.

Figure 4.3: a). Seed photon luminosity as a function of truncation radius, for disc seed photons (red dotted line) and cyclo-synchrotron seed photons (green dashed line). For $R_{trunc} \gtrsim 60\,R_g$ ($\dot{m} \lesssim 0.01$) the dominant source of seed photons is cyclo-synchrotron emission from the hot flow. b). Photon index as a function of 2–200 keV X-ray luminosity, showing softening of the X-ray spectrum at low luminosities as cyclo-synchrotron seed photons begin to dominate. c). Hot flow electron temperature as a function of mass accretion rate, where $\dot{m} = \dot{M}/\dot{M_E}$.

$\dot{m}$ spectra.

However, the ratio between $L_{seed,disc}/L_{hot}$ changes by much more than L_{disc}/L_{hot} as the fraction of seed photons intercepted by the hot flow drops as r_t increases. The seed photons from the disc which illuminate the hot flow are integrated over the entire disc, but both the disc luminosity and the fraction which are intercepted by the flow will peak at r_t. Hence $L_{seed,disc} \approx L_{disc}(r_h/r_t) \arcsin(r_h/r_t) \approx \frac{\dot{m}}{r_t}(r_h/r_t)^2 \propto r_t^{-5} \propto \dot{m}^{2.5}$. Thus $L_{seed,disc}/L_{hot} \propto \dot{m}^{2.5}/\dot{m}^2 \propto \dot{m}^{1/2} \propto r_t^{-1}$. Thus while L_{disc}/L_{hot} increases by a factor 5 as $\dot{m}$ decreases, $L_{seed,disc}/L_{hot}$ decreases by a factor 5 (see red line in Fig.4.3a). If this were the only source of seed photons, the spectrum should harden substantially. However, there are also seed photons from the cyclo-synchrotron emission. These have $L_{seed,cyc} \propto n\nu_{csa}^2 \propto n(B\theta_e^2)^2 \propto \dot{m}^2\theta_e^4$, so $L_{seed,cyc}/L_{hot} \propto \theta_e^4$. This increases as $\dot{m}$ decreases, as θ_e increases as accretion rate drops (see below). The green line on Fig.4.3a shows the internally generated cyclo-synchrotron emission starts to dominate over seed photons from the disc at $r_t > 60$ (equivalently $\dot{m} \le 10^{-2}$). Thus the total L_{seed}/L_{hot} reaches a minimum at this point, and then starts to increase. This change in dominant seed photons can also be seen in Fig.4.2 as the Compton spectrum extends to lower energies reflecting the lower seed photon energy of the cyclo-synchrotron photons.

The Comptonisation spectral slope is set by L_{seed}/L_{hot}, so this also shows a minimum corresponding to the minimum L_{seed}/L_{hot}. Fig.4.3b shows the resulting 2–10 keV power law index and 2–200 keV bolometric luminosity, L_{2-200}. Our minimum in photon index occurs at $2 - 3 \times 10^{-3} L_{Edd}$, a factor 2–3 below that shown by the data in S11 (Fig.4.1). Given the simple assumptions made about the structure of the flow, this is probably not significant. In particular, changing the efficiency of the hot flow from a simple $\propto \dot{m}$ to the more complex behaviour calculated by Xie & Yuan (2012) specifically for an ADAF model would make this discrepancy smaller. Thus the model is able to quantitatively describe a key observation of the LHS, namely that the X-ray spectrum hardens with decreasing $\dot{m}$ and then softens again by the change in seed photons from the disc to internally generated cyclo-synchrotron. This softening of the Comptonised emission can be seen by eye in the spectra of Fig.4.2 by comparing the slope of the tail at highest and lowest luminosity.

Qiao & Liu (2013) find a similar trend in photon index using a full ADAF calculation. However their disc-corona geometry is rather different from that considered here. They focus on the residual inner disc which can remain after considering thermal conduction from the hot flow (Liu et al., 2006; Liu et al., 2011). At the highest mass accretion rates, the disc extends all the way down to the innermost stable circular orbit. Then as accretion rate drops a gap opens up between the inner and outer disc at $\sim 200\,R_g$, and this gap extends inwards and outwards as $\dot{m}$ decreases until the entire inner disc evaporates. Thus their drop in seed photons comes from a decreasing outer extent of the inner disc, whereas in our model it comes from the increasing inner radius of the outer disc. Nonetheless, both models have a drop in seed photons with mass accretion rate, so the spectra harden, and then both models show the characteristic minimum as self-generated cyclo-synchrotron photons take over as the dominant seed photons in the hot flow.

Fig.4.3c shows the resulting electron temperature, set from the balance of heating and cooling. The heating rate is $\propto \dot{m}^2$, and cooling is predominantly Compton cooling so is $\propto 4\theta_e\tau L_{seed}$. At high $\dot{m}$, the seed photons are from the disc so the cooling rate is $\propto 4\theta_e\tau\frac{\dot{m}}{r_t}(\frac{r_h}{r_t})^2$. Hence $\theta \propto \dot{m}^{-3/2}$, quite close to the observed dependence. Conversely, when seed photons from cyclo-synchrotron cooling dominate, the Compton cooling rate is $\propto 4\theta_e\tau n\nu_{csa}^2$, where $\nu_{csa} \propto B\theta_e^2$ so $\theta_e \propto \dot{m}^{-0.2}$. The strong increase in seed photons with increasing temperature leads to increasing cooling with decreasing mass accretion rate, which counteracts much of the decrease in cooling from the decrease in optical depth. Thus the electron temperature increases much more slowly as the mass accretion rate decreases. Again this can be seen in the spectra of Fig.4.2, where the electron temperature (marked by the high energy rollover of the tail) first increases markedly with decreasing mass accretion rate, then stabilises. This changing temperature dependence on accretion rate is a testable prediction of the model. Current observations already show that the brightest LHS spectra do indeed show the predicted decrease in temperature with increasing $\dot{m}$ (Motta et al., 2009; Torii et al., 2011), but future observations with the more sensitive Soft Gamma Ray detector (60–600 keV bandpass) on ASTRO-H (Takahashi et al., 2012) will be able to constrain the temperature down to much lower $\dot{m}$.

The other obvious change in the Comptonised emission is that it is progressively less well described by a power law as $\dot{m}$ decreases and $\tau \ll 1$. At such low optical depths the individual scattering orders become visible, giving a more complex spectral shape. Data are rarely fit with such low optical depths, as X-ray observations do not show the strong first Compton peak, which would be clearly visible in the X-ray regime if the seed photons were provided by a disc. From our model it is clear the dominant source of seed photons at these mass accretion rates is cyclo-synchrotron emission. This brings the first peak out of the X-ray regime, leaving the X-ray spectrum to be dominated by higher order scattering with less extreme curvature. Nevertheless, this is still not visible in X-ray spectra from low L/L_{Edd} flows (e.g. Corbel et al., 2006). We suggest the reason for this is that our model assumes that the electrons in the hot flow completely thermalise. An initially non-thermal acceleration process will probably thermalise via self-absorption of its own cyclo-synchrotron radiation in a bright LHS (Malzac & Belmont, 2009; Poutanen & Vurm, 2009). However, the thermalisation timescale increases as the source luminosity drops, so the electron distribution retains more of its initially non-thermal character, giving a non-thermal power law Compton spectrum (e.g. Veledina et al., 2011).

4.3 Fiducial Conical Jet

The radio jet is an important part of the energy budget of the black hole accretion flow, with kinetic energy comparable to the hard X-ray luminosity at $\dot{m}_c$ (e.g. Cyg X-1: Gallo et al., 2005; Russell et al., 2007; Malzac et al., 2009). Hence we take the jet $L_{KE,max} = \dot{m}_c L_{Edd} = 1.3 \times 10^{38}$ erg s^{-1}.

We add a standard conical jet model onto our accretion flow (e.g. Blandford & Königl, 1979; hereafter BK79; Merloni et al., 2003; Falcke et al., 2004), assuming that some acceleration process operates continuously down the jet, so that a small fraction of the electrons in the jet form a relativistic particle distribution. The electrons radiate via synchrotron to produce a broad band spectrum from radio to X-rays. If these radiative losses are high then this will affect the self-similar jet

structure. Hence we limit the radiative luminosity to 10% of the kinetic luminosity of the jet i.e. we make a maximally radiatively efficient, self-similar jet.

Parameters for a standard conical jet include the distance from the black hole at which the material is accelerated, Z_0 (the jet base). We make the standard assumption that the energy is transported by Poynting flux from r_h in the hot flow, where the jet is presumably launched, to z_0 without any radiative losses. The self-similar behaviour then extends out from Z_0 to a distance of $Z_{max} = 10^6 Z_0$, where $Z = zR_g$ is distance along the jet. Distance perpendicular to the jet is $R_j = \rho R_g = \phi Z$, where ϕ is constant for a conical jet.

We use observations to set the bulk Lorentz factor $\Gamma = 1.2$ and opening angle $\phi = 0.1$ (e.g. Gallo et al., 2005). We assume that these stay constant with $\dot{m}$. We transform all specific luminosities, L_ν, from jet frame to observer frame by multiplying by $\delta^3 = (\Gamma - \sqrt{\Gamma^2 - 1}\cos\psi)^{-1}$, assuming a mean inclination angle $\psi = 60°$, and boost all frequencies by δ.

In such a geometry, the magnetic field energy density $U_B(z) \propto z^{-2}$ (BK79). Turbulence in the field probably results in scaling between the relativistic particle and magnetic pressures, so $U_{rel}(z) = m_e c^2 \int_1^{\gamma_{max}} N(z,\gamma)\gamma d\gamma = f_{rel}U_B(z)$ where $N(z,\gamma)$ is the electron distribution at each point z of the jet. We make the standard assumptions that $N(z,\gamma) = K(z)\gamma^{-p}$ with $p = 2.4$ between $\gamma_{min} = 1$ and $\gamma_{max} = 10^5$. Hence the optically thin synchrotron emission has energy index $\alpha = (p-1)/2 = 0.7$, i.e. it rises in νf_ν with energy output peaking at the highest frequency $\nu_{max} = 4/3\gamma_{max}^2\nu_B$.

The power law synchrotron emission becomes optically thick to self-absorption below $\nu_{ssa} \propto K^{2/7}B^{5/7}R_j^{2/7} \propto (z/z_0)^{-1}$ (Ghisellini et al., 1985), i.e. decreases with larger distance along the jet. The flux at this point is therefore:

$$L_{sync}(\nu_{ssa}) \propto (z/z_0)^{-1}(\nu_{ssa}/\nu_B)^{-(p-1)/2}dZ \propto (z/z_0)^{-1}dZ \propto d\log Z \qquad (4.3.1)$$

Thus the self-absorbed spectra from each part of the jet sum together to produce the characteristic 'flat spectrum' (i.e. energy index $\alpha = 0$: BK79) at low frequencies.

We do not include the self-consistent inverse Compton emission from the jet,

since this contributes only at higher energies (e.g. Zdziarski et al., 2012) and the baseline model we are testing is one where the X-rays are produced by synchrotron from the jet.

4.3.1 Jet at $\dot{m}_c$

We anchor the jet at $\dot{m}_c = 0.1$ using observational constraints. The observed break from optically thick to optically thin synchrotron in a bright LHS from GX 339–4 is $\nu_{ssa,0} \sim 10^{13.5}$ Hz, and the 10 GHz radio luminosity from the sum of self-absorbed jet components is six orders of magnitude below the X-ray emission (Gandhi et al., 2011), i.e. $\nu L_\nu \sim 10^{31}$ erg s^{-1} at 10 GHz. This sets $B(z_0) \sim 3.5 \times 10^4$ G (i.e. $K(z_0) = 2.4 \times 10^{12}$ cm^{-3} for $f_{rel} = 0.1$) and $z_0 \sim 5300$. This gives a total radiated luminosity of 10% of the kinetic jet power, as described above for a maximal radiatively efficient jet.

We note that in these standard conical jets the fraction of radiative power to kinetic power is not constant down the jet, as the radiation depends on both magnetic and electron energy density so is $\propto (z/z_0)^{-4}dV$, while the jet kinetic energy is simply $\propto (z/z_0)^{-2}dV$.

4.3.2 Jet Scaling with Mass Accretion Rate – Transition to a Jet Dominated State?

We then assume that the jet kinetic power $\propto \dot{m}$ and that all energy densities scale as $\dot{m}/\dot{m}_c$ (Heinz & Sunyaev, 2003), so that $B(z, \dot{m}) = B_0(z_0, \dot{m}_c)(z/z_0)^{-1}(\dot{m}/\dot{m}_c)^{1/2}$ and $K(z, \dot{m}) = K_0(z_0, \dot{m}_c)(z/z_0)^{-2}(\dot{m}/\dot{m}_c)$. Fig.4.4a shows a sequence of spectra for decreasing $\dot{m}$ using this coupled accretion flow-jet model. Our model reproduces the $L_R \propto L_X^{0.7}$ radio–X-ray correlation, as shown in Fig.4.4b. The X-rays come from a radiatively inefficient accretion flow and are therefore proportional to $\dot{m}^2$. The radio is from the optically thick jet, where it has a flat spectrum, so $L_R \propto B_0^{1.2}K_0^{0.8} \propto \dot{m}^{1.4}$ for any model where the magnetic and relativistic particle energy densities scale with $\dot{m}$, hence $L_R \propto L_X^{0.7}$ for a radiatively inefficient X-ray flow (e.g. Heinz & Sunyaev, 2003; Merloni et al., 2003). The model slightly deviates from this relation at low $\dot{m}$,

Figure 4.4: a). Model SEDs, including synchrotron jet emission (dotted line) for increasing truncation radius: $20\,R_g$ (black), $35\,R_g$ (blue), $50\,R_g$ (green), $70\,R_g$ (red) and $100\,R_g$ (magenta). b). $L_R \propto L_X^{0.7}$ radio–X-ray correlation (black line), together with correlation from coupled accretion flow/jet model (solid red line). Dashed red line shows radio–X-ray correlation using $L_{20-100\,\mathrm{keV}}$ X-ray luminosity instead of $L_{3-9\,\mathrm{keV}}$. c). X-ray luminosity as a function of mass accretion rate, where $\dot{m} = \dot{M}/\dot{M}_E$, for X-rays from the radiatively inefficient accretion flow (solid red line) and X-rays from the jet (dotted blue line).

because of the spectral curvature in our model at low $\dot{m}$ which changes the scaling over a small bandpass ($L_{3-9\text{kev}}$). The dashed line in Fig.4.4b shows that a wider bandpass recovers the relation even down to the lowest luminosities. As discussed in Section 2.1, this detailed issue can probably be circumvented by a proper treatment of the self-consistent electron distribution in the hot flow (as in Veledina et al., 2011).

The jet emission does not ever dominate the total hard X-ray emission, but remains an approximately constant factor below the hot flow (Fig.4.4c). This is because the optically thin synchrotron luminosity $L_{X,jet} \propto K_0 U_{B,0} \propto \dot{m}^2$ so it also follows a radiatively inefficient scaling. This is in contrast to the jet *kinetic* luminosity, which does scale as $\dot{m}$. Thus while the *kinetic* luminosity of the jet can easily dominate the radiative energy of the flow, the *radiated* energy of the jet drops as fast as that from the flow. Thus there is no transition in the X-ray spectrum from being dominated by the radiatively inefficient hot flow to being dominated by the synchrotron emission of a conical, self-similar jet (see also Merloni et al., 2003; Falcke et al., 2004). There are instances in the literature where there is a flow-jet transition in the X-ray flux, but these use models where the flow is radiatively efficient (no advection: Fender et al., 2003) and/or have a jet kinetic power which does not scale as $\dot{m}$ (Yuan & Cui, 2005).

4.3.3 Jet Scaling with Electron Cooling

The discussion above assumes that the energy density in relativistic particles scales as the energy density in the jet, i.e. $\int N(\gamma)\gamma d\gamma \propto \dot{m}/M$ (Heinz & Sunyaev, 2003). However, a better approach is to say that it is the injected electron distribution, $Q(\gamma) = Q_0 q(\gamma)$, which scales, and then cools into a steady state distribution $N(\gamma) = Kn(\gamma)$ (e.g. for Blazar jets: Ghisellini et al., 2010). In this case, the injected distribution is normalised to the available power (i.e. $Q_0 \propto \dot{m}$), so that K now scales as $K \propto Q_0/U_{seed}$. For synchrotron cooling, $U_{seed} = U_B \propto \dot{m}$ giving K constant with accretion rate. Hence $L_{X,jet} \propto K_0 U_{B,0} \propto \dot{m}$, such that it is possible for the jet X-rays to overtake the X-ray luminosity from the flow as accretion rate decreases. However, $L_R \propto B_0^{1.2} K_0^{0.8}$, so the radio luminosity no longer scales as $L_R \propto \dot{m}^{1.4}$ but

Type of Jet	L_R	$L_{X,jet}$	$L_{X,flow}$	Dominant Source of X-ray Emission	α_{flow} $(L_R \propto L_{X,flow}^\alpha)$	α_{jet} $(L_R \propto L_{X,jet}^\alpha)$
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Standard	$\dot{m}^{1.4}$	$\dot{m}^2$	$\dot{m}^2$	Either Hot Flow or Jet	0.7	0.7
Cooling	$\dot{m}^{0.6}$	$\dot{m}$	$\dot{m}^2$	Hot Flow with a transition to Jet X-rays as $\dot{m}$ decreases	0.3	0.6

Table 4.1: Scalings with accretion rate for a standard self-similar conical synchrotron jet, and the same synchrotron jet including cooling. (2) Dependence of radio luminosity on Eddington scaled accretion rate. (3) Dependence of jet X-ray luminosity on accretion rate. (4) Dependence of flow X-ray luminosity on accretion rate, assuming a radiatively inefficient hot flow. (5) Dominant source of X-ray emission as a function of accretion rate. (6) Index of the radio–X-ray correlation for X-rays from the flow. (7) Index of the radio–X-ray correlation for X-rays from the jet.

scales as $L_R \propto \dot{m}^{0.6}$ (Table 4.1). When the X-rays come from the jet, this gives $L_R \propto L_{X,jet}^{0.6}$, which is not inconsistent with the data. But for higher accretion rates, when the X-rays come from the flow (and we know $L_{X,flow}$ must be proportional to $\dot{m}^\beta$ where $\beta > 1$ for a transition to occur at all), this becomes $L_R \propto L_X^{0.3}$, which does not match the observed correlation. The optically thick synchrotron from the jet must still drop as $\dot{m}^{1.4}$ to make the observed $L_R - L_X$ relation when the hot flow dominates.

4.3.4 Composite Jet with Electron Cooling Break

Since energetic electrons cool faster, the electron distribution is expected to be composite, with the electron distribution above some break energy, γ_{cool}, being dominated by cooling, while below this energy it reflects instead the injected electron distribution (e.g. Markoff et al., 2005; Yuan et al., 2005; Zdziarski et al., 2012). The cooling break energy $\propto 1/(U_B Z) \propto \dot{m}^{-1}$ (Zdziarski et al., 2012, equation 36), so does not depend on mass but increases linearly with decreasing mass accretion rate. For our parameters, the cooling break is $\gamma_{cool} \sim 1.3$ for $\dot{m} = \dot{m}_c$, increasing to $\gamma_{cool} \sim 26$ for our lowest $\dot{m}$. Even the highest γ_{cool} is mostly below the synchrotron self-absorption break, so makes very little difference in the spectrum (see e.g. Zdziarski et al., 2012 Fig. 5a) or in the predicted X-ray scaling from the completely cooling dominated jet described above.

4.3.5 Arbitrary Jet Scaling

It is possible to contrive situations for a synchrotron jet where the radio scales as $\dot{m}^{1.4}$ but the X-rays scale as $\dot{m}$ (e.g. by allowing z_0 to change with accretion rate). But it is clear any transition from a flow where $L_X \propto \dot{m}^2$ to this jet, where $L_X \propto \dot{m}$, will cause a steepening of the observed $L_R - L_X$ relation (Yuan & Cui, 2005). Changing the X-ray behaviour with $\dot{m}$, without a simultaneous change in the behaviour of the optically thick radio emission, necessarily changes the $L_R - L_X$ correlation in a way which is not observed (Corbel et al., 2013). Since we do not observe a change in the radio–X-ray correlation down to quiescence in BHBs (e.g Corbel et al., 2013 and references therein), we know there can be no transition in the dominant X-ray production mechanism down to these luminosities. Whatever dominates the X-rays in the brightest LHS spectra, at the top of the correlation, must dominate at the bottom. This rules out all plausible models in which the X-rays switch from being dominated by the flow to being dominated by the jet.

We note that a break has been observed in the radio–X-ray correlation in active galactic nuclei (AGN) at very low accretion rates ($\dot{m} \sim 10^{-6}$, Yuan et al., 2009), but crucially the observed minimum in photon index occurs where the correlation is unbroken, implying that jet emission taking over cannot be the cause.

4.3.6 Jet Dominated Models

Since a switch from flow dominated to jet dominated X-ray flux is ruled out by the radio–X-ray correlation, the final possibility is that the X-rays are always dominated by the jet (Falcke et al., 2004). However our fiducial jet model is already very efficient at producing radiation. To make the jet dominate at $\dot{m}_c$ would require that almost all of the jet kinetic energy was transformed to radiation, which seems unlikely. It would also impact on our assumptions that adiabatic and radiative losses are negligible, and leaves unanswered the question of what causes the change in X-ray behaviour (hardening then softening of the spectral index) as $\dot{m}$ decreases. Jet dominated models are also unable to fit the rollover at high energies seen in the bright LHS ($\dot{m} \sim \dot{m}_c$: Ibragimov et al., 2005; Torii et al., 2011). More spectral

and energetic constraints are discussed in Zdziarski et al. (2003) and Malzac et al. (2009).

4.4 Conclusions

The observed change in X-ray spectral index as $\dot{m}$ decreases in the LHS (first hardening then softening) can be quantitatively explained by a truncated disc/radiatively inefficient hot inner flow. Seed photons from the disc drop as the disc recedes with decreasing $\dot{m}$ so that self-generated cyclo-synchrotron seed photons in the flow become dominant in Compton cooling. This model can also produce the observed radio–X-ray correlation with the addition of a standard, conical self-similar jet. These standard jets are as radiatively inefficient as the hot flow, so there is no transition from the X-rays being dominated by the flow to being dominated by the jet, which was the alternative explanation for the X-ray spectral behaviour (Russell et al., 2010, S11). Including the effects of cooling allows the jet X-rays to drop more slowly with accretion rate and hence overtake the X-rays from the hot flow, however such a transition would also necessarily distort the radio–X-ray correlation in a way which is not observed.

Thus we show that the truncated disc/radiatively inefficient hot flow/standard conical jet model can quantitatively explain the broad band spectral evolution of BHBs in the LHS, with the X-rays always being dominated by the flow, and the radio by the jet. However, at low luminosities, the optical depth in the hot flow is rather small, so the X-ray spectra are no longer a power law and the individual Compton scattering orders can clearly be seen. Yet the observed BHB spectra at these low mass accretion rates (quiescence) are well described by a power law (e.g. Gallo et al., 2006). This discrepancy is even more evident in the low mass accretion rate AGN (e.g Yu et al., 2011), which has again led to models where the X-rays are dominated by the jet. Since these are inconsistent with the observed radio–X-ray correlation, we suggest instead that this points to a more complex flow, where the electron acceleration process produces an intrinsically non-thermal distribution. Thermalisation via cyclo-synchrotron emission and absorption produces the domi-

nant thermal electron population of the bright LHS, while the dramatic increase in seed photons from the disc in the high/soft state means that the power law distribution is seen (Malzac & Belmont, 2009; Vurm & Poutanen, 2009). We suggest at very low accretion rates the cooling is so inefficient that thermalisation does not happen and the electron distribution remains non-thermal. A non-thermal electron distribution will emit a power law spectrum, as observed. Hybrid thermal/non-thermal models, especially combined with a multi-zone approach (Veledina et al., 2013), hold out the possibility to understand the broad band spectral variability both on long timescales, with changing accretion rate, and on short timescales, to understand how fluctuations in the flow can make the observed IR/optical/X-ray correlations (Kanbach et al., 2001; Malzac et al., 2004; Gandhi et al., 2008).

Chapter 5

What Powers the Most Relativistic Jets? I. BL Lacs

5.1 Introduction

In this chapter we focus on supermassive black holes in the same accretion regime as low/hard state black hole binaries, whose X-ray spectra definitely are dominated by jet emission.

Relativistic jets are the most dramatic consequence of accretion onto stellar mass black hole binaries (BHBs) and supermassive black holes (SMBHs). Blazars are extreme examples of this; they are a class of Active Galactic Nuclei (AGN) where the jet is viewed very close to the line of sight so its emission is maximally boosted by the relativistic bulk motion and can dominate the spectrum of the AGN from the lowest radio energies up to TeV. However, despite years of study, the fundamental issues of powering and launching these jets are not understood. There is general agreement only that it requires magnetic fields, but whether these can be generated solely from the accretion flow or whether the jets also tap the spin energy of the black hole (Blandford & Znajek, 1977) is still an open question. It is also difficult to test this observationally as neither black hole spin nor total jet power are easy to measure, leading to divergent views, e.g. in BHBs compare Russell et al. (2013) with Narayan & McClintock (2012) and in SMBHs compare Sikora et al. (2007) with Broderick & Fender (2011).

By contrast, the radiation emitted from the jet is fairly well understood, with spectra separating blazars into two types: BL Lacs and Flat Spectrum Radio Quasars (FSRQs). BL Lac spectra are typically completely dominated by jet emission, showing a double humped synchrotron self-Compton (SSC) spectrum. The FSRQs are more complex, showing clear signatures of a 'normal' AGN disc and broad line region (BLR), in contrast to BL Lac spectra which lack both a disc and broad lines. The presence of the disc and BLR in FSRQs means that there is an additional source of seed photons for cooling of relativistic particles in the jet, so their jet emission includes both SSC and external-Compton (EC) components, making their Compton component more luminous (Dermer et al., 1992; Sikora et al., 1994).

Thus the nature of the accretion flow itself is different in BL Lacs and FSRQs. This is linked to a clear distinction in Eddington ratio between BL Lacs and FSRQs, with the BL Lacs all consistent with $\dot{m} = \dot{M}/\dot{M}_{Edd} < 0.01$ (where $\eta \dot{M}_{Edd} c^2 = L_{Edd}$ and efficiency η depends on black hole spin), while the FSRQs have $\dot{m} > 0.01$ (see e.g. Ghisellini et al., 2010, hereafter G10). Thus the BL Lacs have low accretion rates and hot, geometrically thick, radiatively inefficient accretion flows, while FSRQs accrete at higher rates and have standard disks (see e.g. Ghisellini & Tavecchio, 2008b; G10; Ghisellini et al., 2011, Best & Heckman, 2012).

Since all BL Lacs are associated with a low $\dot{m}$ accretion flow, we test here the hypothesis that all low accretion rate flows can launch a BL Lac type jet whose properties are determined simply by the black hole mass and mass accretion rate. We use the simplest possible scalings for how the jet (emission region size, magnetic field and injected power) scales with these parameters (Heinz & Sunyaev, 2003; Heinz, 2004), anchoring our scalings onto the fits to individual BL Lac objects of G10, and use cosmological simulations to predict the number densities of black holes with $\dot{m} < 0.01$, assuming that each of these will produce an appropriately scaled, randomly orientated, BL Lac type jet. In this way, we can predict how many BL Lacs should be detected by the Fermi Large Area Telescope and compare our results with observations, so giving some statistical constraints on what conditions may be required to power highly relativistic BL Lac jets.

5.2 Synchrotron Self-Compton Jets

We adopt a single zone SSC model of the type used by Ghisellini & Tavecchio (2009), which self-consistently determines the electron distribution from cooling. We briefly summarise our model here, with full details in the Appendix.

We assume a spherical emission region of radius R_{diss}. We neglect the contribution from regions further out along the jet, as these only make a difference to the low energy (predominantly radio) emission. We assume material in the jet moves at a constant bulk Lorentz factor (Γ), and that a fraction of the resulting jet power is used to accelerate electrons in the emission region. The power injected into relativistic electrons is then $P_{rel} = 4/3\pi R_{diss}^3 \int \gamma m_e c^2 Q(\gamma) d\gamma$, where the accelerated electron distribution is a broken power law of the form:

$$Q(\gamma) = Q_0 \left(\frac{\gamma}{\gamma_b}\right)^{-n_1} (1 + \gamma/\gamma_b)^{n1-n2} \text{ for } \gamma_{min} < \gamma < \gamma_{max} \qquad (5.2.1)$$

These electrons cool by emitting self-absorbed synchrotron and synchrotron self-Compton radiation, so the seed photon energy density $U_{seed} = U_B + g(\gamma)U_{sync}$ and includes both the magnetic energy density ($U_B = B^2/8\pi$) and the fraction $g(\gamma)$ of the energy density of synchrotron seed photons (U_{sync}) which can be Compton scattered by electrons of energy γ within the Klein-Nishina limit. This gives rise to a steady state electron distribution, $N(\gamma) = -\dot{\gamma}^{-1} \int_\gamma^{\gamma_{max}} Q(\gamma')d\gamma'$, where the rate at which an electron loses energy $\dot{\gamma}m_e c^2 = 4/3\gamma^2\sigma_T c U_{seed}$. However, this assumes that the electrons can cool within a light crossing time, but the cooling timescale $t_{cool} = \gamma/\dot{\gamma}$ itself depends on γ, with high energy electrons cooling fastest. We calculate the Lorentz factor that can just cool in a light crossing time of the region, γ_{cool}, and join smoothly onto the accelerated electron distribution below this. The full self-consistent electron distribution can be characterised by $N(\gamma) = Kn(\gamma)$, where K is the number density of electrons at $\gamma = 1$ and $n(\gamma)$ incorporates all the spectral shape. We calculate the resulting (self-absorbed) synchrotron and self-Compton emission using the delta function approximation, as this is much faster than using the full kernel but is accurate enough for our statistical analysis (Dermer & Menon, 2009).

This jet frame emission is boosted by the bulk motion of the jet, with the amount of boosting depending on both Γ and the orientation of the jet. The emission is then cosmologically redshifted and attenuated due to pair production on the extragalactic infrared background light (though this is generally small for the Fermi bandpass) to produce the observed flux.

The parameters of our model are therefore:

- Physical parameters of the jet: Γ and radius of emission region R_{diss}.

- The magnetic field of the emission region and power injected into relativistic electrons (B and P_{rel}).

- Parameters of the injected electron distribution: γ_{min}, γ_b, γ_{max}, n_1 and n_2.

We adopt the cosmology used in the Millennium simulations: $h = 0.72$, $\Omega_m = 0.25$, $\Omega_{vac} = 0.75$ (Springel et al., 2005; Fanidakis et al., 2011).

5.3 Scaling Jets

We assume that the acceleration mechanism is the same for all BL Lacs, giving the same injected electron distribution, regardless of mass and accretion rate. We also assume all jets are produced with the same Γ. This leaves three remaining parameters: R_{diss}, B and P_{rel}.

We scale $R_{diss} \propto M$, since all size-scales should scale with the mass of the black hole (Heinz & Sunyaev, 2003). We assume the jet power is a constant fraction of the total accretion power, $P_j \propto \dot{m}M$. This assumption is valid whether the jet is powered by the accretion flow or the spin energy of the black hole, since extraction of black hole spin energy relies on magnetic fields generated in the accretion flow, which will be affected by accretion rate. A constant fraction of the total jet power is then injected into relativistic particles and magnetic fields. Hence $P_{rel} \propto P_j \propto \dot{m}M$.

Energy density in the jet frame is related to power in the rest frame via $P = \pi R_{diss}^2 \Gamma c U$, so $P_B \propto R_{diss}^2 U_B \propto \dot{m}M$, hence $B \propto U_B^{1/2} \propto (\dot{m}/M)^{1/2}$. Therefore all energy densities should scale as $U_B \propto U_{rel} \propto (\dot{m}/M)^{1/2}$.

We anchor this with parameters from the fit to the classic low peaked BL Lac (LBL) object, 1749+096 from G10, which is relatively near to the top of the BL Lac accretion rate range. This gives $M_0 = 7 \times 10^8 \, M_\odot$, $R_{diss,0} = 172 \times 10^{15}$ cm, $B_0 = 1$ G, $P_{rel,0} = 3.5 \times 10^{42}$ erg s^{-1}, $\Gamma = 15$, $\gamma_{min} = 1$, $\gamma_b = 2 \times 10^3$, $\gamma_{max} = 1 \times 10^5$, $n_1 = 0.9$, $n_2 = 2.8$, and we scale R_{diss}, P_{rel} and B as:

$$R_{diss} = R_{diss,0} \frac{M}{M_0} \tag{5.3.2}$$

$$P_{rel} = P_{rel,0} \frac{\dot{m}}{\dot{m}_0} \frac{M}{M_0} \tag{5.3.3}$$

$$B = B_0 \left(\frac{\dot{m}}{\dot{m}_0} \frac{M_0}{M} \right)^{1/2} \tag{5.3.4}$$

We calculate the accretion rate of 1749+096 following the method of G10. Assuming the jet is maximal ($P_j = \dot{M}c^2$), we sum the power in magnetic fields, relativistic electrons and the bulk motion of cold protons to calculate P_j, giving:

$$\dot{m}_0 = \frac{P_j \eta}{1.38 \times 10^{38}(M/M_\odot)} \sim 3.5 \times 10^{-3} \tag{5.3.5}$$

for their value of $\eta = 0.08$. Fig.5.1 shows our model spectrum, together with the model and data from G10. The two models differ slightly due to our use of the delta function approximation to speed up calculation time. Nevertheless the two models are in agreement within ~ 0.3 dex, and crucially our model reproduces the correct level of Fermi flux (red bow tie).

5.4 Transition from High Frequency Peaked to Low Frequency Peaked BL Lacs with Accretion Rate

We limit our model to a maximum accretion rate of $\dot{m} = 10^{-2}$, since above this the accretion flow is expected to make a transition to a radiatively efficient thin disc. The strong UV and consequent broad line region emission provide additional seed

Figure 5.1: Model spectrum using parameters for 1749+096 (z=0.322) from Ghisellini et al. (2010) (solid black line). Dotted line shows their spectrum for the same parameters and red points show their data.

photons, switching the main Fermi radiation process from SSC (BL Lacs) to EC (FSRQs).

Fig.5.2a shows a sequence of spectra with R_{diss}, P_{rel} and B scaling as described above for $\dot{m} = 10^{-2}$ (black) – 10^{-4} (magenta) for constant mass. This shows the systematic decrease in luminosity, coupled to a change in spectral shape from a low synchrotron peak energy (optical: LBL) to a high synchrotron peak energy (X-ray: HBL), as shown by Ghisellini & Tavecchio (2008b) and Ghisellini & Tavecchio (2009).

We can compare the Fermi flux levels of our lower accretion rate spectra with observed HBLs. The HBL 1959+650 (see Tavecchio et al., 2010 for a spectrum, G10 for spectral fitting parameters) has a mass of $2 \times 10^8 \, M_\odot$, so only slightly larger than the $10^8 \, M_\odot$ system shown in Fig.5.2a. 1959+650 has an injected P_{rel} of $7 \times 10^{40} \, \mathrm{erg \, s^{-1}}$ (G10), corresponding in our scalings to $\dot{m} = 2.45 \times 10^{-4}$. So it should have a similar Fermi flux to the red spectrum of Fig.5.2a, which corresponds to $\dot{m} = 3 \times 10^{-4}$, $M = 10^8 \, M_\odot$. The observed $\log(\nu L(\nu))$ Fermi flux of 1959+650 at

Figure 5.2: a). BL Lac model SEDs for fixed black hole mass and increasing accretion rate ($\dot{m} = 10^{-4}$ (magenta), 3×10^{-4} (red), 10^{-3} (green), 3×10^{-3} (blue) and 10^{-2} (black), $M_{BH} = 10^8 \, M_\odot$). b). Corresponding steady state electron distributions. c). Fermi flux as a function of accretion rate, using mass and distance of model spectrum.

10^{23} Hz is 43.5 Hz erg s^{-1}, which is consistent with our red spectrum.

The changing shape of the emitted spectrum with accretion rate is due to the decrease in seed photons for electron cooling at lower $\dot{m}$, as shown explicitly by the corresponding self-consistent electron distributions in Fig.5.2b. The Lorentz factor of electrons which can cool in a light crossing time is $\gamma_{cool} \propto 1/(R_{diss}U_{seed}) \propto (1/M)(M/\dot{m}) \propto 1/\dot{m}$. The lowest mass accretion rate ($\dot{m} = 10^{-4}$, magenta) shows cooling only for the highest Lorentz factors, with $\gamma_{cool} \sim 10^{4.5}$. Below this the shape of the electron distribution is the same as the injected distribution, with a smooth break at $\gamma_b \sim 10^3$. As $\dot{m}$ increases, the sharp break at γ_{cool} moves to lower Lorentz factors. For the black electron distribution corresponding to $\dot{m} = 10^{-2}$, γ_{cool} is comparable to γ_b. This is clear from the black spectrum in Fig.5.2a, where the spectral peak is now produced by the cooled electron distribution above γ_{cool}.

Increasing cooling, as a result of increasing accretion rate, therefore provides a natural explanation for the existence of high frequency peaked (HBL) and low frequency peaked (LBL) BL Lacs (Ghisellini & Tavecchio, 2008b; Ghisellini & Tavecchio, 2009). In the context of our model, HBLs correspond to black holes with very low accretion rates. There is very little cooling and the bulk of the synchrotron emission is produced by electrons with Lorentz factors close to γ_{max}. Their electron distributions most closely resemble the original injected distributions. LBLs correspond to black holes with higher accretion rates, where cooling becomes increasingly important and the bulk of the energy is produced by electrons close to γ_{cool}.

Since HBLs are at lower accretion rates they are intrinsically fainter and so should be observed at lower redshifts than LBLs. This is indeed observed (Shaw et al., 2013). Fermi sensitivity is also a strong function of spectral index, decreasing with spectral hardness (Nolan et al., 2012). Since LBLs have softer spectra this suggests Fermi will preferentially select LBLs over HBLs due to spectral shape as well as flux.

Fig.5.2a also shows that as $\dot{m}$ increases, the ratio of the Compton to synchrotron luminosities changes. With our scalings, $L_{sync} \propto R_{diss}^3 U_B K$ and $L_{comp} \propto R_{diss}^3 U_{sync}K \propto R_{diss}^3(R_{diss}U_B K)K$, i.e. $L_{sync}/L_{comp} \propto 1/(R_{diss}K)$, where K is the normalisation of the steady state electron distribution. If there is complete cooling,

i.e. $\gamma_{cool} < \gamma_{min}$, then $K \propto Q_0/U_{seed} \propto (\dot{m}/M^2)(\dot{m}/M)^{-1} \propto 1/M$ which is independent of accretion rate. However the BL Lac spectra do not show complete cooling (see Fig.5.2b). If there is no cooling, $K \sim R_{diss}Q_0/c \propto \dot{m}/M$. The BL Lac spectra lie in this regime where the cooling is incomplete, hence $L_{sync}/L_{comp} \propto 1/\dot{m}$. This can be seen in Fig.5.2b, where the normalisation of the electron distribution at $\gamma = 1$ increases with $\dot{m}$. The scaling is not exactly $K \propto \dot{m}$, since there is an additional dependence on $\dot{m}$ introduced by γ_{cool} decreasing through the intermediate regime.

Fig.5.2c shows how the flux in the Fermi band drops with accretion rate. For higher accretion rates, cooling is efficient, so $L_{comp} \propto \dot{m}$. For low accretion rates, cooling is inefficient so $L_{comp} \propto \dot{m}^3$.

However, a more detailed comparison of Fig.5.2a to the data in the 'blazar sequence' shows evidence that the Compton flux changes more slowly with decreasing mass accretion rate due to an increase in the maximum Lorentz factor of the accelerated electron distribution (Ghisellini & Tavecchio, 2008b; Ghisellini & Tavecchio, 2009). Again, this can be a consequence of the different cooling environment, where electrons are accelerated to a maximum energy which is set by a balance between the acceleration timescale and the cooling timescale. Thus the accelerated electron distribution may itself change with cooling, such that $\gamma_b \propto \gamma_{max} \propto 1/U_{seed}$. We will consider such models later in the chapter.

5.5 BL Lac Visibility

The visibility of a BL Lac is strongly affected by viewing angle. Fig.5.3a shows how sharply the observed luminosity decreases for our assumed $\Gamma = 15$ with increasing viewing angle, where θ is measured in radians from the jet axis. Thus there is a difference of 10^{12} between the observed flux from a face on jet compared to an edge on jet.

The more distant the source, the more closely aligned to our line of sight the jet must be in order to boost the observed flux to a visible level. We define a flux limit of $F_{100\text{MeV}-100\text{GeV}} > 5 \times 10^{-12}\,\text{erg cm}^{-2}\,\text{s}^{-1}$ from the Fermi 2-year catalogue (Nolan et al., 2012), and show in Fig.5.3b the limiting redshift, z_{limit}, at which a

Figure 5.3: a). Decrease in observed flux with increasing viewing angle, where θ is measured in radians from the jet axis, for $\Gamma = 15$. b). Redshift limits for Fermi visible BL Lacs as a function of black hole mass, for increasing viewing angle and $\dot{m} = 10^{-2}$. c). Redshift limits for Fermi visible BL Lacs as a function of accretion rate, for $M_{BH} = 10^7$ (red), 10^8 (magenta), 10^9 (blue) and $10^{10}\ M_{\odot}$ (black) and $\theta = 0$.

BL Lac at $\dot{m} = 10^{-2}$ with different masses can be detected by Fermi at different inclination angles. We include the effects of absorption from pair production on the extragalactic IR background using the model of Kneiske & Dole (2010), though this is negligible. Only the most massive black holes ($\sim 10^{10}\,M_\odot$) which are most closely aligned to our line of sight can be seen out beyond $z = 4$. z_{limit} drops by a factor of ~ 3 if the inclination angle is increased from 0 to the more statistically likely $1/\Gamma$. This represents a change of just $\sim 4°$ for $\Gamma = 15$ used in our calculations. For a typical BL Lac mass of $10^9\,M_\odot$ viewed at $1/\Gamma$ the maximum observable redshift is $z \sim 1$, increasing to 4 only for the most face on jets.

Fig.5.3c shows how the redshift limit drops as a function of accretion rate for each mass (10^{10} (black), 10^9 (blue), 10^8 (magenta) and $10^7\,M_\odot$ (red)) black hole for $\theta = 0$. If LBLs correspond to BL Lacs at $\dot{m} \sim 10^{-2}$ and HBLs at $\dot{m} < 10^{-3}$ this shows how the redshift limits for the two populations should differ, with the majority of HBLs being observed below $z = 3$. Shaw et al. (2013) find this to be the case, with the distribution of LBLs extending to higher z, although they find the means of both populations are well below $z = 3$.

5.6   Predicted BL Lac Population from Cosmological Simulations

Cosmological simulations predict the number of supermassive black holes accreting at different redshifts, together with their masses and accretion rates. These simulations have been found to agree well with the observed number densities of broad line and narrow line AGN in the local universe (Fanidakis et al., 2011; Fanidakis et al., 2012). Combining our spectral code with the black hole data from these simulations allows us to predict the number of AGN that should be detected as BL Lacs by Fermi.

We combine our code with the black hole number densities predicted by the Millennium Simulation (Springel et al., 2005; Fanidakis et al., 2011; Fanidakis et al., 2012), binned as a function of both mass and mass accretion rate. We define a luminosity density from the number density multiplied by the luminosity at that

Figure 5.4: Predicted mass and accretion rate distribution of accreting black holes at increasing redshift from Millennium simulation. Colours trace luminosity density, with red showing the mass and accretion rates at which the maximum accretion luminosity is emitted at each redshift.

mass and mass accretion rate, i.e. $L = \eta \dot{M} c^2$ for the thin disc regime $10^{-2} < \dot{m} < 1$, joining smoothly onto a radiatively inefficient regime at lower $\dot{m}$ where $L \propto \dot{m}^2$ (Narayan & Yi, 1995) and onto a super-Eddington flow at higher $\dot{m}$ where $L \propto \ln(1 + \dot{m})$ (Shakura & Sunyaev, 1973). The luminosity density in each $(z, M, \dot{m})$ bin therefore depends on the mass, accretion rate, spin (which sets η), the inferred accretion regime and the number of black holes in that bin.

Fig.5.4 shows the evolution of the luminosity density of accretion power across cosmic time showing the features described by Fanidakis et al. (2011) and Fanidakis et al. (2012). At high redshift there is plenty of gas to fuel accretion. The black holes accrete close to the Eddington limit and grow rapidly. Comparing the snapshots for $z \sim 9$ and $z \sim 5$, the typical black hole mass producing the bulk of the accretion luminosity increases from $\sim 10^6$ to $\sim 10^8 \, M_\odot$. As the black holes gradually run out of gas, their accretion rates drop (compare $z \sim 2$ and $z \sim 1$). By redshift 2, accretion rates are beginning to drop below $\dot{m} = 10^{-2}$, into the regime at which BL Lac type jets should be produced. This suggests no BL Lacs should be observed much above $z \sim 2$, not just because the flux becomes too faint, but because the typical accretion rate is too high for the production of BL Lac jets.

We initially assume all black holes accreting inefficiently will produce a BL Lac type jet. We can then calculate the number of AGN hosting a BL Lac type jet in each $(z, M, \dot{m})$ bin by using the number densities of black holes in the radiatively inefficient regime ($\dot{m} < 10^{-2}$). If this number is less than 1 we use Poisson statistics to randomly determine whether a black hole is present or not. Each black hole in each $(z, M, \dot{m})$ bin is then assigned a random distance within this redshift bin and random θ_{obs}, assuming $\cos \theta_{obs}$ is distributed uniformly. We then calculate the observed spectrum to determine whether or not the jet would be visible to Fermi.

Fig.5.5a shows the predicted redshift distribution of Fermi visible BL Lacs (black). The predicted distribution peaks at $z \sim 0.5$ and drops gradually to $z \sim 2$. No BL Lacs are observed above this point, not because they are not visible (see Fig.5.3b), but because there simply are not enough SMBHs accreting below 10^{-2} in the cosmological simulations to produce SSC jets, due to the higher activity expected at earlier times. The low redshift distribution of BL Lacs is a direct result of cosmic

Figure 5.5: a). Predicted redshift distribution of Fermi visible BL Lacs, assuming black holes of all spins accreting below $\dot{m} = 10^{-2}$ produce a BL Lac type jet (black). Red solid line shows observed redshift distribution of Fermi detected BL Lacs. Red dashed line shows observed redshift distribution $\times 100$. b). Predicted mass distribution of Fermi visible BL Lacs. c). Predicted accretion rate distribution of Fermi visible BL Lacs.

downsizing and the requirement of an $\dot{m} < 10^{-2}$ to produce a SSC jet.

However, comparing this to observations shows a huge discrepancy (red solid line from Shaw et al., 2013, which almost merges with the x-axis at this scale). Our expected number of ~ 100000 BL Lacs dramatically overpredicts the observed number of Fermi detections (~ 500). A clear illustration of the problem can be seen from simply the number density of massive ($8 < \log M < 9$) black hole accretion flows with $10^{-3} < \dot{m} < 10^{-2}$ in the cosmological simulations in the redshift bin centred around $z \sim 0.5$ (Fanidakis et al., 2011). This number is $6.8 \times 10^{-4}\,\mathrm{Mpc}^{-3}$ and the volume of this bin, from $0.509 < z < 0.564$ is $\sim 9\,\mathrm{Gpc}^3$ so this gives 6101659 objects which should host similar jets to 1749+096 (Fig.5.1), i.e. have Fermi flux of $10^{-10.5}(0.5/0.322)^{-2} \sim 10^{-11}\,\mathrm{erg\ s}^{-1}\,\mathrm{cm}^{-2}$ if viewed at the same angle (roughly $1/\Gamma$). The probability that we see the source within this angle is $1 - \cos(1/\Gamma)$, so the expected number of Fermi detections of these sources alone is ≈ 13554, similar to the full calculation results. The large discrepancy clearly points to a fundamental breakdown of one of the assumptions.

On a more subtle level, the shape of the redshift distribution for the Fermi predictions is also mismatched to the observations. The red dashed line shows the observed number of BL Lacs scaled by a factor of 100 so it can be compared to the predicted distribution. We define redshift from Shaw et al. (2013) as either spectroscopic redshift, spectroscopic lower limit, the mean of their redshifts derived from host galaxy fitting, or their redshift upper limits, in that order of preference. The dashed line shows this observed redshift distribution ($\times 100$). It is clear that not only is the total predicted number wrong, but we are also over estimating the proportion of BL Lacs in the range $z = 0.5 - 2$.

The mass distribution and mass accretion rate distributions are as expected (Fig.5.5b and c), with higher luminosity SSC flows (i.e. higher $\dot{m}M$) being more likely to be observed, so simple energetics selects the highest mass and mass accretion rate objects, so the typical predicted mass of a Fermi visible BL Lac of $\sim 10^{8.5} - 10^9\,M_\odot$ is a combination of three factors:

1. Very few black holes accreting at $\dot{m} < 10^{-2}$ above $z = 2$.

2. Most black holes at $z < 2$ accreting with $\dot{m} < 10^{-2}$ have $M > 10^8$.

3. Black holes with $M > 10^9$ are increasingly rare in the local universe so we are less likely to observe one favourably orientated to our line of sight.

5.7   Another Factor Affecting Jet Scaling?

Our results predict $\sim 10^5$ BL Lacs should have been detected in the Fermi 2-year catalogue. In contrast, ~ 500 objects in the 2nd Fermi LAT catalogue are classed as BL Lacs. Even allowing for Galactic centre emission limiting sky coverage ($|b| > 10°$ means that only 80% of the sky is included), this is still 3 orders of magnitude larger than observed. Clearly jets do not simply scale with accretion power.

We assumed the injected electron distribution was independent of mass accretion rate. This may not be the case. Ghisellini & Tavecchio (2009) approximately use $\gamma_{max} \propto \gamma_b \propto 1/\dot{m}$ to fit their blazar sequence. Acceleration of electrons is affected by the ambient photon field which depends on the amount of cooling and ultimately on accretion rate. In the efficient cooling regime, $L_{sync} \propto L_{comp} \propto \dot{m}M$, making $\gamma_{max} \propto 1/\dot{m}$ not unreasonable. However, an increase of γ_{max}, and γ_b in particular, only serves to increase the Fermi band luminosity for lower $\dot{m}$ systems, and increase the discrepancy between the predictions and observations.

The discrepancy could instead be explained if every black hole accreting below 10^{-2} has the potential to produce a BL Lac type jet, but only does so 1/1000th of the time. This seems unlikely, since Fanaroff Riley Type I (FRI) AGN, the misaligned versions of BL Lacs (Padovani & Urry, 1990, Padovani & Urry, 1991, Urry et al., 1991), show large scale extended radio jets. This suggests these jets are persistent, analogous to the steady low/hard state jet seen in BHBs at low $\dot{m}$, not transient events.

Another possibility is that the jet depends on magnetic flux being advected onto the black hole from the extremely large scale hot halo gas around the galaxy. Sikora & Begelman (2013) suggest that if there is magnetic flux in this gas, then it could be dragged down close to the black hole by cold gas from a merging spiral galaxy. However, this does not address the fundamental question as to where the magnetic

flux in the halo gas comes from, and using cold gas from a spiral merger to drag this field down to the black hole is unlikely to be applicable in the BL Lacs, as they have low ongoing mass accretion rates.

The bulk Lorentz factor of the jet is the biggest factor affecting its visibility. We rerun our calculations with a reduced $\Gamma = 10$ instead of 15 and find this roughly halves the predicted number of BL Lacs, but still wildly overpredicts the observations.

We have assumed all jets are produced with the same value of bulk Lorentz factor but this is clearly not the case — BHBs at low $\dot{m}$ have $\Gamma \sim 1.2$ (Fender et al., 2004). The most obvious way to reduce the number of visible BL Lacs is to allow a distribution of Γ. Yet there must be some physical parameter which controls the jet acceleration. The acceleration region, where the magnetic (Poynting) flux of the jet is converted to kinetic energy, is very close to the black hole, so it seems most likely that this is set by the black hole itself, in which case black hole spin is the only remaining plausible parameter. A potential explanation for the lower number of observed BL Lacs is that if only black holes with the highest spin produce highly relativistic jets, and high spin is rare.

The cosmological simulations include the growth of SMBH spin via accretion processes and black hole-black hole coalescence following galaxy mergers (Volonteri et al., 2005; Volonteri et al., 2007; Volonteri, 2012: Fanidakis et al., 2011; Fanidakis et al., 2012). The mass accumulated onto the central SMBH in an accretion event is tied in the simulation to a fixed fraction (0.5%) of the mass of gas in a star formation episode in the host galaxy. If this mass is all accreted in a single event (prolonged accretion) then this is sufficient to spin most black holes up to maximal (Volonteri et al., 2005; Volonteri et al., 2007). However, the mass accreting onto the central black hole in any single event may be limited by self-gravity. This splits the accreting material up into multiple smaller events, each of which can be randomly aligned since the star formation scale-height is large compared to the black hole sphere of influence even in a disc galaxy (King et al., 2008). Such chaotic accretion flow models result in predominantly low spin black holes (Volonteri et al., 2007: King et al., 2008: Fanidakis et al., 2011; Fanidakis et al., 2012) and high spins are

rare, as they are produced not via accretion but via black hole mergers (Fanidakis et al., 2011; Fanidakis et al., 2012).

We use the spin distribution from the chaotic accretion flow model simulations, and introduce a spin cut to our results, so that only black holes with spin greater than $a_{cut} \sim 0.8$ produce a BL Lac type jet. This reduces the predicted number of Fermi visible BL Lacs to ~ 900. Fig.5.6a shows the resulting redshift distribution together with the observed distribution. Not only does this reduce the discrepancy between predicted and observed total numbers, it also gives a better match to the shape of the distribution. Limiting production of BL Lac type jets to black holes with high spin causes the redshift distribution to peak slightly later and drop off more sharply above $z = 0.5$. This is because high spins arise from black hole mergers. Production of BL Lac jets is already limited to black holes accreting below 10^{-2}, i.e. $M > 10^8\,M_\odot$. The black hole mergers which make the most massive black holes, occur at the latest times.

Fig.5.6b and c show how this affects the predicted mass and accretion rate distributions of Fermi visible BL Lacs. The scaled down distributions including black holes of all spins are shown by the dashed lines for comparison. Requiring high spin increases the peak of the mass distribution to $\sim 10^{8.5} - 10^{9.5}\,M_\odot$, because it is the most massive black holes that are former by mergers and are consequently more likely to have high spin. The peak of the accretion rate distribution is actually slightly reduced. This is because the more massive black holes have lower accretion rates; the spin cut has excluded lower mass black holes with lower spins which tend to have slightly higher accretion rates.

The low spin, low accretion rate black holes, which generally have smaller masses $(10^7 - 10^8\,M_\odot)$ correspond to the LINERs, which are not observed to have jets as relativistic as those in BL Lacs. If they are low spin, as expected, and high spin is required for a highly relativistic jet, then this naturally explains why LINERs are observed to have weaker radio jets.

Figure 5.6: a). Predicted redshift distribution of Fermi visible BL Lacs, assuming only black holes with spins $a > 0.8$ produce a BL Lac type jet (blue). Red solid line shows observed redshift distribution of Fermi detected BL Lacs. b). Predicted mass distribution including only black holes with $a > 0.8$ (blue). Black dashed line shows predicted distribution including all black holes $\times 0.01$. c). Predicted accretion rate distribution including only black holes with $a > 0.8$ (blue). Black dashed line shows predicted distribution including all black holes $\times 0.01$.

5.8   Implications of Scaling Jet Power with Spin: FRI Sources

The parent population of BL Lacs is probably the Fanaroff Riley Type I (FRI) sources (e.g. the review by Urry & Padovani, 1995). These show 'fluffy' radio jets whose surface brightness decreases with distance from the central source, contrasting with the classic lobe and hotspot radio emission seen in the more powerful FRII sources which are the parent population of the FSRQ (Padovani & Urry, 1992). Thus the FRI sources should also correspond to high spin black holes, and indeed are similarly powered by high mass SMBHs (Woo & Urry, 2002).

However, we might then expect some difference in jet radio emission between the FRIs and lower mass LINERs as the cosmological simulations predict that the lower mass SMBHs have lower spin (Fanidakis et al., 2011; Fanidakis et al., 2012). Sikora et al. (2007) claim that this difference is indeed seen, with radio emission being ~ 3 orders of magnitude higher in the FRIs.

However, some of this difference disappears when only core radio luminosity (rather than core plus lobes) is used (Broderick & Fender, 2011). This is clearly an issue as the extended radio emission also depends on environment. The black holes in FRIs are more massive than those in LINERs, hence live in richer cluster environments, with larger dark matter halos which trap more hot cluster gas. The jet then emerges into a denser, higher pressure environment, which means that a much larger fraction (potentially all) of the jet kinetic energy is converted to radiation and/or heating of the cluster gas (Bîrzan et al., 2004). Conversely, any jet from the lower mass LINERs emerges into a poorer group environment, so adiabatic losses can predominate and the extended radio emission is much smaller (e.g. Krause et al., 2012).

Some of the remaining difference in core radio power is *expected* due to the difference in mass (Broderick & Fender, 2011). However, even accounting for this there is still a factor of ~ 10 in mass corrected, core radio emission. The LINERs lie on the Fundamental plane (Merloni et al., 2003; Falcke et al., 2004), i.e. have the expected core radio emission for their black hole mass and mass accretion rate, so the

FRIs are a factor 10 brighter in mass corrected, core radio emission than expected from the same jet models which produce the low bulk Lorentz factor BHB and LINER jets, consistent with the idea that the jet is intrinsically more powerful/has higher Lorentz factor due to black hole spin.

It is difficult to predict the difference in core radio emission with black hole spin in our models as synchrotron self-absorption means that the observed radio emission does not arise in the same region as produces the Fermi flux. It may be produced either at larger radii, perhaps where the jet has decelerated, or in a lower density, lower bulk Lorentz factor layer surrounding the $\Gamma = 15$ spine of the jet. Either of these could explain the lower Lorentz factor ($\Gamma \sim 2 - 10$) of the radio jet observed in FRIs (Chiaberge et al., 2000), though the spine-layer structure may additionally be able to explain the very fast variability timescales seen in some BL Lacs (Ghisellini & Tavecchio, 2008a).

5.9   Conclusions

We have taken a statistical approach to constrain the conditions necessary to produce the highly relativistic jets seen in BL Lac objects. We combine SMBH number densities from cosmological simulations, known to reproduce the optical luminosity function of AGN, with spectral models of jet emission and simple jet scaling functions which depend only on mass and accretion rate. The key assumption is that every black hole accreting with $\dot{m} < 10^{-2}$, i.e. in the radiatively inefficient accretion flow regime, should produce a BL Lac type jet.

Our calculation of the expected number of BL Lacs detectable by Fermi over-predicts the observations by three orders of magnitude. This clearly shows that our fundamental assumptions are incorrect, and that the jet power and properties do not scale simply with mass and mass accretion rate. The only other parameter which a black hole can have is spin. We can reproduce the observed numbers of BL Lacs if SMBHs grow predominantly via chaotic (randomly aligned) accretion episodes and BL Lac type jets are restricted to black holes with spin $a > 0.8$. These are rare as they form from black hole-black hole coalescence following a major merger

event which is not then overwhelmed by further chaotic accretion, i.e. this requires a gas poor major merger event and only the most massive galaxies, which host the most massive black holes, are gas poor in the local Universe (Fanidakis et al., 2011; Fanidakis et al., 2012).

A spin cut is in line with the longstanding speculation that these most relativistic jets require high spin black holes (Maraschi et al., 2012), and also gives a good match to the observed redshift distribution of BL Lacs which peaks at $z = 0.5$ and then drops off sharply, with no objects above $z \sim 2$. This is a consequence of three factors:

1. BL Lac jets are restricted to black holes with $\dot{m} < 10^{-2}$, and there are no black holes accreting at $\dot{m} < 10^{-2}$ above $z \sim 2$.

2. Only the most massive black holes have high spin through mergers, which happen at late times, causing the bulk of the population to fall below $z = 1$

3. These most massive objects are rare in the local universe causing the distribution to decrease again below $z = 0.5$.

Since FRI sources are consistent with being the misaligned analogs of BL Lacs, they should also have high spin. They are indeed offset from the Fundamental Plane, i.e. have higher (mass corrected) core radio emission than the lower mass and presumably lower spin LINERs (Broderick & Fender, 2011), however only by a factor ~ 10. This suggests the radio emission, which is not predominantly produced from the same region as the Fermi flux, may not be as sensitive to the difference in spin.

Chapter 6

What Powers the Most Relativistic Jets? II. Flat Spectrum Radio Quasars

Flat Spectrum Radio Quasars (FSRQs) are the high accretion rate counterparts of BL Lacs ($\dot{m} = \dot{M}/\dot{M}_{Edd} > 0.01$, where $\eta \dot{M}_{Edd} c^2 = L_{Edd}$). At high accretion rates the accretion flow around the black hole takes the form of a radiatively efficient accretion disc, which can often be seen dominating the optical–UV spectra of FSRQs, despite the strongly boosted jet emission (Ghisellini et al., 2010, hereafter G10). The strong UV disc emission illuminates material above the disc, which, at a particular radius set by the gas density and illuminating flux, re-emits the radiation in the form of broad emission lines (the 'broad line region', BLR). Both the direct accretion flow emission and BLR emission (and reprocessed emission from the torus) provide a source of external seed photons for FSRQ jets, in addition to the internal synchrotron radiation which predominates in BL Lacs. Crucially the radius at which the broad lines are produced is normally at a larger distance from the black hole than the region of the jet where the high energy emission is produced (Ghisellini & Tavecchio, 2009). As a result the jet electrons moving with the relativistic bulk motion of the jet see the stationary BLR seed photons strongly Doppler boosted. This enhances the luminosity of the external-Compton (EC) emission, causing the Compton humps of EC FSRQs to be much more luminous than those of synchrotron self-Compton

(SSC) BL Lacs.

In this chapter we extend the SSC model of the previous chapter to include external sources of seed photons and now use the number densities of black holes accreting with $\dot{m} > 0.01$ to try to predict the observed population of Fermi detected FSRQs. In addition, we test our assumed jet scalings using observations of gamma-ray loud Narrow Line Seyfert 1s (γNLS1s), which should be lower mass versions of FSRQs.

6.1 External-Compton Jets

We extend the SSC model of the previous chapter to include sources of external seed photons, since these are important in the higher accretion rate FSRQs, which have UV bright accretion discs. We include seed photons from the accretion disc and X-ray corona, emission from the BLR and torus, and reflection of coronal X-rays off the BLR.

We assume the distance of the BLR and infra-red (IR) torus (R_{BLR} and R_{IR}) from the central black hole scales with the accretion disc luminosity as (Ghisellini & Tavecchio, 2009):

$$R_{BLR} = 10^{17} (\frac{L_d}{10^{45}\,\mathrm{erg\,s^{-1}}})^{1/2}\,\mathrm{cm} \tag{6.1.1}$$

$$R_{IR} = 2.5 \times 10^{18} (\frac{L_d}{10^{45}\,\mathrm{erg\,s^{-1}}})^{1/2}\,\mathrm{cm} \tag{6.1.2}$$

And we calculate the external seed photon energy densities following the prescriptions in Ghisellini & Tavecchio (2009) (full calculation details are given in the Appendix). The total seed photon energy density $U'_{seed} = U'_B + g(\gamma)(U'_{sync} + U'_{ex})$ and now includes the magnetic energy density ($U'_B = B^2/8\pi$) and the fraction $g(\gamma)$ of the energy density of synchrotron (U'_{sync}) and external (U'_{ex}) seed photons which can be Compton upscattered by electrons of energy γ within the Klein-Nishina limit.

As before we assume a spherical emission region of radius R_{diss}, where R_{diss} is related to the distance of the emission region from the black hole by $Z_{diss} = \phi R_{diss}$,

ϕ is the half opening angle of the jet and we assume a conical jet. The electrons are accelerated into an intial distribution $(Q(\gamma))$, and cool by emitting self-absorbed synchrotron and synchrotron self-Compton radiation, and by upscattering seed photons from external sources of radiation, into the steady state distribution $N(\gamma)$. This jet frame emission is then boosted by the bulk motion of the jet, cosmologically redshifted and attenuated due to pair production on the extragalactic infra-red background light (though this is generally small for the Fermi bandpass) to produce the observed flux.

The parameters of our extended model are therefore:

- Parameters of the accretion flow: black hole mass and Eddington scaled accretion rate (M and $\dot{m}$), for calculating the density of external seed photons.

- Physical parameters of the jet: Γ, radius of emission region R_{diss} and half opening angle of the jet ϕ.

- The magnetic field of the emission region and power injected into relativistic electrons (B and P_{rel}).

- Parameters of the injected electron distribution: γ_{min}, γ_b, γ_{max}, n_1 and n_2.

We adopt the cosmology used in the Millennium simulations: $h = 0.72$, $\Omega_m = 0.25$, $\Omega_{vac} = 0.75$ (Springel et al., 2005; Fanidakis et al., 2011).

6.2 Standard Jet Scalings

As in the previous chapter, we assume that the acceleration mechanism is the same for all FSRQs, giving the same injected electron distribution, regardless of mass and accretion rate. We also assume all jets are produced with the same Γ and the same half opening angle (which we fix to $\phi = 0.1$). As before, this leaves three remaining parameters: R_{diss}, B and P_{rel}.

We initially assume that FSRQs follow the same standard jet scalings that BL Lacs appear to follow. Hence we scale $R_{diss} \propto M$, since all size-scales should scale with the mass of the black hole (Heinz & Sunyaev, 2003), and assume the jet power

is a constant fraction of the total accretion power, $P_j \propto \dot{m}M$. We again stress that this assumption is valid whether the jet is powered by the accretion flow or the spin energy of the black hole, since extraction of black hole spin energy relies on magnetic fields generated in the accretion flow, which will be affected by accretion rate. A constant fraction of the total jet power is then injected into relativistic particles and magnetic fields, giving $P_{rel} \propto \dot{m}M$ and $B \propto (\dot{m}/M)^{1/2}$ (Gardner & Done, 2014b).

We choose the mean FSRQ parameters from G10 to scale from, which are the logarithmic average values from their sample of 53 Fermi detected FSRQs. This gives $M_0 = 1 \times 10^9 \, M_\odot$, $\dot{m} = 0.1$, $R_{diss,0} = 1.89 \times 10^{16}$ cm, $B_0 = 2.6$ G, $P_{rel,0} = 2 \times 10^{43}$ erg s^{-1}, $\Gamma = 13$, $\gamma_{min} = 1$, $\gamma_b = 300$, $\gamma_{max} = 3 \times 10^3$, $n_1 = 1$, $n_2 = 2.7$, and we scale R_{diss}, P_{rel} and B as:

$$R_{diss} = R_{diss,0}\frac{M}{M_0} \tag{6.2.3}$$

$$P_{rel} = P_{rel,0}\frac{\dot{m}}{\dot{m}_0}\frac{M}{M_0} \tag{6.2.4}$$

$$B = B_0\left(\frac{\dot{m}}{\dot{m}_0}\frac{M_0}{M}\right)^{1/2} \tag{6.2.5}$$

The distance to the BLR and IR torus are both $\propto L_d^{1/2}$. This implies R_{BLR} and R_{IR} should also scale with the mass and accretion rate of the black hole, since $L_d \propto \dot{m}M$. Hence $R_{BLR} \propto R_{IR} \propto (\dot{m}M)^{1/2}$.

6.2.1 Spectral Changes with Mass

Fig.6.1a shows a sequence of FSRQ spectra with increasing black hole mass. The accretion rate is fixed to $\dot{m} = 0.1$ and R_{diss}, B and P_{rel} are scaled as described above.

As black hole mass decreases, so does the size of the emission region, since $R_{diss} \propto M$. This can be seen in the increase in synchrotron self-absorption frequency, from $\sim 10^{10.5}$ (black spectrum) to $\sim 10^{11.5}$ Hz (magenta spectrum). The total luminosity also decreases, since the power injected into relativistic electrons decreases with decreasing mass ($P_{rel} \propto M$).
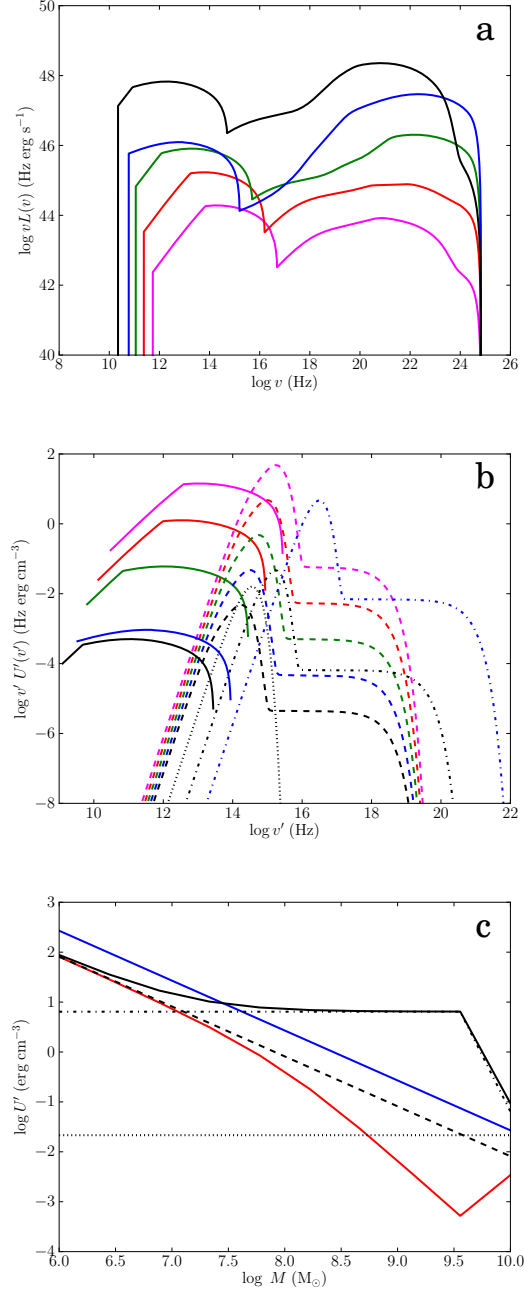
Figure 6.1: Spectral changes with mass using standard jet scalings. a). FSRQ model SEDs for fixed accretion rate and increasing black hole mass ($M_{BH} = 10^6$ (magenta), 10^7 (red), 10^8 (green), 10^9 (blue) and $10^{10} \, M_\odot$ (black), $\dot{m} = 0.1$). b). Corresponding seed photon energy density spectra as seen in the jet frame. Solid lines show synchrotron seed photons, dashed lines show accretion disc plus coronal seed photons, dot-dashed lines show seed photons from the BLR plus coronal flux reflected by the BLR, and dotted lines show seed photons from the torus. The seed photon energy density from the torus is the same for all masses. The seed photon energy density from the BLR is the same for all masses except $10^{10} \, M_\odot$, where $Z_{diss} > R_{BLR}$. c). Seed photon energy densities in the jet frame as a function of black hole mass. Blue line shows U'_B, red line shows U'_{sync}, black lines show energy densities of external seed photons: dashed line shows $U'_d + U'_X$, dot-dashed line $U'_{BLR} + U'_{XBLR}$, dotted line U'_{IR}, and solid line shows total U'_{ex}.

The relative strengths of the synchrotron and Compton humps also changes with mass. The blue spectrum corresponds to the mean FSRQ model of G10, with $\dot{m} = 0.1$, $M = 10^9 M_\odot$. It shows a strong Compton hump at 10^{22} Hz due to Compton up-scattering of external seed photons, predominantly from the BLR. The low energy synchrotron hump is roughly an order of magnitude less luminous ($\sim 10^{46}$ erg s^{-1}). As the black hole mass drops from 10^9 (blue) - $10^6 M_\odot$ (magenta), the relative luminosity of the Compton hump decreases until at the lowest masses the two humps show comparable luminosity. The relative strength of the two humps depends on the relative strength of the energy density in magnetic fields compared to the energy density of external seed photons.

Fig.6.1b shows the spectral energy density of seed photons in the jet frame. As mass drops so does the emission region size and hence its distance from the black hole, since $Z_{diss} = R_{diss}/\phi \propto M$. Smaller R_{diss} and Z_{diss} increases the energy density of accretion disc seed photons, despite the drop in L_d with M, showing an increase of ~ 3 orders of magnitude (blue dashed line to magenta dashed line). However the dominant source of seed photons is U'_{BLR} and this stays constant, since $U'_{BLR} \propto L_d/R^2_{BLR} \propto L_d/(L_d^{1/2})^2 = const$ for $Z_{diss} < R_{BLR}$ (blue dot-dashed line). In contrast the magnetic field, which determines the amount of synchrotron emission, increases as black hole mass decreases, since $B \propto M^{-1/2}$. As a result U'_{sync} (solid lines) increases by more than 4 orders of magnitude, becoming comparable to U'_{BLR} at the lowest masses. As a result the Compton humps of the lowest mass spectra are dominated by up-scattering of synchrotron radiation causing them to look more like low accretion rate SSC BL Lacs than FSRQs, despite their higher accretion rates.

The lack of external seed photons means less efficient cooling in lower mass objects. For the $10^9 M_\odot$ spectrum (Fig.6.1a, blue line), the cooling is almost complete with $\gamma_{cool} = 7$, where γ_{cool} is the minimum Lorentz factor of electrons that can cool in one light crossing time. For the $10^6 M_\odot$ spectrum (magenta), γ_{cool} has increased to 106, resulting in a clear spectral break at $\sim 10^{13.5}$ Hz in the synchrotron emission. The decreasing frequency of this spectral break tracks the decrease in γ_{cool} and increase in cooling from $10^6 - 10^9 M_\odot$.

Above $10^9 M_\odot$, γ_{cool} increases again (Fig.6.1a, black spectrum, $\gamma_{cool} = 39$). This

is because for a $10^{10} M_\odot$ black hole the emission region has gone beyond R_{BLR}, since $Z_{diss} \propto M$ while R_{BLR} (and R_{IR}) $\propto M^{1/2}$. This causes U'_{BLR} to drop dramatically (black dot-dashed line in Fig.6.1c), reducing the amount of cooling. The next strongest source of seed photons is the torus (Fig.6.1b, black dotted line). U'_{IR} is constant for all masses since like U'_{BLR}, $U'_{IR} \propto L_d/R_{IR}^2 \propto L_d/(L_d^{1/2})^2 = const$ for $Z_{diss} < R_{IR}$, which is the case for all five masses. Consequently, above $10^9 M_\odot$ the ratio between synchrotron and Compton peaks drops again.

Fig.6.1c shows the total energy densities of seed photons in the jet frame as a function of black hole mass. This shows clearly for masses around $10^9 M_\odot$, where the energy density of BLR seed photons dominates (black dot-dashed line), the energy density of synchrotron radiation is suppressed (red line) due to the strong cooling. U'_{sync} recovers at higher masses, as $Z_{diss} > R_{BLR}$, and dominates over U'_{BLR} at low masses ($< 10^7 M_\odot$), where the magnetic field is strongest (blue line).

6.2.2 Spectral Changes with Accretion Rate

Fig.6.2a shows a sequence of FSRQ spectra with increasing accretion rate. We fix $M = 10^9 M_\odot$ and increase the accretion rate from $\log \dot{m} = -2$ to 0.5.

As accretion rate increases, the synchrotron self absorption frequency increases from $\sim 10^{10}$ (magenta) to 10^{12} Hz (blue). This is because the size of the emission region stays constant (R_{diss} does not depend on $\dot{m}$) while the magnetic field is increasing ($B \propto \dot{m}^{1/2}$).

As accretion rate increases the total luminosity also increases since $P_{rel} \propto \dot{m}$ and $B \propto \dot{m}^{1/2}$. However the synchrotron emission increases faster than the Compton emission so that for the highest accretion rate ($\log \dot{m} = 0.5$) the two peaks show comparable luminosity (blue spectrum) while at $\log \dot{m} = -1.5$ (red) the Compton peak is ~ 2 orders of magnitude brighter than the synchrotron peak.

This is because the increase in synchrotron emission comes from both the increase in P_{rel} and the increase in its seed photons from the magnetic field. In contrast the main source of seed photons for the Compton hump is the BLR and the energy density of these remains constant while $Z_{diss} < R_{BLR}$. Hence most of the increase in L_{comp} is due to P_{rel}. Only for the highest accretion rates (blue and cyan spectra) do
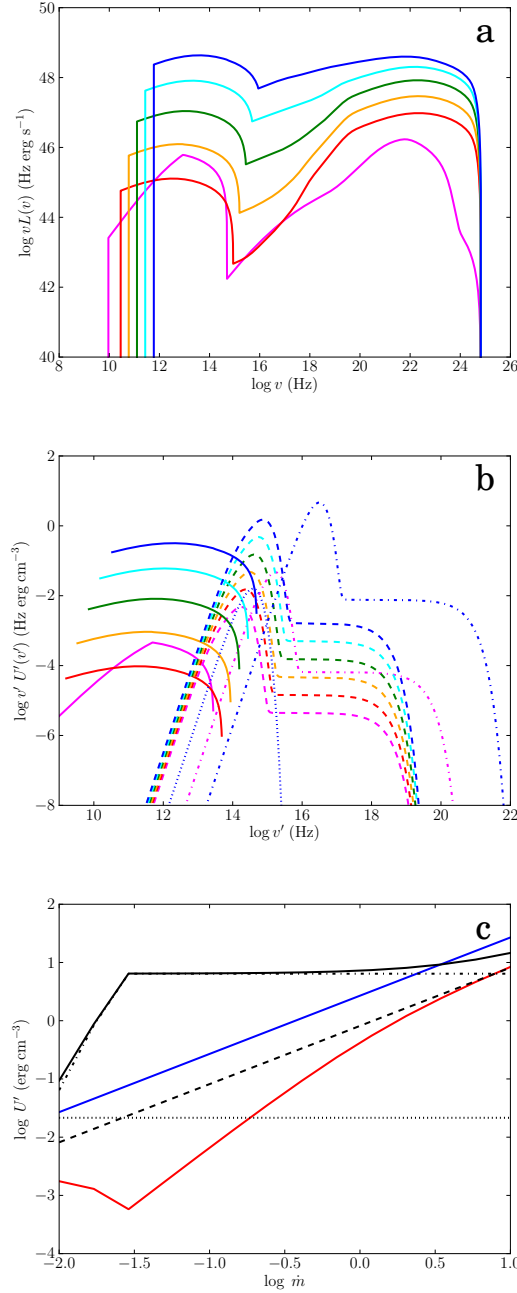
Figure 6.2: Spectral changes with accretion rate using standard jet scalings. a). FSRQ model SEDs for fixed black hole mass and increasing accretion rate ($\log \dot{m} = -2$ (magenta), -1.5 (red), -1 (orange), -0.5 (green), 0 (cyan) and 0.5 (blue), $M_{BH} = 10^9 \, M_\odot$). b). Corresponding seed photon energy density spectra as seen in the jet frame. Solid lines show synchrotron seed photons, dashed lines show accretion disc plus coronal seed photons, dot-dashed lines show seed photons from the BLR plus coronal flux reflected by the BLR, and dotted lines show seed photons from the torus. The seed photon energy density from the torus is the same for all accretion rates. The seed photon energy density from the BLR is the same for all accretion rates except $\log \dot{m} = -2$, where $Z_{diss} > R_{BLR}$. c). Seed photon energy densities in the jet frame as a function of accretion rate. Blue line shows U'_B, red line shows U'_{sync}, black lines show energy densities of external seed photons: dashed line shows $U'_d + U'_X$, dot-dashed line $U'_{BLR} + U'_{XBLR}$, dotted line U'_{IR}, and solid line shows total U'_{ex}.

the other sources of seed photons (U'_{sync} and U'_{acc}, solid and dashed lines in Fig.6.2b) become comparable with U'_{BLR}. These are much lower energy photons than the blue shifted BLR emission (Fig.6.2b). Consequently the Compton hump at the highest accretion rates is much broader as well as being more similar in luminosity to the synchrotron peak. Again this gives a spectral shape much more typical of low accretion rate BL Lacs, except now it is the result of an extremely high $\dot{m}$ causing U'_{sync} to dominate over U'_{ex} rather than an extremely low $\dot{m}$ where U'_{ex} is absent.

The red spectrum in Fig.6.2a ($\log \dot{m} = -1.5$) shows the greatest luminosity difference between synchrotron and Compton peaks. The difference lessens again for the lowest accretion rate spectrum (magenta, $\log \dot{m} = -2$). This is because $R_{BLR} \propto L_d^{1/2} \propto \dot{m}^{1/2}$. For a $10^9 M_\odot$ black hole at $\dot{m} = 10^{-2}$, the BLR radius has shrunk so much that it is now less than the distance to the jet emission region. Once the BLR is behind the jet emission region its seed photons are de-boosted and U'_{BLR} drops significantly (compare magenta and blue dot-dashed lines in Fig.6.2b). The amount of cooling drops, shown by the appearance of a cooling break at 10^{13} Hz in the magenta spectrum. Synchrotron and accretion flow seed photons become more important, broadening the Compton hump again. But even these cannot help for long; $\dot{m} = 10^{-2}$ is the rate at which accretion flows make the transition from radiatively efficient to inefficient. Below $\dot{m} = 10^{-2}$ UV bright accretion discs can no longer be sustained and give way to ADAF-like flows. This severely reduces the available accretion flow seed photons and effectively switches off the BLR since there are no UV photons to illuminate it. This final magenta spectrum represents the transition from dimming FSRQ to a low accretion rate SSC BL Lac.

Fig.6.2c shows the total energy densities of seed photons in the jet frame as a function of accretion rate. For $-1.5 < \log \dot{m} < -0.5$, U'_{BLR} dominates, suppressing U'_{sync} and giving the luminous Compton hump and much smaller synchrotron peak typical of FSRQs. Only at the extremes of accretion rate does U'_{sync} dominate. At super Eddington accretion rates U'_{sync} and U'_{acc} start to overtake U'_{BLR} producing a pseudo-BL Lac type spectrum but with extremely high luminosity. And at the lowest accretion rates U'_{sync} recovers when R_{BLR} has shrunk below Z_{diss} and the object is making the transition to a low accretion rate SSC BL Lac.

Figure 6.3: Redshift limits using standard jet scalings. a). Redshift limits for Fermi visible FSRQs as a function of black hole mass, for increasing viewing angle ($\theta = 0$ (black), $1/\Gamma$ (blue) and $1/2\Gamma$ (magenta), where $\Gamma = 13$) and $\dot{m} = 0.1$. b). Redshift limits for Fermi visible FSRQs as a function of accretion rate, for $M_{BH} = 10^7$ (red), 10^8 (magenta), 10^9 (blue) and $10^{10} \, M_{\odot}$ (black) and $\theta = 0$.

6.2.3   FSRQ Visibility

Having shown how the spectrum of a FSRQ might change with mass and accretion rate, we now investigate the redshift limits at which FSRQs of different masses and accretion rates should be visible to Fermi, according to these standard jet scalings. We define a flux limit of $F_{1\mathrm{GeV}-100\mathrm{GeV}} > 5 \times 10^{-10}$ photons cm^{-2} s^{-1} from the Fermi 1-year catalogue (Abdo et al., 2010). If a FSRQ of a given mass and accretion rate has $F > F_{limit}$ in the Fermi band we assume it will be detected.

Fig.6.3a shows the redshift limits for Fermi visible FSRQs as a function of black hole mass. We fix $\dot{m} = 0.1$ and show three different inclination angles: $\theta = 0$ (black), $1/\Gamma$ (blue) and $1/2\Gamma$ (magenta). Clearly more closely aligned FSRQs are seen out to higher redshifts. The limiting redshift increases with mass, since L_{comp} increases with mass (see Fig.6.1a), until $\sim 10^{9.5} \, M_{\odot}$. A highly aligned FSRQ with a $10^{9.5} \, M_{\odot}$ black hole can be detected out beyond $z = 6$. However above $10^{9.5} \, M_{\odot}$, the redshift limits drop sharply to $z \leqslant 2$ for a $10^{10} \, M_{\odot}$ black hole. The reason for this can be seen in Fig.6.1a. For the most massive $10^{10} \, M_{\odot}$ black holes, $Z_{diss} > R_{BLR}$, because Z_{diss} grows $\propto M$ while $R_{BLR} \propto M^{1/2}$. BLR photons are still the dominant source of seed photons, however they are now behind the emission region. Consequently they are de-boosted, so that the peak energy of BLR seed photons is lower. This shifts

the peak of the Comptonised emission to lower energies. The flux in the Fermi band ($1 - 100 \, \text{GeV}$, corresponding to $23.38 < \log \nu < 25.38$) drops significantly. The luminosity of a $10^{10} \, M_{\odot}$ FSRQ at $10^{24} \, \text{Hz}$ is almost 2 orders of magnitude less than a $10^9 \, M_{\odot}$ black hole at the same accretion rate (compare black and blue lines, Fig.6.1a). Redshifting the spectrum only exacerbates the shift of the Compton peak to lower energies and further reduces the Fermi flux. Consequently the redshift limits of $10^{10} \, M_{\odot}$ FSRQs are nearer those of $10^{7-8} \, M_{\odot}$ black holes.

Fig.6.3b shows the redshift limits for FSRQs as a function of accretion rate for four different black hole masses ($M = 10^7$ (red), 10^8 (magenta), 10^9 (blue) and $10^{10} \, M_{\odot}$ (black) and $\theta = 0$). z_{limit} increases with $\dot{m}$, however the rate of increase differs with mass.

The redshift limits for $10^{7-8} \, M_{\odot}$ FSRQs increase very slowly with accretion rate (magenta and red lines). The redshift limit for a $10^7 \, M_{\odot}$ FSRQ is ~ 0.5 at $\log \dot{m} = -2$ and ~ 0.75 at $\log \dot{m} = 1$. $10^8 \, M_{\odot}$ FSRQs show a similarly small factor ~ 3 increase over the same range in accretion rate. This is because for small masses the dominant cooling is through SSC, due to the small emission region size and high magnetic field. As a low mass FSRQ ($10^{7-8} \, M_{\odot}$) increases its accretion rate from $\log \dot{m} = -2$ to 1, its Compton spectrum changes from being high peaked (at $\sim 10^{24} \, \text{Hz}$) to low peaked ($\sim 10^{21} \, \text{Hz}$), analogous to the change in BL Lac spectra from high peaked to low peaked. The reason is the same: increasing $\dot{m}$ increases the cooling, shifting all the peak energies to lower frequency, because low mass FSRQs are similarly dominated by SSC cooling (plus low energy accretion disc seed photons), which always dominates over BLR emission. Even though the total luminosity is increasing, the shift of the peak emission to lower energies means the Fermi band flux increases more slowly and hence z_{limit} shows a very gradual increase.

In contrast, $10^9 \, M_{\odot}$ FSRQs show a much faster increase in z_{limit} with $\dot{m}$ (blue line, Fig.6.3b). This is because they are almost always dominated by BLR Compton scattering. The spectral energy density of BLR seed photons peaks at higher energy (see Fig.6.1b, blue dot-dashed line), than the synchrotron and disc seed photons which dominate in lower mass systems, hence the Compton peak is at higher energy (10^{23} compared to $10^{20} \, \text{Hz}$, compare blue and magenta lines, Fig.6.1a), so more of the

luminosity increase can be seen in the Fermi band. Only at the very lowest accretion rates ($\log \dot{m} \sim -2$) does the Fermi visibility of a $10^9 \, M_\odot$ FSRQ dip below that of a $10^8 \, M_\odot$ object. This is because the Compton cooling is slightly more efficient in the larger mass object, shifting its Compton peak to slightly lower energy and hence giving it a lower Fermi band flux.

The $10^{10} \, M_\odot$ example (black line, Fig.6.3b) shows a similar effect, with the z_{limit} increasing slowly at first and then more rapidly for $\log \dot{m} > -1$. This is because $R_{BLR} \propto (\dot{m}M)^{1/2}$ while $Z_{diss} \propto M$, hence for larger mass a higher $\dot{m}$ is needed for $R_{BLR} > Z_{diss}$, i.e. $R_{BLR} < Z_{diss}$ up to higher $\dot{m}$. While $R_{BLR} < Z_{diss}$, the Compton hump is produced from synchrotron, accretion flow and de-boosted BLR seed photons, so its peak is at lower frequency and the Fermi band flux ($\sim 10^{24}$ Hz) is significantly reduced. Once $R_{BLR} > Z_{diss}$ (at $\log \dot{m} \sim -1$ for $M = 10^{10} \, M_\odot$), Doppler boosted BLR seed photons dominate and z_{limit} increases dramatically.

6.2.4 Predicted FSRQ Population from Cosmological Simulations

Combining our scaled jet emission model with the results from cosmological simulations allows us to predict the population of FSRQs that should be detected by Fermi. As in the previous chapter, we use the black hole number densities from the Millennium Simulation (Springel et al., 2005; Fanidakis et al., 2011; Fanidakis et al., 2012), which predicts the number of SMBHs accreting at different redshifts together with their masses and accretion rates ($n(z, M, \dot{m})$).

We initially assume that all black holes accreting efficiently will produce a FSRQ jet. We can then calculate the number of Active Galactic Nuclei (AGN) hosting a FSRQ jet in each ($z, M, \dot{m}$) bin by using the number densities of black holes in the radiatively efficient accretion regime ($\dot{m} > 10^{-2}$). Again, if this number is less than 1 we use Poisson statistics to randomly determine whether a black hole is present or not. Each black hole in each ($z, M, \dot{m}$) bin is then assigned a random distance within this redshift bin and random θ_{obs}, assuming $\cos \theta_{obs}$ is distributed uniformly, and we calculate the observed spectrum to determine whether or not the jet would be visible to Fermi. We choose the flux limit of the Fermi 1-year catalogue, in order

to compare our simulation results with the observations presented in Shaw et al., 2012.

Fig.6.4a shows the predicted redshift distribution of Fermi visible FSRQs (black line). The predicted distribution peaks between redshifts $1 < z < 2.5$. This corresponds to the peak in quasar activity at $z \sim 2$. At later times ($z < 1$), typical black hole accretion rates drop below 10^{-2} due to systems running out of gas to accrete. At low accretion rates, the accretion flow becomes radiatively inefficient and no longer illuminates the BLR, effectively switching off the source of external seed photons, so that the black holes produce BL Lac rather than FSRQ type jets. A few systems remain at high accretion rates — these typically host smaller black holes ($\sim 10^7 \, M_\odot$), which haven't yet used up their gas supplies. These correspond to Seyfert galaxies in the local Universe. According to our criteria ($\dot{m} > 10^{-2}$), these black holes should host EC jets. However Fig.6.3a shows that the Fermi visibility of jets from such small black holes is poor, so their contribution to the number of FSRQs at late times ($z < 1$) is small.

Whilst the predicted redshift distribution peaks at $1 < z < 2.5$, there is a tail out to high redshifts, with the most distant FSRQs being detected out to $z \sim 5$. As redshift increases, the typical black hole mass decreases. At $z = 2$, the bulk of the accretion luminosity is produced by $10^8 \, M_\odot$ black holes (see Fig.5.4, previous chapter). For $z > 2$, the typical black hole mass producing the bulk of the accretion luminosity is dropping below $10^8 \, M_\odot$. Fig.6.3a shows how sharply the Fermi visibility drops with mass, more than halving for a decade drop in mass from $10^9 - 10^8 M_\odot$. Fig.6.3b shows that for small black hole masses ($\leqslant 10^8 M_\odot$), the increase in accretion rate at early times does not compensate for the drop in mass in terms of Fermi visibility (compare magenta and blue lines, Fig.6.3b). The decreasing tail of the predicted redshift distribution from $2.5 < z < 5$ is therefore due to the decreasing number density of $10^{8-9} \, M_\odot$ black holes and the increasing reliance on strongly beamed sources ($\theta_{obs} \sim 0$) to reach the Fermi flux limit.

The total number of Fermi visible FSRQs predicted by our simulation is ~ 26000, while the actual number of FSRQs detected in the Fermi 1 year catalogue is ~ 300 (Abdo et al., 2010). Our simulation overpredicts the number of Fermi visible FSRQs
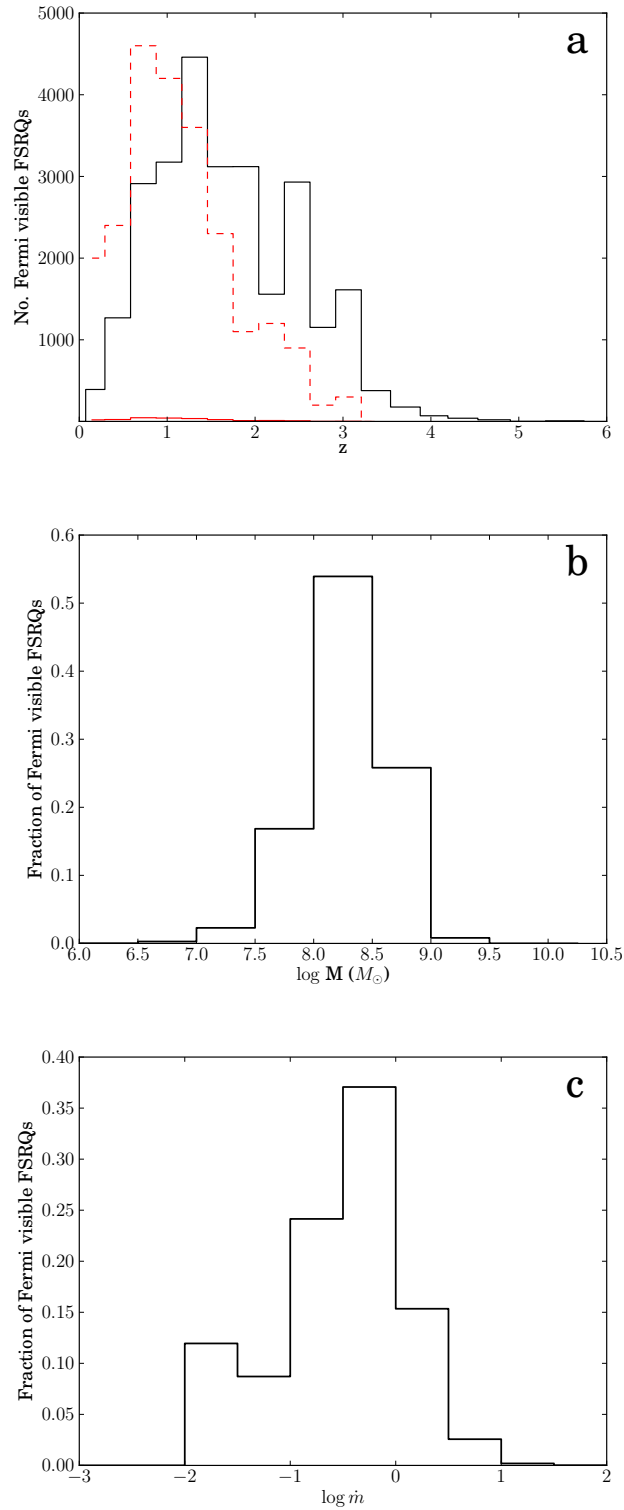
Figure 6.4: Simulation results using standard jet scalings. a). Predicted redshift distribution of Fermi visible FSRQs, assuming black holes of all spins accreting above $\dot{m} = 10^{-2}$ produce a FSRQ jet (black). Red solid line shows observed redshift distribution of Fermi detected FSRQs from Fermi 1-year catalogue. Red dashed line shows observed redshift distribution $\times 100$. b). Predicted mass distribution of Fermi visible FSRQs. c). Predicted accretion rate distribution of Fermi visible FSRQs.

by ~ 2 orders of magnitude. This is one order of magnitude less than the 3 orders of magnitude discrepancy found in the previous chapter using the same method to predict the Fermi population of BL Lacs. Nevertheless, a 2 order of magnitude discrepancy is still too large to be explained by the sky coverage limit of Fermi ($|b| > 10°$ implying 80% of the sky is included).

In Fig.6.4a we show the observed redshift distribution of Fermi detected FSRQs from Shaw et al. (2012) (solid red line). The dashed red line shows the observed redshift distribution $\times 100$ for better comparison of the shapes of the two distributions. The observed redshift distribution peaks later ($0.5 < z < 1.5$ rather than $1 < z < 2.5$), with no FSRQs detected in the 1LAT catalogue with $z > 3.5$. Not only is the total number of FSRQs overpredicted, but the shape of the redshift distribution also does not match the observations.

Fig.6.4b and c show the predicted mass and accretion rate distributions of Fermi visible FSRQs from the simulation. The typical predicted FSRQ accretion rate is $-1 < \log \dot{m} < 0$, since Fermi visibility increases with accretion rate. Above Eddington, the increase in Fermi flux with $\dot{m}$ becomes progressively less (see Fig.6.2a) and the number density of super-Eddington sources of sufficient mass ($> 10^8 \, M_\odot$) drops off sharply, both of which result in the typical FSRQ accretion rate being just sub-Eddington. This is in agreement with the findings of Shaw et al. (2012), where most FSRQs are observed to have $-1 < \log \dot{m} < 0$. The typical predicted mass is $10^{8-8.5} \, M_\odot$, since these FSRQs are bright in the Fermi band and most numerous at $1 < z < 2$ where quasar activity peaks. The results of Shaw et al. (2012) indicate that the observed peak extends to slightly higher mass $10^{8-9} \, M_\odot$.

6.2.5 Dependence on Black Hole Spin?

In assuming every black hole accreting with $\dot{m} > 10^{-2}$ is capable of producing a FSRQ jet our simulation overpredicts the number of Fermi detected FSRQs by two orders of magnitude. Clearly another factor is reducing the number of FSRQs detected by Fermi. In the previous chapter we found that the number of Fermi detected BL Lacs was similarly overpredicted (by 3 orders of magnitude) when the same technique was applied to predict the observed population of BL Lacs (i.e. all

Figure 6.5: Effect of imposing a spin cut on simulation results using standard jet scalings. Blue line shows predicted redshift distribution of Fermi visible FSRQs, assuming only black holes with $\dot{m} > 0.01$ and spin $a > 0.77$ produce a FSRQ jet. Red line shows observed redshift distribution of Fermi detected FSRQs. Black line shows predicted redshift distribution with no spin cut.

black holes accreting *below* $\dot{m} = 10^{-2}$ produce BL Lac type jets). The observed numbers of BL Lacs, and their redshift distribution, are much better reproduced assuming that only high spin black holes with $\dot{m} < 10^{-2}$ produce BL Lac type jets, specifically $a > 0.8$. This suggests black hole spin might be important in the production of highly relativistic $\Gamma = 15$ jets in BL Lacs. Maraschi et al. (2012) find that the efficiency of spin-powered jet production drops off sharply below 0.8, which provides additional support for an effective spin threshold for relativistic jet production at $a \sim 0.8$. We investigate whether this may also be the case for FSRQs.

The chaotic accretion model for the evolution of black hole spin is required to match the population of Fermi detected BL Lacs. In this model, accretion events are randomly aligned, producing predominantly low spin black holes, while high spin is rare, being only produced through black hole-black hole mergers, so requiring high spin reduces the predicted number of BL Lacs, in better agreement with the observed numbers. Motivated by this, we continue to use the black hole spin distributions predicted by the chaotic accretion model.

Fig.6.5 shows the results of applying a spin cut to our simulation results for FSRQs. The black line shows our original results, where all black holes with $\dot{m} >$

Figure 6.6: Redshift distribution of high spin black holes ($a > 0.8$) from the Millennium simulation. Red line shows black holes accreting with $\dot{m} < 10^{-2}$ (corresponding to BL Lacs), blue line shows black holes accreting at $\dot{m} > 10^{-2}$ (corresponding to FSRQs), black line shows total.

10^{-2} produce a FSRQ jet (now with a logarithmic y axis). The red line shows the observed FSRQ redshift distribution. The blue line shows the simulation results after imposing a spin cut, so that only black holes with $\dot{m} > 10^{-2}$ and $a > a_{cut}$ produce a FSRQ jet. We find a lower spin cut is required to reproduce the observed number of FSRQs compared to the BL Lacs: $a_{cut} \sim 0.77$ rather than 0.8.

Closer comparison of the red and blue lines shows that, although the total number of FSRQs is better reproduced, the simulation cannot reproduce the tail out to high redshifts (> 2). Imposing a spin cut limits the maximum expected FSRQ redshift to ~ 2.3. This is clearly in conflict with observations.

Fig.6.6 shows the number density of high spin black holes ($a > 0.8$) as a function of redshift from the Millennium Simulation (black line). This has two peaks, one at $z = 0$ and one at $z = 5$. The red line shows the number density of high spin black holes with low accretion rates ($\dot{m} < 10^{-2}$). These are responsible for the peak at $z = 0$. They get their high spins from late gas poor mergers, so represent the most massive black holes. This is the population of black holes responsible for the production of BL Lac jets. The blue line shows the number density of high spin black holes with high accretion rates ($\dot{m} > 10^{-2}$). These are responsible for the peak

at $z = 5$. They acquire their high spins through early mergers of much smaller black holes. At early times the black holes still have a plentiful gas supply (hence their high accretion rates) and subsequent chaotic accretion gradually spins down the black holes, so that the number density of high spin, high accretion rate black holes drops with decreasing redshift, not only because typical accretion rates drop, but also because most black holes are losing their earlier high spins. If we require high spin as well as high accretion rate to produce a FSRQ jet, then these are the black holes that should be responsible for FSRQs. However, when we include a spin cut in our simulation our results cannot replicate the observed FSRQ population between $2 < z < 3$. Fig.6.6 (blue line) shows that the number of high spin, high $\dot{m}$ black holes has dropped significantly by $2 < z < 3$. Many of the high spin black holes that remain are still small ($10^{7-8} \, M_\odot$), because if they had grown significantly since their last merger the process of chaotic accretion would have reduced their spins. As a result they are not massive enough to be Fermi visible in our simulation. Yet the observations show there are some relatively massive black holes with FSRQ jets at these redshifts. Our simulation accounts for spin-ups due to mergers, so these black holes cannot have acquired their spins through mergers, but must have maintained them whilst growing by accretion. This suggests that early accretion may be more organised than late accretion, i.e. gas from a star formation episode is accreted in one go from a single direction so that accretion alone can spin up the black hole. Perhaps there is a transition from prolonged accretion to more chaotic accretion as gas supplies diminish and redshift decreases. In assuming a chaotic accretion mode throughout, we have underestimated the number of high spin black holes at higher redshifts ($z > 2.5$), where accretion may be more ordered. Although in order to reproduce the observed BL Lac population, chaotic accretion must have taken over and reduced many black holes to low spins by $z = 2$ when accretion rates drastically drop.

This may suggest why a lower spin cut is required in our simulation to match the FSRQ population compared to the BL Lacs. Alternatively spin may genuinely be less important in high accretion rate FSRQs, since FSRQs have more accretion power available to power a jet than the low accretion rate BL Lacs. However given the

uncertainties in tracking black hole spin evolution in the cosmological simulations no firm conclusions on this can be drawn.

An alternative explanation is that we have simply underestimated the Fermi visibility of the small mass FSRQs which dominate at high redshifts, by relying on standard jet scalings to predict their luminosity. In the next section we test this assumption.

6.3 The Gamma-Ray Loud NLS1 as a Test of Jet Scalings

A key assumption affecting the results of our simulations are the scalings we choose to calculate R_{diss}, P_{rel} and B. Gamma-ray loud NLS1s (γNLS1s) provide a unique opportunity to test these scalings. NLS1s are small mass ($10^{6-8}\,M_\odot$), high accretion rate AGN. Most are radio quiet, however a small subset are found to be radio loud (Komossa et al., 2006). Of these radio loud NLS1s, 7 have been detected by Fermi. These γNLS1s show Doppler boosted jet emission with a weaker synchrotron hump and strong Compton emission so that their SEDs appear like 'mini FSRQs' (Abdo et al., 2009). The high accretion rates of NLS1s ($\dot{m} \sim 1$) mean that their accretion flows are in the radiatively efficient regime, giving them a BLR (albeit with relatively narrow broad lines) so they should correspond to low mass FSRQs. In which case we should be able to replicate their spectra simply by turning down the mass in our mean FSRQ spectral model.

PMN J0948+0022 was the first γNLS1 to be discovered. A multi-wavelength monitoring campaign was carried out on the source in 2009. Abdo et al. (2009) subsequently fitted its broadband spectrum (Fig.6.7, top panel) with the jet model of Ghisellini & Tavecchio (2009). In Table 6.1 we show their derived values of R_{diss}, P_{rel} and B. We also show the expected values of R_{diss}, P_{rel} and B for this source that result from scaling the mean FSRQ parameters of G10 as $R \propto M$, $P_{rel} \propto \dot{m}M$ and $B \propto (\dot{m}/M)^{1/2}$ ('standard scalings') according to the mass and accretion rate of PMN J0948+0022 ($M = 1.5 \times 10^8$, $\dot{m} = 0.5$). These scalings correspond to assuming both the power injected into relativistic electrons and the power in magnetic fields

Figure 6.7: Top: Observed spectrum of γNLS1 PMN J0948+0022, from Abdo et al. (2009). Bottom: Predicted γNLS1 spectra. Dashed black line shows the expected jet spectrum for a black hole with mass and accretion rate of the γNLS1 PMN J0948+0022 ($M = 1.5 \times 10^8 \, M_\odot$, $\dot{m} = 0.5$), from scaling the mean FSRQ spectrum of G10 according to standard jet scalings ($R_{diss} \propto M$, $P_{rel} \propto \dot{m}M$ and $B \propto (\dot{m}/M)^{1/2}$). Red line shows resulting spectrum replacing R_{diss}, P_{rel} and B with the observed values found by Abdo et al. (2009) to fit PMN J0948+0022. Solid black line shows expected spectrum using the alternative jet scalings ($R_{diss} \propto \dot{m}M$, $P_{rel} \propto \dot{m}M$ and $B \propto (\dot{m}M)^{-1/2}$). See Table 6.1 for parameter values.

Parameter	$< FSRQ >$	γ NLS1		
		Observed	§6.2 Standard Scalings	§6.4 Alternative Scalings
M ($M_\odot$)	10^9	1.5×10^8		
$\dot{m}$ (L/L_{Edd})	0.1	0.5		
R_{diss} ($\times 10^{15}$ cm)	18.9	6.75	2.835	14.2
P_{rel} ($\times 10^{43}$ erg s^{-1})	2.0	2.3	1.5	1.5
B (G)	2.6	4.1	15	3.0

Table 6.1: Comparison between the observed jet parameters for the γNLS1 PMN J0948+0022 (Abdo et al., 2009) with those expected from scaling the mean FSRQ jet parameters from G10 according to the standard jet scalings of Section 6.2 ($R_{diss} \propto M$, $P_{rel} \propto \dot{m}M$, $B \propto (\dot{m}/M)^{1/2}$) and using the alternative jet scalings of Section 6.4 ($R_{diss} \propto \dot{m}M$, $P_{rel} \propto \dot{m}M$, $B \propto (\dot{m}M)^{-1/2}$).

is a fixed fraction of the accretion power, i.e. $P_{rel} \propto P_B \propto P_{acc} \propto \dot{m}M$. In Table 6.1 we also list the mean FSRQ values for reference.

The values of R_{diss} and P_{rel} found from fitting the observed spectrum of PMN J0948+0022 are both slightly larger than expected scaling from the mean FSRQ parameters. However the biggest difference is in the magnetic field strength. Scaling the mean FSRQ magnetic field strength of 2.6 G as $B \propto (\dot{m}/M)^{1/2}$ implies PMN J0948+0022 should have a jet magnetic field of 15 G. In reality, the magnetic field required to fit its spectrum is much smaller (4 G). This is larger than the mean FSRQ value, as expected for its smaller mass, but not nearly as large as the standard scalings predict.

In Fig.6.7 (bottom panel) we show the effect of this on the predicted spectrum. The dashed black line shows the spectrum produced taking the mean FSRQ parameters and scaling R_{diss}, P_{rel} and B to the mass and accretion rate of PMN J0948+0022 according to the standard scaling relations, i.e. standard scaling values from Table 6.1. The red line shows the same spectrum, but replacing R_{diss}, P_{rel} and B with the values found by Abdo et al. (2009) from fitting the observed spectrum of PMN J0948+0022 (i.e. observed values in Table 6.1). The red line is not a fit to PMN J0948+0022, since we have kept the other parameters (ϕ, Γ etc.) the same as the mean FSRQ, in order to show just the effect of correcting R_{diss}, P_{rel} and B (although we note that the injected electron distribution parameters of PMN J0948+0022 are not very different to those of the mean FSRQ). It is clear that replacing the values predicted by the standard scalings with the observed values has a big effect on the shape of the spectrum. The high magnetic field predicted by the standard scal-

ing relations causes synchrotron and SSC emission to dominate the dashed black predicted spectrum, resulting in synchrotron and Compton peaks of similar luminosity. The resulting spectral shape is similar to that of a BL Lac, where the only source of emission is synchrotron and SSC emission. In contrast, when we use the observed values, where the magnetic field is much lower, the Compton emission is dominated by up-scattering of external seed photons and the synchrotron emission is suppressed (red spectrum). As a result the Compton peak is much brighter than the synchrotron peak — a spectral shape typical of FSRQs. However, according to the standard scaling relations, γNLS1s *shouldn't* look like mini FSRQs — SSC should dominate their Compton humps not EC. Yet they do (Fig.6.7, top panel). This suggests the standard scaling relations break down between normal mass FSRQs and γNLS1s. Or that FSRQs simply do not follow them. This is surprising, given that these scaling relations are based on just two assumptions: that the size-scales of the jet should scale with black hole mass, and that the power injected into relativistic electrons and the power in magnetic fields is a fixed fraction of the accretion power.

Fig.6.7 (bottom panel) shows that replacing the parameters predicted by standard jet scalings with those observed for a γNLS1 increases the Fermi flux (at $\nu \sim 10^{22}$ Hz) by half an order of magnitude. This suggests the small mass FSRQs in our simulation should be brighter and hence more visible than we have estimated. Consequently our original predicted population, which already overestimates the observed population by 2 orders of magnitude, should be an underestimate. This only increases the need for some other factor, such as a limit on black hole spin, to reduce the predicted numbers. It would also increase the visibility of high redshift FSRQs, which are typically lower mass, so potentially making it easier for our spin cut simulation to reproduce the high redshift tail of the observed redshift distribution.

6.4 Alternative Jet Scaling

Comparison of our scaled spectra with the observed spectrum of a γNLS1 suggests FSRQs do not follow the standard jet scalings. We now investigate the effects of

trying an alternative jet scaling.

The biggest discrepancy between the observed γNLS1 values and those predicted by standard jet scalings is in the magnetic field. This discrepancy is further amplified by the fact the emitted radiation depends on $U_B \propto B^2$.

Our standard scalings assume $B \propto (P_B/R_{diss}^2)^{1/2}$, where $R_{diss} \propto M$ and $P_B \propto P_{acc}$. The relativistic electrons are accelerated by inhomogeneities in the magnetic field, so P_B and P_{rel} should scale the same way. If $P_{rel} \propto \dot{m}M$ then P_B, which caused the particle acceleration in the first place, must be scaling as $\dot{m}M$. The only way to change how B scales is then to change the scaling of R_{diss}.

It seems reasonable that jet size-scales should scale with the mass of the black hole, however we have so far neglected any scaling with accretion rate. If the jet base, where the bulk of the high energy emission comes from, represents a standing shock caused by some pressure balance between high pressure material outflowing into less dense material, this could well depend on accretion rate. In which case a higher accretion rate means high central pressure, with pressure equilibrium occurring further out. R_{diss} would then have a dependence on $\dot{m}$ as well as M, assuming $R_{diss} = \phi Z_{diss}$ with constant ϕ. If $R_{diss} \propto \dot{m}M$ and $P_{rel} \propto P_B \propto \dot{m}M$, then this gives $B \propto (\dot{m}M)^{-1/2}$.

In Table 6.1 we show the predicted R_{diss}, P_{rel} and B values for PMN J0948+0022 using these alternative scalings. The predicted magnetic field of $3\,\mathrm{G}$ is now much nearer the observed value than the estimate of $15\,\mathrm{G}$ from using standard jet scalings. We have not altered how the values scale with mass so this is simply a result of the NLS1's high accretion rate partially compensating for its lower mass. The solid black line in Fig.6.7 (bottom panel) shows the predicted spectrum using these new values. The Compton dominance and the level of Fermi flux is much better reproduced (compare black and red solid lines). The Compton dominance is actually slightly overestimated — comparison of the synchrotron humps shows that the synchrotron emission is more suppressed using these alternative scalings. This is because R_{diss} is now much larger (14.2×10^{15} compared to $6.75 \times 10^{15}\,\mathrm{cm}$) due to the additional dependence on $\dot{m}$. We have slightly overcorrected the underestimated R_{diss} of the previous standard scalings. However the effect of this is less in the EC dominated

Fermi band, which we are interested in.

In the following sections we test the effect of these alternative scalings ($R_{diss} \propto \dot{m}M$ and $B \propto (\dot{m}M)^{-1/2}$, with P_{rel} still $\propto \dot{m}M$) on the visibility of different FSRQs and the predicted Fermi visible population.

6.4.1   Spectral Change with Accretion Rate

We have not changed how the FSRQs scale with mass, so for constant accretion rate the alternative jet scalings cause the spectrum to change as in Fig.6.1, i.e. the same way as the standard jet scalings. By allowing R_{diss} to also scale with $\dot{m}$ we have only changed how the spectrum should scale with accretion rate.

Fig.6.8 shows three sequences of spectra, analogous to Fig.6.2a, demonstrating how the spectrum changes with accretion rate for this alternative jet scaling. The three panels represent three different black hole masses — 10^9, 10^8 and $10^7\,M_\odot$ (top to bottom). In each the mass is kept constant and accretion rate increases from $\log\dot{m} = -2$ (magenta) to 1 (black).

It is clear that the accretion rate at which the FSRQ is most Compton dominant (i.e. appears most like a typical FSRQ spectrum) changes with mass. For a $10^9\,M_\odot$ black hole the spectrum is most Compton dominant at 'low' accretion rates ($-1 < \log\dot{m} \leqslant -0.5$). For a $10^8\,M_\odot$ black hole it's at roughly Eddington ($0 < \log\dot{m} \leqslant 0.5$), and for a $10^7\,M_\odot$ black hole it's super-Eddington ($\log\dot{m} \geqslant 0.5$).

This is a direct result of $Z_{diss} \propto \dot{m}M$ (since we assume $Z_{diss} \propto R_{diss} \propto \dot{m}M$). For a $10^9\,M_\odot$ FSRQ (Fig.6.8a), Z_{diss} is already large due to its large mass. Only a moderate increase in $\dot{m}$ is required to push Z_{diss} beyond the radius of the BLR ($R_{BLR} \propto (\dot{m}M)^{1/2}$). The lack of blue shifted BLR seed photons results in a much lower frequency Compton peak (cyan, blue and black spectra), firmly shifting the peak out of the Fermi band.

For a $10^8\,M_\odot$ FSRQ (Fig.6.8b), Z_{diss} is inherently smaller due to the smaller black hole mass. Only the very highest accretion rates are enough to push Z_{diss} beyond the BLR (black spectrum). Hence $10^8\,M_\odot$ FSRQs appear 'FSRQ-like' up to higher accretion rates. However at the lowest accretion rates ($\log\dot{m} \lesssim -1.5$) they start to lose their Compton dominance again. This is because, for a $10^8\,M_\odot$
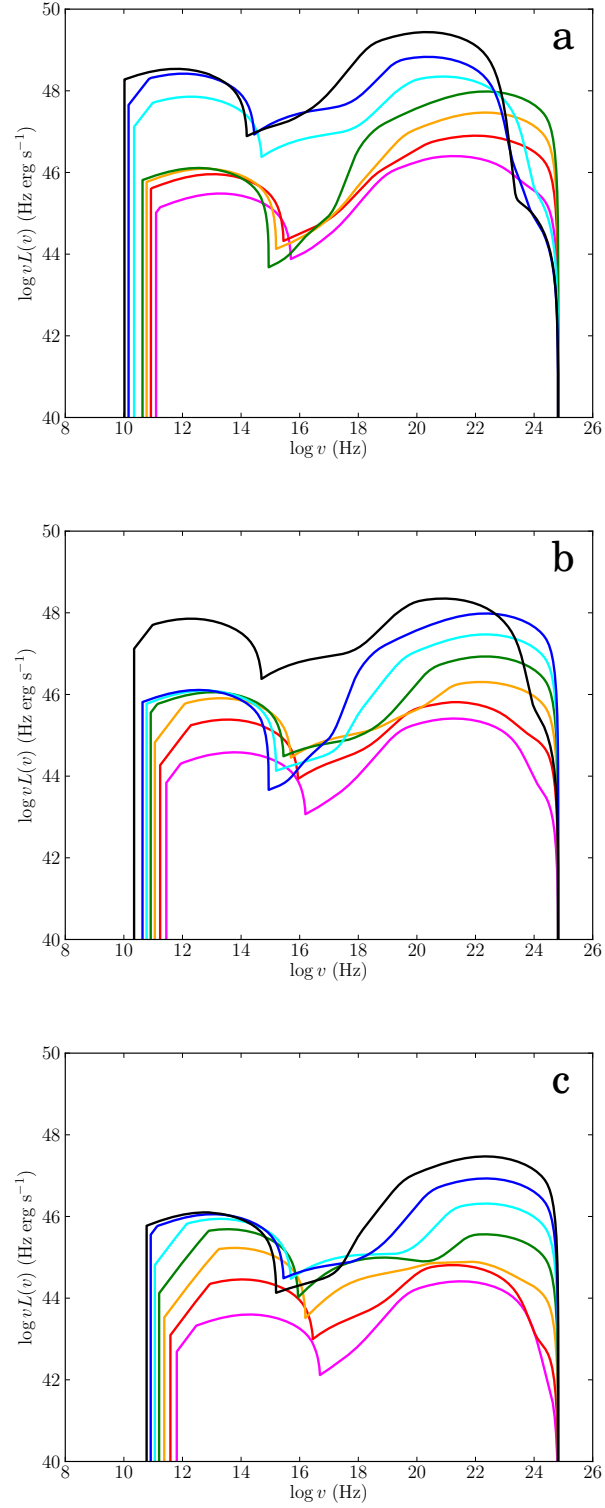
Figure 6.8: FSRQ model SEDs for fixed black hole mass and increasing accretion rate ($\log \dot{m} = -2$ (magenta), -1.5 (red), -1 (orange), -0.5 (green), 0 (cyan), 0.5 (blue) and 1.0 (black)) for a). $M_{BH} = 10^9$, b). 10^8 and c). $10^7 \, M_\odot$, using the alternative jet scalings.

FSRQ, Z_{diss} is now getting so close to the black hole that direct accretion flow seed photons and SSC begin to dominate over BLR seed photons. This reduces both the frequency of the Compton peak and the Compton dominance.

For the smallest $10^7 M_\odot$ FSRQs (Fig.6.8c), Z_{diss} never goes beyond R_{BLR}, even for the highest accretion rates. Synchrotron and direct disc seed photons now dominate up to higher accretion rates ($\log \dot{m} \sim -0.5$), since the very small black hole mass means Z_{diss} is already very compact through $Z_{diss} \propto M$ before reducing it further by $Z_{diss} \propto \dot{m}$.

When the Compton peak is shifted to lower frequencies, either through Z_{diss} becoming larger than R_{BLR} or Z_{diss} becoming very small, this greatly reduces the Fermi flux. This means FSRQ type jets are clearly most likely to be detected when they have strong Compton dominance, i.e. when they are most 'FSRQ-like'. Fig.6.8 shows that, selecting only the most FSRQ-like spectra picks out different accretion regimes for different mass FSRQs. So that a $10^9 M_\odot$ FSRQ is likely to be accreting at 10% of Eddington, while a $\lesssim 10^8 M_\odot$ FSRQ is likely to be accreting at Eddington. This would explain why γNLS1s, which like all NLS1s accrete at roughly the Eddington limit, still have jets that resemble FSRQs, despite the typical FSRQ being more massive and at lower accretion rate. The existence of γNLS1s therefore requires that R_{diss} ($\propto Z_{diss}$) scales not only with mass but also with accretion rate.

6.4.2 Effect on FSRQ Visibility

Fig.6.9 shows how scaling $R_{diss} \propto \dot{m}M$ (and by extension $Z_{diss} \propto \dot{m}M$ and $B \propto (\dot{m}M)^{-1/2}$, with P_{rel} still $\propto \dot{m}M$) affects the Fermi visibility of FSRQs as a function of accretion rate. We show the redshift limits as a function of accretion rate for four different masses: 10^{10}, 10^9, 10^8 and $10^7 M_\odot$ (in black, blue, magenta and red respectively), analogous to Fig.6.3b, which shows the same plot but using the standard jet scalings.

Each mass shows the same trend in z_{limit} with accretion rate — specifically z_{limit} increases with $\dot{m}$ until a peak z_{limit} of ~ 8, after which the visibility drops sharply. This is a result of the effect shown in Fig.6.8a. As $\dot{m}$ increases, so does the total luminosity and Z_{diss} until eventually $Z_{diss} > R_{BLR}$, at which point the peak of

Figure 6.9: Redshift limits for Fermi visible FSRQs as a function of accretion rate, for $M_{BH} = 10^7$ (red), 10^8 (magenta), 10^9 (blue) and $10^{10} M_\odot$ (black) and $\theta = 0$, using the alternative jet scalings.

the Compton hump shifts to lower frequencies, severely reducing the Fermi band flux. As a consequence the limiting redshift out to which the FSRQ is visible drops sharply.

The accretion rate at which the Fermi visibility peaks (i.e. where z_{limit} is greatest) is higher for smaller masses. Again this is a result of the effect shown in Fig.6.8. The Fermi band flux is a maximum when the FSRQ is strongly Compton dominant, i.e. when Z_{diss} is neither too close nor too far from the black hole. Since $Z_{diss} \propto \dot{m}M$, this optimum region will be at a higher accretion rate for a smaller mass. A FSRQ will be visible out to the highest redshifts when its Fermi flux is brightest, hence the maximum limiting redshift occurs at a higher accretion rate for a smaller mass FSRQ.

Comparison of Fig.6.9 with Fig.6.3b shows that the predicted redshift limits are very different for these alternative scalings compared with using the standard jet scalings. The highest mass black holes ($10^{10} M_\odot$) are now most visible at low accretion rates, since this is when $Z_{diss} < R_{BLR}$. In contrast, using the standard jet scalings Z_{diss} was always $> R_{BLR}$ at this mass. This meant $10^{10} M_\odot$ FSRQs always had low frequency peaked Compton humps and were only visible at the highest accretion rates, because their enormous luminosities gave sufficient Fermi

flux despite the peak of the Compton hump being below the Fermi band.

However the biggest difference is in the visibility of low mass FSRQs ($10^{7-8} \, M_\odot$). Using the standard jet scalings, Z_{diss} was always close to the black hole and direct disc seed photons and SSC dominated, again lowering the frequency at which the Compton hump peaked and limiting the Fermi band flux. Consequently these objects were never visible beyond $z \sim 3$. In contrast the alternative scalings, with $Z_{diss} \propto \dot{m}M$, allow an increase in accretion rate to compensate their small mass, so that at high accretion rates they can be Compton dominant and Fermi bright. As a result, for a limited range of accretion rates, these objects are Fermi visible beyond $z = 4$.

This will have implications for our simulated population of Fermi visible FSRQs. Our spin cut results using standard jet scalings fail to reproduce the observed number of high redshift FSRQs ($z > 2.5$). As redshift increases the typical black hole mass decreases and the typical accretion rate increases. Using these alternative jet scalings means FSRQ jets from these small mass, high accretion rate black holes should be more visible.

6.4.3 Effect on Predicted FSRQ Population

We now rerun our simulation using these alternative jet scalings ($R_{diss} \propto \dot{m}M$, $P_{rel} \propto \dot{m}M$ and $B \propto (\dot{m}M)^{-1/2}$). Fig.6.10 shows the predicted redshift distribution of Fermi visible FSRQs using the alternative jet scalings (solid lines) compared with our previous results using standard jet scalings ($R_{diss} \propto M$, $P_{rel} \propto \dot{m}M$ and $B \propto (\dot{m}/M)^{1/2}$, dashed lines).

Using the alternative jet scalings almost triples the predicted number of Fermi detected FSRQs (~ 73000, compared to 26000 using standard jet scalings), assuming every black hole with $\dot{m} > 10^{-2}$ produces a FSRQ jet. The difference between the numbers predicted using the alternative scalings versus the standard scalings increases as redshift increases (compare solid and dashed black lines, Fig.6.10). This is because the alternative scalings increase the visibility of small mass, high mass accretion rate black holes and as redshift increases an increasing fraction of the active black holes fall into this category.

Figure 6.10: Comparison of redshift distributions of Fermi visible FSRQs predicted using the alternative jet scalings (solid lines) with those predicted using standard jet scalings (dashed lines). Black lines compare results where all black holes with $\dot{m} > 0.01$ produce a FSRQ jet. Blue lines compare results assuming only black holes with $\dot{m} > 0.01$ and spin $a > 0.77$ produce a FSRQ jet. Red solid line shows observed redshift distribution of Fermi detected FSRQs.

Fig.6.11a and b compare the predicted mass and accretion rate distributions of Fermi visible black holes using the two different scalings, assuming every black hole with $\dot{m} > 10^{-2}$ produces a FSRQ jet. Again solid lines show results using the alternative scalings and dashed lines show the standard scaling results. Fig.6.11a shows that small mass FSRQs ($< 10^8\,M_\odot$) make up a much larger fraction of the Fermi detections when the alternative jet scalings are used (~ 0.45 as opposed to ~ 0.2), with the largest fraction of detections having masses between $10^{7.5-8.5}\,M_\odot$. Fig.6.8b shows that a $10^8\,M_\odot$ FSRQ is Fermi brightest for $-0.5 < \log\dot{m} < 0.5$. Consequently these are the accretion rates that form the largest fraction of Fermi detections in the simulation using the alternative jet scalings (Fig.6.11b, solid line). Low accretion rate FSRQs ($\log\dot{m} < -1$) now make up a much smaller fraction of the total number of detections than when we used the standard jet scalings (~ 0.05 as opposed to ~ 0.2). This is partly due to the much larger number of high accretion rate FSRQs detected and also because low mass, low mass accretion rate FSRQs are slightly less Fermi visible using the alternative jet scalings (compare red and magenta lines, Fig.6.9 and Fig.6.3b). In contrast, high mass, low mass accretion
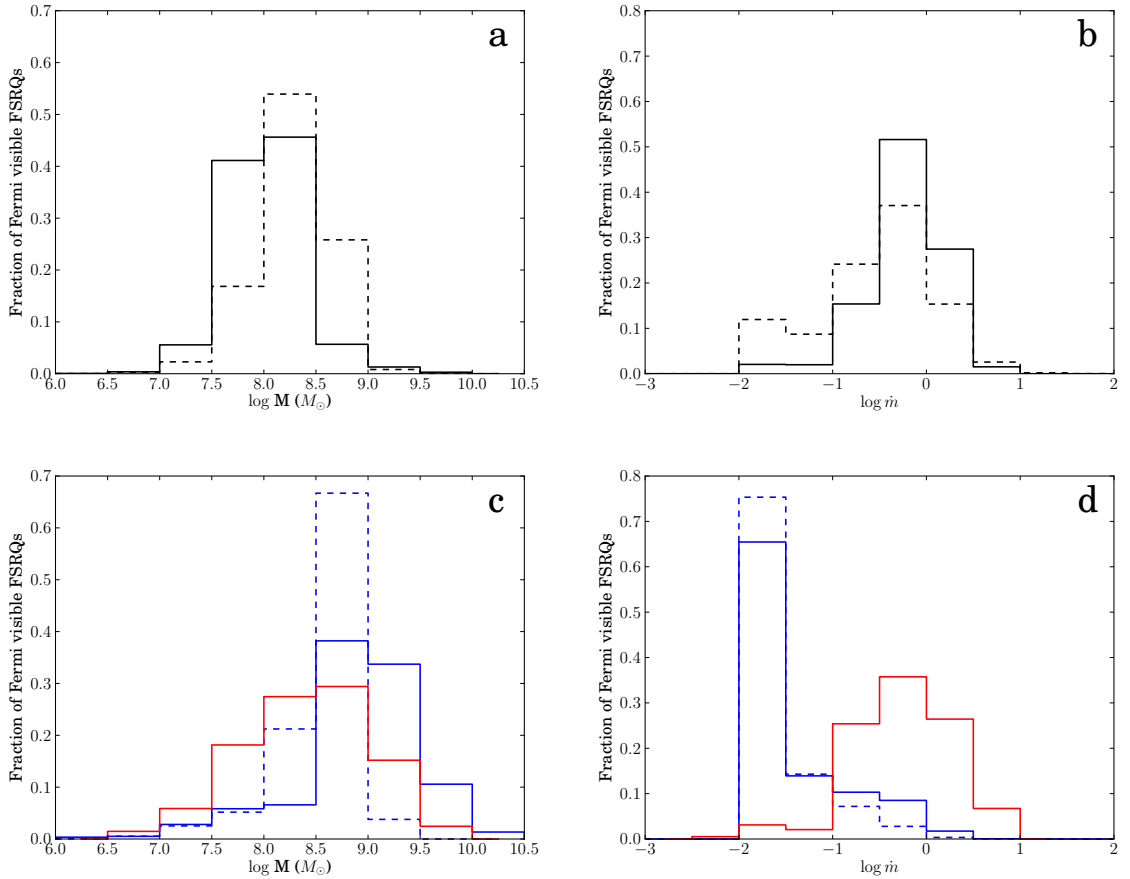
Figure 6.11: a). Predicted mass distribution of Fermi visible FSRQs using the alternative jet scalings (solid line) compared with standard jet scalings (dashed line), assuming all black holes with $\dot{m} > 0.01$ produce a FSRQ jet. b). Predicted accretion rate distribution of Fermi visible FSRQs using the alternative jet scalings (solid line) compared with standard jet scalings (dashed line), assuming all black holes with $\dot{m} > 0.01$ produce a FSRQ jet. c). & d). are same as a) and b) but compare the results after imposing a spin cut of $a_{cut} = 0.77$ so that only high spin, high accretion rate black holes produce a FSRQ jet. Red lines show observed mass and accretion rate distributions from the data of Shaw et al. (2012).

rate FSRQs are more Fermi visible using the alternative jet scalings (compare black lines, Fig.6.9 and Fig.6.3b), which is why the predicted mass distribution extends to slightly higher masses when using the alternative jet scalings, however the number density of these is very low.

In Fig.6.10 we also show the predicted population of Fermi detected FSRQs using the alternative jet scalings, assuming only high spin, high accretion rate black holes produce FSRQ jets ($a > 0.77$ and $\dot{m} > 0.01$, solid blue line). Despite increasing the visibility of small mass, high accretion rate black holes, which dominate at high redshifts, we still cannot reproduce the tail of the observed redshift distribution at high redshifts ($z > 2$). This is due to a lack of high spin black holes between $2 < z < 3$ in the cosmological simulation (see Fig.6.6).

In Fig.6.11c and d we show the effect of imposing a spin cut on the predicted mass and accretion rate distributions (blue lines). Again solid lines show results using the alternative scalings and dashed lines show the standard scaling results. The red lines show the observed mass and accretion rate distributions of Fermi detected FSRQs from the data of Shaw et al. (2012). Where they have multiple mass estimates for individual objects we use in order of preference masses derived from Hβ line width, MgII and CIV, and as in Shaw et al. (2012) convert the local continuum luminosity to a bolometric luminosity using the factors from Richards et al (2006) to estimate accretion rate. Comparing Fig.6.11a with Fig.6.11c shows that restricting FSRQ jet production to high spin objects has skewed the predicted mass distribution towards higher masses for both scalings. Instead of the majority of detections having masses in the range $10^{7.5-8.5} \, M_\odot$ (alternative scalings, solid black line, Fig.6.11a) the majority of detections now have masses between $10^{8.5-9.5} \, M_\odot$ (solid blue line, Fig.6.11c). This is because the cosmological simulation assumes chaotic accretion at all redshifts so the only highly spinning black holes are produced by mergers and are necessarily the most massive. In contrast the observed mass distribution (red line, Fig.6.11c) is much more symmetric. It peaks at roughly the same mass ($10^{8.5-9} \, M_\odot$) as the spin cut distributions but shows at least as many Fermi detected FSRQs have masses below this as above. This implies more low mass black holes are highly spinning than the cosmological simulations account for,

if high spin is required to produce a FSRQ type jet.

Fig.6.11b and d compare the predicted accretion rate distributions before and after imposing a spin cut, and with the observed distribution (red line, Fig.6.11d). Before imposing a spin cut, the predicted accretion rate distribution using both standard jet scalings and the alternative scalings peaks at $-1 < \log \dot{m} < 0.5$, in agreement with the observations (compare black lines Fig.6.11b with red line Fig.6.11d). However after imposing a spin cut the predicted distributions are heavily distorted towards the lowest accretion rates ($-2 < \log \dot{m} < -1.5$, since below $\log \dot{m} = -2$ we assume the jet appears as a BL Lac not a FSRQ). This is clearly in disagreement with the observations (compare blue and red lines, Fig.6.11d). This is again because the cosmological simulation relies only on black hole-black hole mergers to produce high spin black holes, and the most visible of these are those produced at late times in the local universe and as a consequence they are slowly accreting. The blue distributions in Fig.6.11d show that imposing a spin cut has merely caused us to select the 'high accretion rate' tail of the population of massive, generally lower accretion rate black holes that are responsible for BL Lac type jets. In reality the observed mass, accretion rate and redshift distributions suggest that this is not the population of black holes primarily responsible for producing FSRQ type jets. Instead there is a population of possibly highly spinning black holes that are smaller and at higher accretion rates and our failure to match the redshift distribution suggests they are at correspondingly higher redshifts. A solely chaotic model of accretion cannot produce this population, as our simulations have shown. This provides support for models where accretion is more ordered and prolonged in the early universe, allowing some highly accreting black holes to acquire and/or maintain high spins and so produce FSRQ type jets. Then as accretion rates drop, accretion becomes more disordered and chaotic, spinning most black holes down to zero. As a result most black holes in the local universe are low spin, expect for a small fraction which have undergone late gas-poor mergers, causing their black holes to spin up again. These are the most massive black holes and since they are at low redshift they have low accretion rates, so the relativistic jets they produce appear as BL Lac type jets, not FSRQs.

Our original simulations show that clearly not every black hole with $\dot{m} > 0.01$ produces a FSRQ jet. There is another condition which must be satisfied in order to produce a highly relativistic $\Gamma \sim 13$ jet. Assuming this condition is a highly spinning black hole can dramatically reduce the predicted number of Fermi detected FSRQs by the two orders of magnitude required to match the observations. The fact that the shape of the observed mass and accretion rate distributions cannot be matched indicates a limit on the ability of the cosmological simulations to track the evolution of black hole spin. In which case the observed mass, accretion rate and redshift distributions of FSRQs could prove a useful constraint for testing the predictions of future cosmological simulations.

6.5 Other Assumptions

Clearly the jet scalings we adopt are one of the biggest factors affecting our results, however the calculation involves a number of other assumptions which we now discuss.

The Fermi band flux from a FSRQ is dominated by Compton up-scattering of seed photons from the BLR. We have approximated the BLR as a spherical shell of radius R_{BLR} centred on the black hole. However studies of line profiles have shown that the BLR geometry may be more similar to that of a disc wind (Kollatschny & Zetzl, 2013), in which case the material should have a more flattened geometry. The BLR clouds which are reprocessing the disc luminosity are then clumps of material thrown off the disc, possibly as part of a UV line driven disc wind. A more flattened geometry means the BLR is behind the jet emission region not in front, in which case the BLR seed photons will be deboosted and redshifted in the jet frame. In order to produce the bright Compton emission of FSRQs requires a source of seed photons in front of the jet emission region, so that the seed photons appear maximally boosted with respect to the jet electrons. If the BLR does have a flattened disc wind geometry then the seed photons we attribute to the BLR may instead be provided by the narrow line region (NLR). The BLR would still provide a source of seed photons but they are likely to scale differently with $\dot{m}$ and M.

We assume R_{BLR} (or R_{NLR}) scales as $L_d^{1/2}$, but for the case of a funnel-like disc wind both the launch radius, the opening angle of the wind and the vertical height material travels before bending outwards are all likely to scale with L_d in a more complex way (Elvis, 2000).

We have also assumed that a fixed fraction of L_d is reprocessed by the BLR (/NLR) and torus. Again this may not be the case. G10 find from spectral fitting of a sample of FSRQs that the fraction does vary slightly, although not by much. Consequently we consider this to be a small effect less important in our statistical sample.

We have used the mean FSRQ spectrum of G10 as our model spectrum from which to scale, however FSRQs are highly variable. During flaring the Fermi flux can increase by more than an order of magnitude. As a result, distant FSRQs that would not normally be detected may become visible. This would extend the tail of the redshift distribution out to higher redshifts than otherwise expected. In only modelling the typical FSRQ emission, such events have not been included in our simulation.

We have also assumed that a FSRQ jet is produced for the entire time a black hole is accreting with $\dot{m} > 10^{-2}$. If instead the jet follows a duty cycle and is only produced for a fraction of that time then this will reduce the number of Fermi detected FSRQs. We overpredict the number of Fermi visible FSRQs by 2 orders of magnitude assuming every black hole with $\dot{m} > 10^{-2}$ produces a FSRQ jet. If each of these black holes only produces a FSRQ jet 100th of the time then we could match the observed numbers without needing any limits on the spin of the black hole. However FSRQs should be the aligned analogues of FRII sources (Padovani & Urry, 1992). The large scale radio lobes of FRII sources indicate that the jet producing them must be persistent, since the time taken to produce such large scale structures is of the order of megayears. One explanation may then be that the large scale structure is produced by a persistent slower jet, while there is a fast central spine with $\Gamma = 13$, which appears as a FSRQ when viewed head on and is intermittent.

One final caveat to our results relates to the Doppler boosting of the jet emission.

We assume the synchrotron, SSC and EC radiation is boosted by a factor δ^3 in luminosity and shifted up by a factor of δ in frequency, where $\delta = (\Gamma - \cos\theta\sqrt{\Gamma^2 - 1})^{-1}$. This is the approach adopted by G10 and crucially allows us to match our model spectra to their observations. However Dermer & Menon (2009) point out that the EC radiation should in fact be boosted by a factor of δ^4 and shifted up in frequency by δ^2, due to the anisotropy of the BLR seed photon field in the jet frame. As a result the beaming cone for EC jet emission is narrower than that of the synchrotron and SSC emission. Since our simulations involve a range of R_{diss}, including cases where the dominant external seed photons move behind the emission region where a δ^4 boosting factor is not appropriate, we keep to the simpler model of G10 where synchrotron, SSC and EC are all boosted by δ^3. We note that an extra factor of δ would increase the visibility of very highly aligned high redshift sources and would improve the Compton dominance of highly aligned small mass objects. However, including this extra factor does not affect the main conclusions of our work: that in order to reproduce the number of Fermi detections not every black hole with $\dot{m} > 0.01$ can produce a FSRQ jet, and if the requirement to produce one is high spin then there must be more high spin black holes at $2 < z < 3$ than chaotic accretion models predict, suggesting accretion at early times may be more ordered than at late times.

6.6 Comparing FSRQ and BL Lac Jets

FSRQs and BL Lacs are typically understood as representing the two ends of the 'blazar sequence'. The transition from low power BL Lac to high power FSRQ can be understood in terms of increasing accretion rate onto the central black hole. The dimmest BL Lacs, produced by the lowest accretion rate black holes ($\dot{m} < 10^{-3}$), appear as high peaked BL Lacs (HBLs). Their low accretion rates mean lower magnetic fields and injected electron powers, which result in less cooling so the synchrotron and SSC emission peak at high frequencies. As accretion rate increases, B and P_{rel} increase, the amount of cooling increases so the electron distribution cools down to lower Lorentz factors and the observed synchrotron and SSC spectra peak

at lower frequencies. Increasing $\dot{m}$ and the corresponding increase in B and P_{rel} switch the observed spectrum from a HBL to a low peaked BL Lac (LBL). As $\dot{m}$ becomes greater than 10^{-2}, the accretion flow around the black hole switches to a radiatively efficient disc, effectively turning on external sources of seed photons and the jet stops being a BL Lac and appears as a FSRQ. In this picture, the jet is the same in both cases, the only difference is the power input (B and P_{rel}) and the presence or absence of external seed photons, both of which are linked by a dependence on the accretion rate. However in reality this is not the case. There are fundamental differences between the two types of jet.

Comparison of the mean injected electron distribution parameters found by G10 for FSRQs and BL Lacs shows that γ_{max} and γ_b are much larger for BL Lacs ($\gamma_{max} \sim 10^5$ for BL Lacs compared to 10^3 for FSRQs and $\gamma_b \sim 10^4$ compared to 10^2). The difference between BL Lacs and FSRQs is not simply that the electrons have a different seed photon field to cool off. The accelerated electron distribution is intrinsically different in FSRQs compared to BL Lacs. This suggests there is some difference in the way electrons are accelerated, presumably by shocks, in FSRQ jets compared to BL Lacs.

A more fundamental difference is in jet opening angle (ϕ). Here and in the previous chapter we have used $\phi = 0.1$, which is typically assumed for calculating blazar spectra. However Krause et al. (2012) have shown that ϕ should be larger for BL Lac jets. They find from hydrodynamic simulations that the jet opening angle sets the large scale morphology of the jet, with FRII jets (corresponding to misaligned FSRQs) being produced for opening angles $< 24°$ ($= 0.4 \, \mathrm{rad}$) and FRI morphologies (corresponding to BL Lacs) being produced for larger opening angles. Since ϕ relates Z_{diss} and R_{diss}, this means that the same size emission region will be located at smaller Z_{diss} for a larger opening angle. Since the calculation of BL Lac spectra does not involve any external seed photons, the only change as a result of accounting for a larger opening angle will be that the Z_{diss} derived from fitting a given BL Lac spectrum will be smaller.

A related factor is that the mean BL Lac BLF is slightly larger than the mean FSRQ BLF (15 compared to 13, G10). The BLF of the jet should influence where

the dissipation region is, if it corresponds to a standing shock at the base of the jet. For larger Γ material will travel further before shocking. The discontinuity in both opening angle and BLF suggests that R_{diss} and Z_{diss} should not scale continuously between FSRQs and BL Lacs.

This break in the jet scalings is emphasised by the fact FSRQs do not appear to follow standard jet scalings, while BL Lacs do. Standard jet scalings (where all size-scales scale with mass and the power injected into relativistic electrons and magnetic fields is a constant fraction of the accretion power) reproduce the observed transition from HBL to LBL with increasing accretion rate in BL Lacs. The alternative jet scalings (where $R_{diss} \propto \dot{m}M$, which give a slightly better approximation to the behaviour of FSRQs) cannot produce this transition. The peak of the synchrotron spectrum from a BL Lac is determined by the minimum electron Lorentz factor that can cool in one light crossing time, γ_{cool}. $\gamma_{cool} \propto (R_{diss}U_{seed})^{-1}$, where the seed photons are provided by the magnetic field $U_B \propto \dot{m}M/R_{diss}^2$, hence $\gamma_{cool} \propto R_{diss}/\dot{m}M$. Following standard jet scalings, $R_{diss} \propto M$ and $\gamma_{cool} \propto \dot{m}^{-1}$, allowing γ_{cool} to decrease as $\dot{m}$ increases, which produces the transition from HBL to LBL. If $R_{diss} \propto \dot{m}M$ then γ_{cool} stays constant and there is no transition from HBL to LBL. Clearly BL Lacs cannot follow the alternative jet scalings, while FSRQs seem not to follow the standard scalings.

6.7 Conclusions

We have combined models of FSRQ spectra together with prescriptions for how they should scale with mass and accretion rate and the number densities of black holes from cosmological simulations to predict the number of FSRQs that should be observed by Fermi. If we assume all black holes accreting with $\dot{m} > 10^{-2}$ produce a FSRQ jet our simulation overpredicts the number of Fermi detected FSRQs by two orders of magnitude. If we restrict the production of FSRQ jets to high spin black holes ($a > 0.77$) we can reproduce the observed numbers, however the predicted redshift distribution does not extend to as high redshift ($2 < z < 3$) as is observed.

An important factor in our simulations are the scaling relations we use to predict

the spectra from FSRQs of different masses and accretion rates. We test these by comparing our scaled models with gamma-ray loud NLS1s, which should be scaled down versions of the more massive FSRQs. We find that standard scaling relations (allowing all size-scales to scale with mass and assuming the power injected into relativistic electrons and magnetic fields is a constant fraction of the accretion power) significantly underpredict the amount of Fermi flux small mass FSRQs ($\lesssim 10^8 \, M_\odot$) should produce. This is because they predict small mass FSRQs should have much higher magnetic fields, while in reality the magnetic fields of γNLS1s are comparable to those of normal mass FSRQs. High magnetic fields produce synchrotron and SSC dominated spectra, with low Compton dominance. This is in contrast to the predominantly EC, highly Compton dominant and consequently much more Fermi visible spectra that are actually observed from γNLS1s.

In light of this we reran our simulations using an alternative jet scaling, where we allowed the size-scales of the jet to scale with both mass *and* accretion rate. If the jet emission region is a standing shock caused by some pressure balance between high pressure material outflowing into less dense material this could plausibly depend on accretion rate, with pressure equilibrium occurring further out and giving larger size-scales for higher accretion rates. This allows the higher accretion rates of small black holes to compensate for their smaller masses, giving magnetic fields for γNLS1s more similar to those of normal mass FSRQs. Increasing the visibility of low mass FSRQs by adopting these alternative scaling relations only increases the predicted population of FSRQs, making the need for some other factor, such as a limit on black hole spin, even more necessary to reduce the predicted numbers. However despite increasing the visibility of low mass FSRQs, which should dominate at high redshifts, we still cannot reproduce the observed population of FSRQs at high redshifts after imposing a spin cut. This highlights a limit on the ability of the cosmological simulations to track the evolution of black hole spin. We suggest this may be due to the cosmological simulation we use only assuming chaotic accretion when tracking black hole spin evolution, while in reality there may be a transition from prolonged to chaotic accretion with decreasing redshift.

Chapter 7

A Physical Model for the X-ray Time Lags of Narrow Line Seyfert Type I Active Galactic Nuclei

7.1 Introduction

In the previous chapter we investigated gamma-ray loud Narrow Line Seyfert 1s as scaled down versions of the more massive flat spectrum radio quasars. If the powerful jets in these objects are the result of a highly spinning central black hole, this suggests the ordinary radio quiet Narrow Line Seyfert 1s (NLS1s) should host low spin black holes. In this chapter we use both the spectral and timing properties of radio quiet NLS1s to investigate whether this is the case.

NLS1s are small mass, high mass accretion rate Active Galactic Nuclei (AGN). One of the defining characteristics of NLS1s is their strong soft X-ray emission (Boller et al., 1996). While the disc spectra of most AGN peak in the UV, the relatively small black hole masses and high accretion rates of NLS1s mean that the disc emission in these objects can extend into the soft X-rays. However a simple combination of disc emission plus coronal power law is not sufficient to explain the soft X-ray spectra of NLS1s. The disc rolls over very rapidly past its maximum temperature, whereas the data show a much more gradual decline. An additional component is required to explain the excess of emission at soft energies.

Since the soft X-ray excess has a smooth shape, with no clear atomic signatures, it can be well fit by an additional Comptonised component. However, the temperature of this is remarkably constant at $\sim 0.2\,\mathrm{keV}$ across a range in mass and mass accretion rates, requiring some unknown fine tuning mechanism (Czerny et al., 2003; Gierliński & Done, 2004; Porquet et al., 2004). An alternative model is that the soft X-ray excess is produced by reflection from a partially ionised disc. However this also requires fine tuning of the ionisation state of the disc in order to consistently generate sufficient change in opacity from partially ionised material at $\sim 0.7\,\mathrm{keV}$ (Done & Nayakshin, 2007) and extreme relativistic effects are required to smear out the resulting emission lines, as these are not seen in the data (Fabian et al., 2004; Crummy et al., 2006). As a result the reflection must arise from the innermost disc radii around a highly spinning black hole. Both models for the soft X-ray excess — the reflection and the additional Compton component — can give equally good fits to the X-ray spectra in the $0.3 - 10\,\mathrm{keV}$ CCD bandpass.

Variability gives another way to distinguish between these two models. The discovery of a lagged signal in the X-ray light curves of AGN has led to a breakthrough in our ability to probe the structure of the emission region on the smallest scales. The simplest interpretation of this lag is that it is due to the light travel time delayed response of the disc to X-ray illumination (Fabian et al., 2009). However, the lag behaviour is complex; at long timescales (low frequency) the hard X-rays lag behind the soft, while at high frequencies the opposite is true (Papadakis et al., 2001; Vaughan et al., 2003; McHardy et al., 2004; Fabian et al., 2009; Zoghbi et al., 2011; Emmanoulopoulos et al., 2011; De Marco et al., 2013). This is interpreted as the interplay of two different processes, with propagation of fluctuations on longer timescales giving rise to the hard lags, while reverberation from the disc takes over on shorter timescales. Reflection from the disc produces an iron line and Compton hump above $10\,\mathrm{keV}$, but can also contribute to the soft X-ray band if the disc is partially ionised, producing a soft lag at high frequencies (Fabian et al., 2009; Zoghbi et al., 2011; Cackett et al., 2013).

However, an alternative explanation for the soft lags at high frequency was suggested by Alston et al. (2014) (hereafter ADV14, also Zoghbi et al., 2011). They

noted that another feature of X-ray illumination of the disc is that the non-reflected flux is reprocessed as thermal emission. This provides another component which reverberates behind the hard X-rays in the same way as reflection, and can contribute to the soft X-ray bandpass even if reflection itself does not produce much soft X-ray flux.

Here we build a fully self consistent spectral and timing model, starting from fluctuations generated intrinsically in the accretion flow which then propagate through the system and including reflection/reprocessing with realistic transfer functions. We focus on explaining the data from the NLS1 PG1244+026, where a range of spectral and timing properties favour an additional Compton component. We explore the lags in this model, and then compare them with the lags expected from reflection models for the soft X-ray excess, including the one specifically for this object (Kara et al., 2014, hereafter K14).

7.2 Spectral Decomposition

We use the long (120 ks) XMM-Newton observation of PG1244+026 (OBS ID: 0675320101, as studied by Jin et al., 2013, hereafter J13; ADV14 and K14). We use a similar spectral decomposition to J13 which assumes that the spectrum is composed of three components: a (colour temperature corrected) blackbody (BB) disc, a low temperature optically thick Compton component to describe the soft excess, and a second optically thin Compton component to describe the hard X-ray power law. We model these together using the publicly available model OPTXAGNF, which assumes that these three components are all powered by the accretion flow, so the luminosity to power the soft excess and power law sets the truncation radius R_{cor} of the standard BB disc emission (Done et al., 2012). This model assumes that the soft excess arises at radii $< R_{cor}$, and a plausible origin is that this represents the inner regions of the standard disc, but that the emission here does not completely thermalise, perhaps because of a larger scale-height (and hence lower density) expected if there are strong winds from the disc (Done et al., 2012).

We also include moderately ionised reflection of the power law off the disc in

Figure 7.1: a). Spectral decomposition for PG1244+026: disc (red), soft excess (green), coronal power law (blue), reflection (magenta), total (black). Data points show time averaged spectrum (OBS ID: 0675320101). b). Power spectrum of intrinsic fluctuations in each model component: disc (red), soft excess (green), coronal power law (blue).

order to account for reflection features at $6 - 7\,\mathrm{keV}$ due to iron. J13 showed two extreme fits, one where the disc was mostly neutral, the other where it was highly ionised. ADV14 used these to explore the shape of the lag spectra, but here we use the more likely (and better fitting) moderate ionisation reflection model. Fig.7.1a shows the model components compared to the data, with full parameters detailed in Table 7.1. We fix the inclination of the reflector to $30°$ and match the seed photon temperature for the high energy Comptonisation to the temperature of the soft excess.

7.3   Time Dependent Model

Ideally, the full (energy and frequency dependent) cross-spectrum should be directly fit to constrain the intrinsic components — both their energy spectra and variability properties. Since this inverse problem has not yet been solved, we do forward fitting instead, using the spectral components from model fitting, and estimating their variability from a combination of the observed power spectra and lag-frequency spectra. We then calculate the various spectral-timing properties, and qualitatively compare these to the observations.

There have been many timing studies of PG1244+026. In order to maximise

Component	Parameter	Value
Galactic absorption	N_h ($10^{22}\,\mathrm{cm}^{-2}$)	0.019
Intrinsic absorption	N_h ($10^{22}\,\mathrm{cm}^{-2}$)	6.1×10^{-11}
	redshift	0.048
BBODY	kT (keV)	0.062
	norm	2.0×10^{-4}
NTHCOMP	Γ	2.4
	kT_e (keV)	100
	norm	1.9×10^{-3}
KDBLUR	index	3.0
	r_{in} (R_g)	11
	i (°)	30
RFXCONV	relative refl norm	-0.65
	Z_{Fe}	1.0
	$\log(x_i)$	1.3
COMPTT	kT_e (keV)	0.15
	τ	190
	norm	0.073

Table 7.1: Parameters for the spectral model shown in Fig.7.1a, with a separate optically thick Comptonisation component producing the soft excess emission: WABS*ZWABS(BBODY+NTHCOMP+KDBLUR*RFXCONV*NTHCOMP+COMPTT).

diagnostic power, we choose to use the power spectra and covariance measured by J13 and the coherence, lag-frequency and lag-energy spectra measured by ADV14. We note that ADV14 used slightly different soft and hard energy bands ($0.3 - 0.7\,\mathrm{keV}$ and $1.2 - 4.0\,\mathrm{keV}$) compared to J13 ($0.3 - 1.0\,\mathrm{keV}$ and $2 - 10\,\mathrm{keV}$). Therefore when calculating power spectra we use the hard and soft bands of J13 and use the energy bands of ADV14 when calculating lags and coherence, in order to match to the data. The equations for these calculations are given in the Appendix. Throughout the figures, points show our model results, while the shaded regions illustrate the range of the error on the values measured for PG1244+026.

We use the spectral components derived from the previous section — BB disc, soft excess, power law and reflection — to determine the relative contributions of each component to a given energy band. This then determines the power spectrum of that energy band. Since the observed power spectra of the hard and soft bands are different, this implies different components dominate each band and these components have fluctuations at different characteristic frequencies.

7.3.1 Intrinsic Fluctuations: No Propagation

We assume fluctuations are generated intrinsically within the BB disc, soft excess and corona. We assume the power spectrum of these fluctuations takes the form of a Lorentzian centred around some characteristic frequency — f_d, f_s or f_p, respectively. This characteristic frequency will be different for each component and can be associated with a viscous timescale in analogy with propagating fluctuation models in black hole binaries (BHBs; Kotov et al., 2001; Arévalo & Uttley, 2006; Ingram & Done, 2011; Ingram & Done, 2012a). Since our physical model describes an outer disc, inner soft excess and corona, the relevant radii for each component decreases from disc to soft excess to corona (and scale-height probably also increases), so the characteristic frequency of fluctuations should increase.

This increase in characteristic frequency can be seen in the power spectra of PG1244+026 (J13, ADV14). The soft band power spectra show more power at low frequencies ($f < 10^{-4}$ Hz), whilst the hard band power spectra show much more high frequency power ($f > 10^{-4}$ Hz). The corresponding lag-frequency spectrum (ADV14) shows the soft band leading the hard band at low frequencies and lagging it at high frequencies. This implies the low frequency fluctuations must be generated in the soft band components (or at least pass through them) before reaching the hard power law component, whilst the hard power law component must be the source of the high frequency fluctuations, which then reverberate in the soft band to produce soft lags at high frequencies. The switch between soft leads and soft lags occurs at $\sim 10^{-4}$ Hz in PG1244+026, implying the characteristic frequency of intrinsic fluctuations in the soft band components (disc and soft excess) are at frequencies below 10^{-4} Hz, while the hard power law generates frequencies above 10^{-4} Hz.

Motivated by the observed power spectra and lag-frequency spectrum, we choose $f_d = 3 \times 10^{-5}$ Hz as the characteristic frequency of the fluctuations in the disc, $f_s = 1 \times 10^{-4}$ Hz for the soft excess, and use two Lorentzians ($f_{p,1} = 3 \times 10^{-4}$ Hz and $f_{p,2} = 1 \times 10^{-3}$ Hz) to describe the breadth of the high frequency variability in the coronal power law. We use the method of Timmer & Koenig (1995) to generate fluctuations in each component, $\dot{M}_{d,s,p}(t)$. Each $\dot{M}_{d,s,p}(t)$ is normalised to a mean of unity and fractional variability $\sigma/I = F_{d,s,p}$. This represents fluctuations of
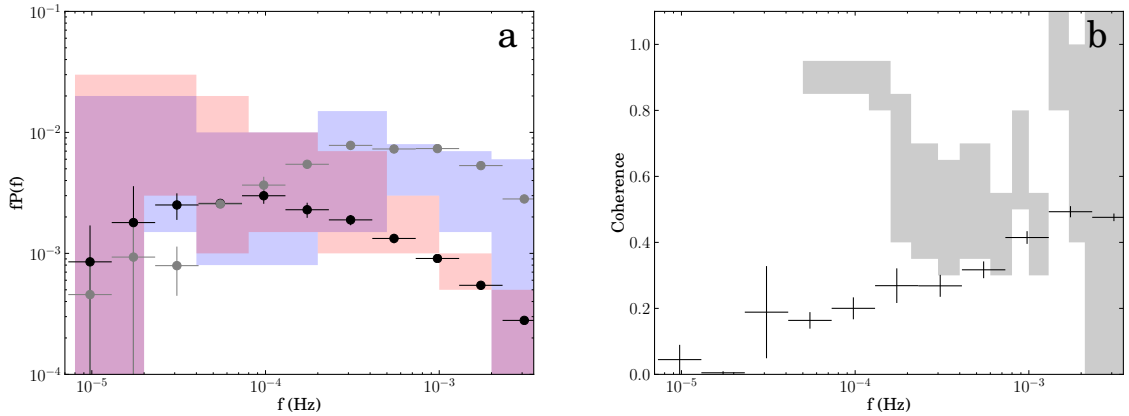
Figure 7.2: Model results with only intrinsic fluctuations, no propagation between components. a). Soft band ($0.3 - 1\,\mathrm{keV}$, black) and hard band ($2 - 10\,\mathrm{keV}$, grey) model power spectra. Red (blue) regions show range of error on soft (hard) band power spectra measured by J13 for PG1244+026. b). Coherence between hard ($1.2 - 4\,\mathrm{keV}$) and soft ($0.3 - 0.7\,\mathrm{keV}$) band model light curves. Grey regions show range of error on coherence measured by ADV14 for PG1244+026.

the component around its time averaged value determined from the fit to the time averaged spectrum.

For each time, t, the time averaged spectral components are multiplied by their fluctuations, $\dot{M}_{d,s,p}(t)$, at that time and summed to give the total spectrum. In order to allow for red noise leakage, we calculate the total spectrum for 2^{20} timesteps with $dt = 100\,\mathrm{s}$ and then split it into ten $102.4\,\mathrm{ks}$ segments. For each segment we calculate the power spectrum in the hard ($2-10\,\mathrm{keV}$) and soft ($0.3-1\,\mathrm{keV}$) bands and average the ten power spectra to get a mean power spectrum.

Fig.7.1b shows the power spectrum of the fluctuations in each model component (red: disc, green: soft excess, blue: power law). Fig.7.2a shows the resulting hard and soft band power spectra in grey and black respectively, analogous to those in J13. The red and blue regions show the error on the soft and hard band power spectra measured by J13 for PG1244+026. The observed hard and soft power spectra are the same within a factor of 2 below $10^{-4}\,\mathrm{Hz}$ and a factor of 10 different at $10^{-3}\,\mathrm{Hz}$, with much less high frequency power in the soft band. In our model, the power in the soft band is much less than that in the individual disc and soft excess components which contribute to it, because the power in both components is incoherent, so summing them reduces the total power (Arévalo & Uttley, 2006). In contrast, the hard band

is dominated by the power law with little dilution from other components, so the hard band power spectrum is very similar to the power spectrum of the power law component.

Fig.7.2b shows the coherence between the hard and soft band model light curves, where 1 is perfect coherence and 0 is incoherence (Vaughan & Nowak, 1997; Nowak et al., 1999). The components themselves are incoherent, so the coherence between the soft excess and power law is 0. Although the hard band is dominated by the power law and the soft band by the soft excess, the soft band also contains a non negligible contribution from the power law. Thus some fraction of the soft band light curve is correlated with the hard band, due to the presence of the power law in both bands. Consequently the coherence at frequencies produced by the power law is equal to the fractional contribution of the power law to the soft band.

The grey regions in Fig.7.2b show the range of the error on the coherence measured by ADV14 for PG1244+026. The measured coherence clearly shows the opposite trend, with the coherence being highest for low frequencies and dropping between $10^{-4} - 10^{-3}$ Hz. In order to replicate this, we must allow fluctuations to propagate inwards from the slowly varying soft components to the faster varying coronal power law which dominates the hard band.

7.3.2 Propagating Fluctuations

We now allow the fluctuations to propagate inwards from the disc into the soft excess and then to the corona. We choose this scenario, rather than allowing fluctuations to pass directly from the disc into the corona, since J13 showed that the soft excess rather than the disc provides the seed photons for Comptonisation in the corona. This suggests that the soft excess and corona are more closely associated. Perhaps the disc and corona are spatially separated by the soft excess. Or more likely in PG1244+026 (where the inner disc radius derived from the spectral fit is $\sim 12\,R_g$), the soft excess lies below the corona, with the corona formed from material that has evaporated from the optically thick soft excess region beneath it. A moderate scale-height corona extended above the soft excess would then receive a greater seed photon flux from the soft excess than from the disc. This is in

agreement with the characteristic frequencies of coronal fluctuations (which cannot vary on timescales faster than the source light crossing time), with the highest frequency fluctuations (3×10^{-3} Hz) being generated at $\sim 6\, R_g$ and the lower frequency fluctuations (10^{-3} Hz) being generated at larger radii $\sim 10\, R_g$.

In this scenario the fluctuations in the soft excess therefore consist of the intrinsic fluctuations ($\dot{M}_{s,int}(t)$) generated in the soft excess, modulated by fluctuations that have propagated inwards from the disc:

$$\dot{M}_s(t) = \dot{M}_{s,int}(t)\dot{M}_{d,int}(t - t_{lag,s}) \qquad (7.3.1)$$

where $t_{lag,s}$ is the time delay for propagation of fluctuations from the disc to the soft excess. We set $t_{lag,s} = 1000$ s, from the results of ADV14. As well as lagging, we also smooth the propagating fluctuations on this timescale using a sliding boxcar of width $t_{lag,s}$. We note that this is much shorter than the expected lag times for propagation, which should be related to the viscous timescale if NLS1s behaved as scaled up BHBs (Ingram & Done, 2011), i.e. $t_{lag,s} \sim \frac{1}{f_{visc,d}} - \frac{1}{f_{visc,s}} \sim 10^4$ s. The light travel time between $20 - 12\, R_g$ is ~ 500 s for PG1244+026 ($M_{BH} \sim 10^7\, M_\odot$), so this could mean $t_{lag,s}$ represents a light travel time from the disc to the soft excess, i.e. that the disc provides the seed photons for the optically thick Compton scattering in the soft excess.

Similarly the fluctuations in the corona consist of the intrinsic coronal fluctuations ($\dot{M}_{p,int}(t)$) modulated by the smoothed and lagged fluctuations propagating inwards from the soft excess:

$$\dot{M}_p(t) = \dot{M}_{p,int}(t)\dot{M}_s(t - t_{lag,p}) \qquad (7.3.2)$$

where $t_{lag,p}$ is the time delay for propagation of fluctuations from the soft excess to the corona. We set $t_{lag,p} = 600$ s, again guided by ADV14, and again note that this is much closer to a light travel time than to a viscous propagation time.

Fig.7.3a shows the power spectra of the individual model components and of the hard and soft band model light curves respectively. Due to accumulation of fluctuations, the hard band power spectrum now has more low frequency power
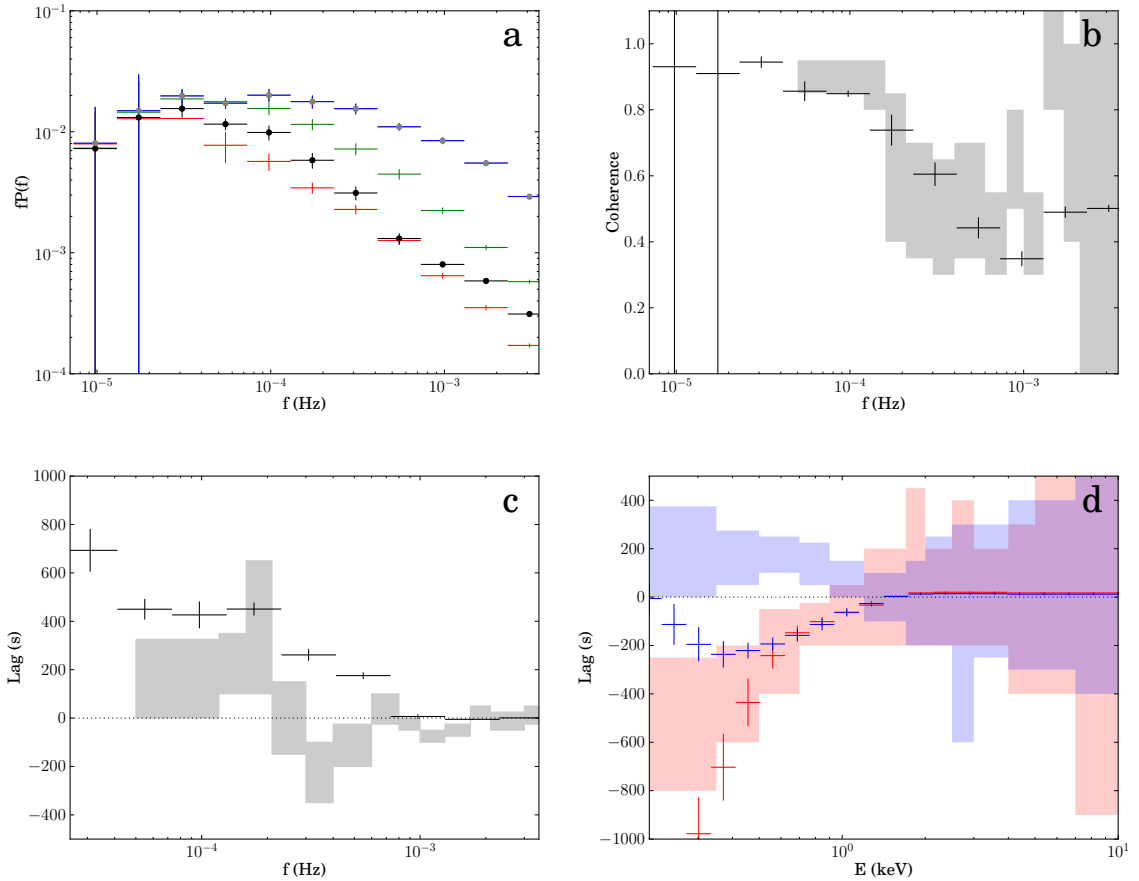
Figure 7.3: Model results allowing fluctuations to propagate from disc to soft excess ($t_{lag,s} = 1000\,\mathrm{s}$) to coronal power law ($t_{lag,p} = 600\,\mathrm{s}$). a). Power spectra: disc (red), soft excess (green), coronal power law (blue), soft band ($0.3 - 1.0\,\mathrm{keV}$, black), hard band ($2-10\,\mathrm{keV}$, grey). b). Coherence between hard ($1.2-4\,\mathrm{keV}$) and soft ($0.3-0.7\,\mathrm{keV}$) bands. c). Lag-frequency spectrum between hard ($1.2-4\,\mathrm{keV}$) and soft ($0.3-0.7\,\mathrm{keV}$) bands. d). Lag-energy spectrum calculated using $1.2-4\,\mathrm{keV}$ reference band: low frequency lag ($2.3 \times 10^{-5}-7.3\times10^{-5}\,\mathrm{Hz}$), red points; high frequency lag ($2.3\times10^{-4}-7.3\times10^{-4}\,\mathrm{Hz}$), blue points. Shaded regions show range of error on values measured by ADV14 for PG1244+026.

than the soft band. The total power in the soft band has also increased, because propagation of fluctuations means more of the power in the soft excess is correlated with the disc power. We note that this model is indeed giving hard and soft power spectra which are within a factor of 2 at 10^{-4} Hz yet different by a factor of 10 at 10^{-3} Hz, as required by the data. Fig.7.3b shows the coherence between the hard and soft band model light curves, which is now highest at low frequencies due to the inward propagation of the slower outer fluctuations. This gives a much better match to the data. The coherence drops above 10^{-4} Hz since this is the maximum typical frequency of fluctuations generated in the soft excess that can propagate down to the coronal power law, which dominates the hard band.

In Fig.7.3c and d we show the lag-frequency and lag-energy spectra from the model. We define the lag as a function of frequency as $\text{lag}(f) = \arg[C(f)]/(2\pi f)$, where $C(f) = S^*(f)H(f)$ is the complex valued cross spectrum, and H is the Fourier transform of the hard band light curve and S is the Fourier transform of the soft band light curve (Vaughan & Nowak, 1997; Nowak et al., 1999). A positive time lag therefore corresponds to the soft band leading the hard, and negative lags to the soft band lagging the hard.

At all frequencies below $\lesssim 6 \times 10^{-4}$ Hz, the model soft band leads the hard, because these are the frequencies that are intrinsically generated in the disc and soft excess, which propagate down to the corona.

This is shown more clearly in the lag-energy spectrum (Fig.7.3d). For each energy bin we calculate the cross spectrum of the light curve of that energy bin with a hard reference band light curve (minus the energy bin light curve if the energy bin lies within the $1.2-4$ keV reference band). We plot the value of the lag at low frequencies $(2.3 \times 10^{-5} - 7.3 \times 10^{-5}$ Hz$)$ in red in Fig.7.3d. A negative lag now represents the energy bin leading the hard reference band. The model light curves at energies dominated by disc and soft excess emission $(< 1$ keV$)$ lead the hard reference band, with a lag that decreases as the energy of the bin increases. This matches the disc contribution decreasing with increasing energy and being replaced by the soft excess, which has a shorter time delay between it and the coronal emission, until ~ 1 keV by which point the coronal emission begins to dominate the total spectrum and the

low frequency lag reduces to 0.

In blue in Fig.7.3d we plot the value of the model lags at high frequencies ($2.3 \times 10^{-4} - 7.3 \times 10^{-4}$ Hz), corresponding to frequencies generated in the soft excess. These show a negative lag at low energies. The energy spectrum of the lag matches the energy spectrum of the soft excess (cf. Fig.7.1a). The lag is a maximum for energy bins corresponding to the peak of the soft excess (~ 0.3 keV). It decreases towards higher energies as the soft excess emission is gradually replaced by power law emission. It decreases towards lower energies because the soft excess emission is replaced by the disc, which doesn't generate strong fluctuations at high frequencies. However, even at its maximum the measured model lag is only ~ 200 s, whilst the actual lag time input into the model for propagation from the soft excess to the power law is 600 s. This is a result of dilution. The soft band contains contributions from the disc and power law in addition to the soft excess component. The small contribution from the power law (which has zero lag with respect to the power law dominated hard band) reduces the net lag that is actually measured. Since we can only ever measure the *net lag*, the lag measured between two bands should never be taken as representing a 'true' delay between components, since it will always be a combination of the lags from each component contributing to that band weighted by their contributions to the total flux in the band.

In Fig.7.3c and d we also show the lag values measured by ADV14 for PG1244+026; the shaded regions show the range of the errors. Clearly the propagating model has no way of producing the observed soft lags (negative grey regions and positive blue regions). These are generally attributed to reverberation but there is an alternative way to produce spectral lags/leads and that is by spectral pivoting. We explore this first below.

7.3.3 The Effect of Power Law Spectral Pivoting

So far we have taken a simple approach of multiplying the time averaged spectral components by their fluctuations. This is justified for components that retain their spectral shape and only change in normalisation with addition of fluctuations, i.e. the disc and soft excess. However the power law has a spectral slope set by the

balance of heating and cooling in the corona, so its shape should be affected by fluctuations in the corona and also in its seed photons, which come from the soft excess (J13). In order to allow for this, we use the Comptonisation code EQPAIR (Coppi, 1999) to replace the simple power law fit from the time averaged spectrum. This calculates the spectrum of the Comptonised emission for a spherical region of size r, optical depth τ, seed photon temperature kT_{seed}, seed photon power l_s and heating power to the electrons l_h. We fix the size of the emission region to $10\,R_g$ and match the seed photon temperature to that of the soft excess ($0.2\,\mathrm{keV}$). We set the mean values of the optical depth, seed photon power and ratio of heating to cooling power to $\tau_0 = 1$, $l_{s,0} = 1000$ and $l_{h,0}/l_{s,0} = 1$ in order to reproduce the slope of the time averaged power law spectrum, and then allow them to vary according to the coronal and soft excess fluctuations as:

$$\tau(t) = \tau_0 \dot{M}_{p,int}(t) \tag{7.3.3}$$

$$l_s(t) = l_{s,0} \dot{M}_s(t - t_{lag,p}) \tag{7.3.4}$$

$$\frac{l_h(t)}{l_s(t)} = \frac{l_{h,0}(t)}{l_{s,0}(t)} \frac{\dot{M}_{p,int}(t)}{\dot{M}_s(t - t_{lag,p})} \tag{7.3.5}$$

where $\dot{M}_{p,int}$ are the intrinsic fluctuations generated in the corona as before. $\dot{M}_p$ therefore consists of a combination of the intrinsic fluctuations in coronal power and the fluctuations in seed photon flux from the soft excess. This assumes that $t_{lag,p}$ is not the lag time for fluctuations to physically propagate from the soft excess into the corona, but instead the light travel time for seed photons. This has clear motivation from the similarity of timescales between the propagation and reprocessing, as determined by ADV14. It is not unlikely that there would be both transmission of fluctuations through varying seed photon flux and also through physical propagation of accretion rate fluctuations, with a slightly longer lag time. However we can only measure one net lag, and since it is closer to a light travel time that suggests seed photon propagation makes a strong contribution. Therefore for simplicity we assume all the transmission of fluctuations is via the seed photon flux in this model. We normalise the resulting spectrum to have a total power equal to the input electron heating power plus the seed photon power. This approach allows us to account

Figure 7.4: As in fig 3, but now for model with propagation of fluctuations from disc to soft excess ($t_{lag,s} = 1000\,\mathrm{s}$) and spectral pivoting of the power law assuming its seed photons are provided by the soft excess ($t_{lag,seed} = t_{lag,p} = 600\,\mathrm{s}$).

for spectral pivoting of the power law resulting from the relationship between the corona and its seed photons from the soft excess.

Fig.7.4 shows the resulting power spectra, coherence, lag-frequency and lag-energy spectra. The pivoting model power spectra, coherence and lag-frequency spectra are very similar to the non-pivoting case. However the lag-energy spectrum at low frequencies has changed (Fig.7.4d, red points). There is a negative lag at low energies, tending to zero at $\sim 1\,\mathrm{keV}$. But instead of remaining at zero, as in the non-pivoting case, a negative lag returns above $\sim 1\,\mathrm{keV}$, increasing in strength with increasing energy. This is due to the pivoting of the power law spectrum.

Poutanen & Fabian (1999) showed how such pivoting from spectral evolution could cause time delays between hard and soft photons. They used a model for short timescale variability in BHB systems where a magnetic reconnection event

caused particle acceleration above the disc, heating the electrons. These are cooled by Compton cooling off the copious seed photons from the disc, resulting in a soft spectrum. If the reconnection event expands upwards then the flux of seed photons drops and the spectrum hardens. This spectral change from soft to hard gives a hard lag as seen in the data, but its origin is from the spectral evolution of a single region rather than delays between spectra emitted from different regions.

In our pivoting model, the coronal spectrum is initially hard and an increase in seed photons from the soft excess causes the coronal emission to soften. Hence this appears as a soft lag (i.e. harder energy bins increasingly 'lead') in the lag-energy spectrum. The hard reference band extends from $1.2 - 4 \, \text{keV}$. Above $4 \, \text{keV}$, the 'lead' increases with energy as the effect of spectral pivoting increases. However, because the softening occurs on timescales associated with fluctuations in the soft excess this does not cause soft lags at high frequencies but at low frequencies. The leads are also at hard energies, whilst the data require soft leads. As a result, the high frequency lag-energy spectrum from the pivoting model looks the same as in the non-pivoting case (cf. Fig.7.3d).

The observed low frequency lag-energy spectrum of PG1244+026 does not appear to show this characteristic pattern of soft lags at high energies (ADV14). In Fig.7.4d we also show the range of the low frequency lag measured by ADV14 shaded in pink. This is more consistent with remaining at zero lag, although the errors are large. This suggests the power law component in PG1244+026 does not pivot, but merely changes in normalisation like the other components. To produce a change in normalisation without pivoting in the spectrum requires l_h and l_s to change together. This suggests a situation where the corona is connected to the soft excess so that a mass accretion rate fluctuation in the soft excess can increases the soft flux, but can also propagate into the corona (perhaps through evaporation) to produce a correlated increase in the power in hot electrons (l_h).

7.3.3.1 Covariance Spectrum

In Fig.7.5 we show the $4 - 10 \, \text{keV}$ covariance spectra from the pivoting model (for comparison we also show the covariance spectra from the non-pivoting model as

Figure 7.5: $4-10\,\text{keV}$ covariance spectra for low frequencies ($2.3 \times 10^{-5} - 7.3 \times 10^{-5}\,\text{Hz}$, red) and high frequencies ($2.3 \times 10^{-4} - 7.3 \times 10^{-4}\,\text{Hz}$, blue) for the propagating model including spectral pivoting of the power law (solid red and blue) and the propagating model without spectral pivoting (dashed red and blue). Black solid line shows total spectrum, dotted lines show model components.

dashed lines). This shows the spectrum of the variability that is correlated with variations in the $4 - 10\,\text{keV}$ band. Red shows the spectrum of the variability that is correlated at low frequencies ($2.3 \times 10^{-5} - 7.3 \times 10^{-5}\,\text{Hz}$) and blue shows the correlated variability at high frequencies ($2.3 \times 10^{-4} - 7.3 \times 10^{-4}\,\text{Hz}$). The covariance is calculated as in J13 (see also Wilkinson & Uttley, 2009). Briefly, a light curve is generated for each energy bin and also for the $4 - 10\,\text{keV}$ reference band (subtracting the light curve of the energy bin if that bin lies within the reference band). The light curves are then Fourier transformed and the power set to zero for all frequencies other than the range of interest. An inverse Fourier transform is then applied to transform the filtered periodogram back into a light curve, now containing variability only in a narrow frequency range. The covariance between the filtered energy bin and reference band light curves is then calculated as in Wilkinson & Uttley (2009).

J13 showed that at low frequencies the spectrum of correlated variability in PG1244+026 has the same shape as the total spectrum, i.e. there is correlated low frequency variability in all components. Our model reproduces this (Fig.7.5, solid red line), by propagation of slower fluctuations from the outer disc down to the inner components.

In contrast the observed high frequency variability correlated with the $4-10\,\text{keV}$ band drops off below $1\,\text{keV}$, following the shape of the time averaged coronal power law. The high frequency covariance spectrum calculated from our model similarly drops off below $1\,\text{keV}$. However the spectral slope of the correlated variability is different depending on whether the power law is allowed to pivot (solid blue line) or not (dashed blue line), with the slope being slightly harder than the time averaged power law for the model including spectral pivoting. The observed high frequency covariance does seem to suggest a slightly harder slope at the highest energies (J13), which would favour a pivoting model. However the low frequency lag-energy spectrum clearly rules out this scenario. This shows that it is not sufficient for a spectral model to match one aspect of the timing observations. A successful spectral model must correctly reproduce them all and matching one does not necessarily mean it will match the others.

In all subsequent simulations we choose to use the non-pivoting model since this is what the data require.

7.3.4 Reflection

In order to fit the time averaged spectrum we require some of the coronal power law to be reflected off the disc (Fig.7.1a, magenta component). The reflected component should therefore respond to variations in the coronal emission, with a light travel time delay. However there will be a range of time delays, as reflection will occur first from inner parts of the disc and take longer to travel to larger radii. If we are viewing the disc at some angle, not face on, the near side of the disc will appear to respond before the far side. The fastest fluctuations in the coronal emission will therefore be smoothed out. In order to account for this we calculate the transfer function of the disc.

We assume reflection occurs from the innermost parts of the disc, between $12-20\,R_g$, since below $12\,R_g$ the disc is replaced by the soft excess, and beyond $20\,R_g$ the solid angle subtended by the disc is rather small. This range of radii matches the range used in our spectral fit.

The time delay (τ) for light reflected off a point on the disc at radius r from a

Figure 7.6: Transfer functions for reflection off the disc ($R = 12 - 20\,R_g$, red) and soft excess ($R = 6 - 12\,R_g$, green).

central source is (Welsh & Horne, 1991):

$$\tau = \frac{r}{c}(1 - \sin i \cos \phi) \tag{7.3.6}$$

where i is the inclination of the axis of the disc to the line of sight and ϕ is the angle between the point on the disc and the projection of the line of sight onto the disc. Since we have not specified the geometry of the corona (beyond confining it to the central regions $< 10\,R_g$) we do not include any general relativistic corrections on photon light travel times. Given that we measure a low black hole spin from our spectral fit this is a reasonable approximation.

We assume the disc of PG1244+026 is inclined at 30° with respect to our line of sight, consistent with it being classed as a 'simple' NLS1, assuming that some of the complexity of the 'complex' NLS1s (Gallo, 2006a) is from absorption/scattering in a disc wind. We calculate the transfer function of the disc by summing the contribution to the time delay from each azimuth of the disc over all relevant radii, and show the result in Fig.7.6 (red line).

The fluctuations in the model's reflected component will therefore be the fluctuations of the coronal emission, convolved with this transfer function, ie:

Figure 7.7: As in fig 3, but now for model with propagation of fluctuations from disc to soft excess ($t_{lag,s} = 1000\,\text{s}$) to coronal power law ($t_{lag,p} = 600\,\text{s}$) including reflection off the disc ($R = 12 - 20\,R_g$, magenta).

$$\dot{M}_{refl}(t) = \int_0^{\tau_{max}} T(\tau)\dot{M}_p(t - \tau)d\tau \qquad (7.3.7)$$

Fig.7.7a shows the power spectrum of the reflected component (magenta), along with the other spectral components. There is a clear drop off in power in the reflected component above $\sim 3 \times 10^{-4}\,\text{Hz}$, corresponding to the highest frequency the transfer function can transmit.

Comparing the model lag-energy spectra in Fig.7.3d with Fig.7.7d, shows that including reflection has caused the highest energy bins ($> 5\,\text{keV}$) to lag slightly behind the $1.2 - 4\,\text{keV}$ reference band at low frequencies (red points). This is the region where reflection makes the largest contribution to the total spectrum (see Fig.7.1a), due to the presence of the iron line at $\sim 6.7\,\text{keV}$.

We find the reflection component alone is not enough to generate the observed high frequency soft lags (Fig.7.7c) as its contribution to the soft band is too small to overcome the soft lead from propagation.

7.3.5 Reprocessing on the Disc

Only a fraction of the coronal flux incident on the disc is reflected. The fraction that is not reflected will thermalise in the disc and be reprocessed. This happens in the same physical location as reflection, so will have the same transfer function, but (unlike reflection) the reprocessed emission is concentrated in the soft X-ray band. We first assume it is reprocessed on the BB disc, so that a fraction of the total flux in the BB disc spectrum will come from reprocessed flux. We assume that half of the coronal emission goes up — the half we observe as power law emission — and the other half goes down, back towards the disc. The total reprocessed flux is then the power law flux minus the flux of the reflected component: $L_{rep} = (\Omega/2\pi)L_p - L_{ref}$, where $\Omega/2\pi = 0.65$ is the solid angle of the reflector as measured from the spectral fit (Table 7.1). The fraction of total disc flux that is due to reprocessed emission is then $f_{rep} = L_{rep}/L_d \simeq 0.3$ for our chosen spectral decomposition. The thermalisation time is of order the Compton time for an AGN disc ($\sim \sigma_T nc \sim 0.5\,\mathrm{s}$ for $n \sim 10^{14}\,\mathrm{cm}^{-3}$, Stepney, 1983), hence we assume it is negligible compared to the timescale of coronal fluctuations.

The fluctuations in the spectrum from the BB disc therefore consist of a sum of the intrinsic disc fluctuations, which have fractional variability F_d around a BB of luminosity $f_{int}L_d = (1 - f_{rep})L_d \simeq 0.7L_d$, and the reprocessed fluctuations, which follow the coronal fluctuations (smoothed out by a transfer function) around a BB of luminosity $f_{rep}L_d$, ie:

$$\dot{M}_d(t) = f_{int}\dot{M}_{d,int}(t) + f_{rep}\int_0^{\tau_{max}} T(\tau)\dot{M}_p(t-\tau)d\tau$$
$$= f_{int}\dot{M}_{d,int}(t) + f_{rep}\dot{M}_{rep}(t) \qquad (7.3.8)$$

Fig.7.8 shows the model power spectra, coherence function, lag-frequency and
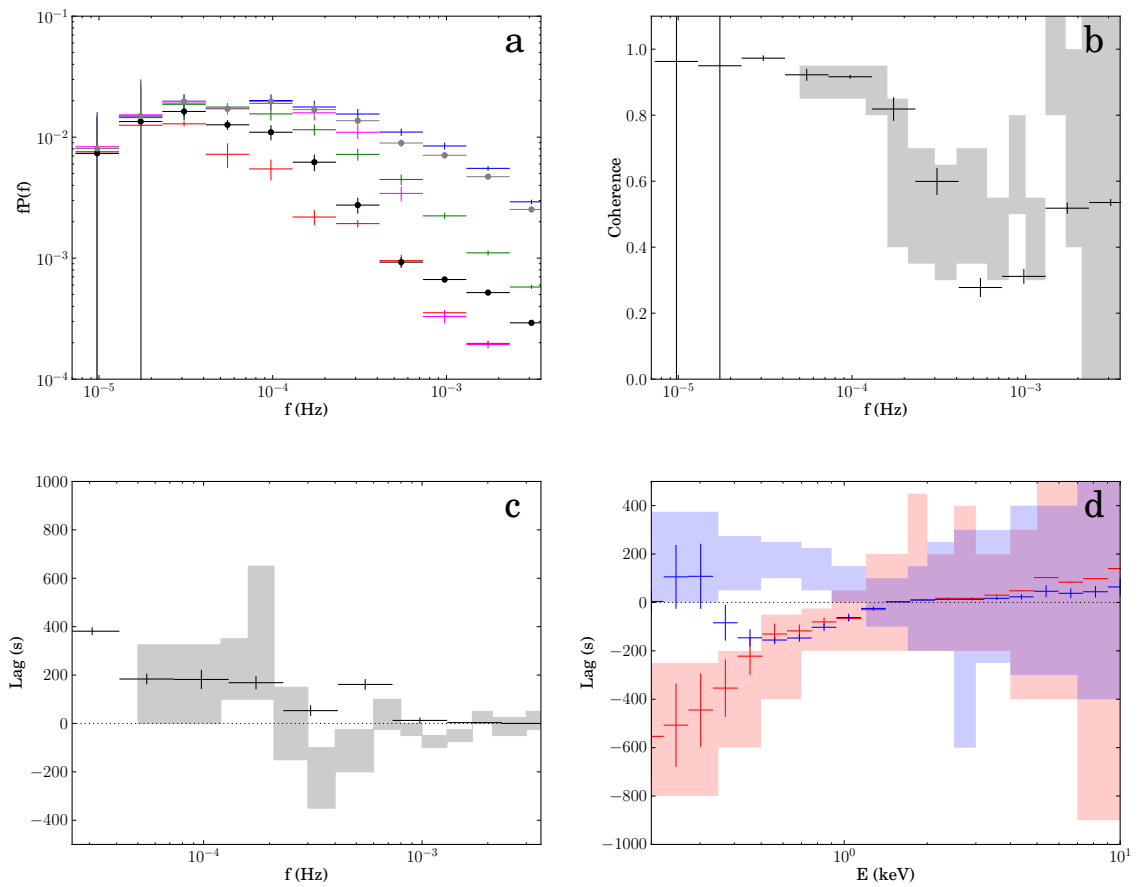
Figure 7.8: As in fig 3, but now for model with propagation of fluctuations from disc to soft excess ($t_{lag,s} = 1000\,\text{s}$) to coronal power law ($t_{lag,p} = 600\,\text{s}$) including reflection and reprocessing on the disc ($R = 12 - 20\,R_g$).

lag-energy spectra, now including reflection and reprocessing on the disc. Comparison of the model lag-frequency spectrum (Fig.7.8c) with Fig.7.7c shows that the low frequency hard lags have been reduced. This is due to dilution of the intrinsic disc fluctuations, which produce the hard propagation lags, by the reprocessed disc component, which produces soft lags. The net lag at low frequencies is simply the intrinsic hard lag minus the reprocessed soft lag, taking into account the relative proportions of each. So for there to be a net hard lag at low frequencies requires that the propagation lags dominate over the reprocessing lags. This is easily achieved if the intrinsic flux dominates over the reprocessed flux (as in this case). However it can still be achieved in situations where the reprocessed flux dominates, if the propagation lag is sufficiently longer than the shorter light travel time reprocessing lag to compensate for the smaller intrinsic fraction. But there must *be* a component capable of producing the propagation lag. Thus the existence of hard lags at low frequencies rules out models where the soft X-ray excess is produced by reflection and a soft jet power law, as we will show explicitly in Section 7.4.1.

Fig.7.8d shows the model lag-energy spectrum. The high frequency lags (blue) are almost identical to Fig.7.7d, except for the lowest three energy bins ($< 0.4 \, \text{keV}$). These show a small lag behind the $1.2 - 4 \, \text{keV}$ hard reference band. Comparison with the spectrum in Fig.7.1a shows that these are the only energies at which the disc emission dominates over the soft excess. Between $0.4 - 1 \, \text{keV}$ the lag-energy spectrum shows the energy bins leading the hard reference band, due to propagation lags generated by the soft excess. In contrast, the *observed* high frequency lag-energy spectrum shows all energy bins from $0.2 - 1 \, \text{keV}$ lagging the hard reference band (Fig.7.8d, shaded blue regions).

ADV14 suggest this could be achieved by allowing the reprocessed fluctuations to propagate from the disc down into the soft excess. The illuminating coronal flux heats up the disc, which then re-emits the radiation as reprocessed flux. The process of disc heating will alter the viscous frequency, allowing matter to propagate faster. In this way, fluctuations in the illuminating continuum can become accretion rate fluctuations in the disc, which can propagate inwards just like the intrinsic fluctuations.

Figure 7.9: Model lag-frequency spectrum, including reflection and reprocessing off the disc and propagation of reprocessed fluctuations. Shaded grey regions show range of error on the lags measured by ADV14 for PG1244+026.

To test the effect of this we now allow both the intrinsic disc fluctuations and fluctuations generated by reprocessing to propagate into the soft excess (and on to the corona). Since the coronal emission is now affected by the delayed propagated fluctuations, which are in turn affected by the time lagged coronal emission, we must iterate to find a self consistent solution. We first calculate the intrinsic disc, soft excess and power law fluctuations at each time. We then allow these fluctuations to propagate inwards and calculate the fluctuations of the reflected/reprocessed component as a function of time ($\dot{M}_{rep}(t)$). The BB component should consist of a sum of these two components: the intrinsic fluctuations (weighted by f_{int}) and the reprocessed fluctuations (weighted by f_{rep}), i.e. $\dot{M}_d(t) = f_{int}\dot{M}_{d,int}(t) + f_{rep}\dot{M}_{rep}(t)$. We repeat the calculation with this new $\dot{M}_d$, now including the reprocessed fluctuations, and again calculate $\dot{M}_{rep}$. This is fed into the next iteration and the process repeated until successive iterations result in a total fractional change in the power law fluctuations of less than 10^{-3}.

Fig.7.9 shows the resulting model lag-frequency spectrum, which is nearly identical to Fig.7.8c and still lacks a high frequency soft lag. We find any propagated reprocessed fluctuations are not strong enough to overcome the intrinsic soft excess fluctuations, which cannot themselves be reduced without losing the hard propaga-
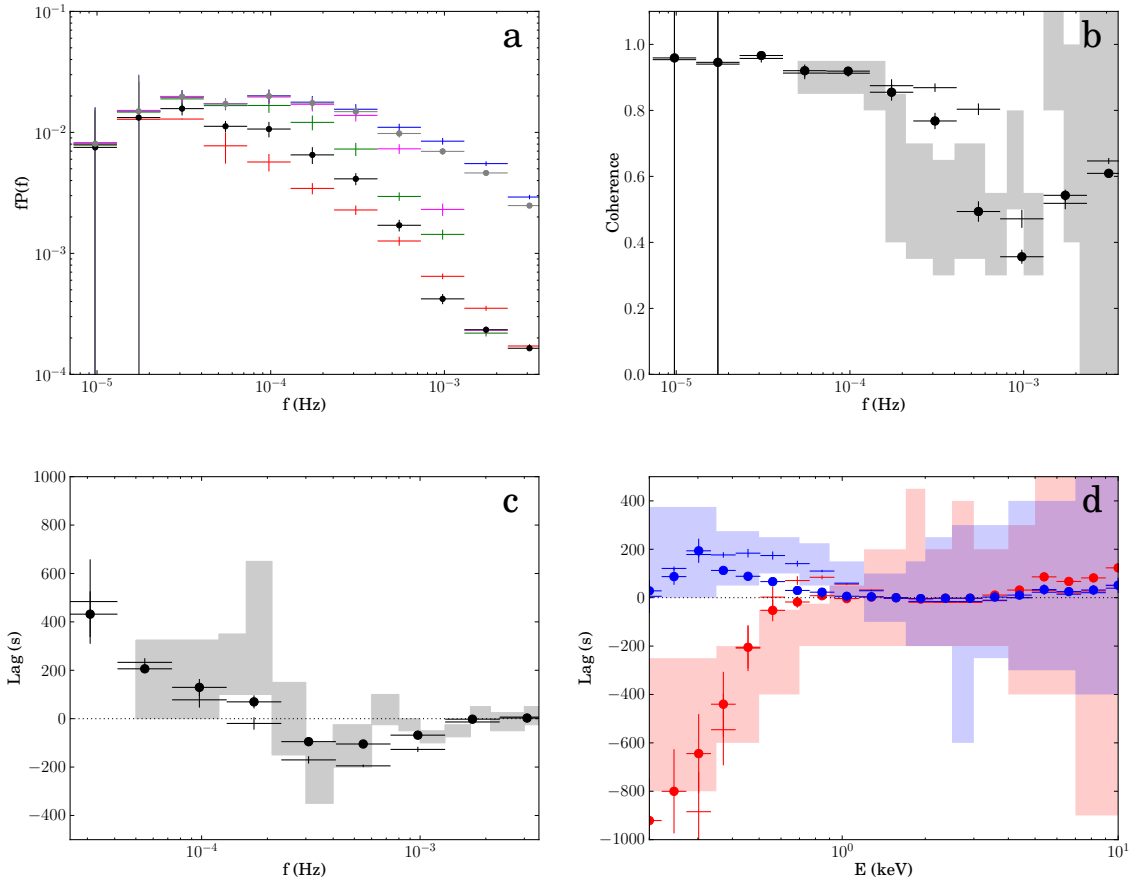
Figure 7.10: As in Fig.7.3, but now for model with propagation of fluctuations from disc to soft excess ($t_{lag,s} = 1000\,\mathrm{s}$) to coronal power law ($t_{lag,p} = 600\,\mathrm{s}$) including reflection and reprocessing on the soft excess ($R = 6 - 12\,R_g$) (crosses). Solid circles show model with half the reprocessing occurring on the disc, half on the soft excess.

tion lags. In order to replicate the observed lag-energy spectrum, the reprocessing must occur not on the disc, but the soft excess.

7.3.6   Reprocessing on the Soft Excess

We rerun our model, this time with the reprocessing occurring in the soft excess. We assume this has a reflection spectrum with the same ionisation as used for the disc, i.e. matching that seen fit to the data in Section 7.2, since the structure of this region is not known. We change our transfer function to match the size of the soft excess region, $\sim 6 - 12\,R_g$ (Fig.7.6, green line). This increases the maximum frequency of the reflected/reprocessed power in the model (Fig.7.10a, magenta points). Consequently the model coherence remains high up to higher

frequencies (compare crosses Fig.7.10b with Fig.7.8b).

Fig.7.10c (crosses) shows the model lag-frequency spectrum, which shows hard propagation lags at low frequencies ($\lesssim 10^{-4}$ Hz) and a soft reprocessing lag between $10^{-4} - 10^{-3}$ Hz. This is in much better agreement with the observed lags of PG1244+026 (ADV14).

Fig.7.10d shows the model lag-energy spectrum. The high frequency lag (blue crosses) shows the soft energy bins lagging the hard reference band for nearly all energies below ~ 1 keV, in much better agreement with the observations. Only the very lowest energy bins, where the disc dominates, now lack a strong soft lag. It is likely there is reprocessing on both the disc and soft excess, though clearly reprocessing in the soft excess region dominates in the band of our observations.

The soft excess is less luminous than the disc, as a consequence the fraction of reprocessed to intrinsic emission is much larger with $f_{int} = 0.3$, $f_{rep} = 0.7$, while when reprocessing is confined to the disc the fractions are roughly reversed. Allowing for some reprocessing on the disc as well as the soft excess would reduce the reprocessed fraction. Nevertheless, even in this limiting case, the model still produces soft leads in addition to reverberation lags. This is because the propagation time lags between soft excess and corona are longer than the light travel time delay for reverberation, so can be more heavily diluted without losing the net lead at low frequencies.

However, a very high reprocessing fraction on the soft excess will affect the model covariance spectrum. In particular, the spectrum of the correlated variability at high frequencies will be extended to lower energies. We calculate the covariance spectra for the soft excess reprocessing model to check they do not disagree with the data and show them in Fig.7.11 (solid red and blue lines). The low frequency correlated variability still matches the shape of the total spectrum, in agreement with the observations. The high frequency correlated variability now rolls over at ~ 0.6 keV rather than 1 keV, due to the amount of reprocessing on the soft excess. The dashed red and blue lines in Fig.7.11 show a model where half of the flux available for reprocessing is reprocessed on the soft excess and the other half is reprocessed on the disc. This increases the energy of the roll-over. Hence both the high frequency

Figure 7.11: $4-10$ keV covariance spectra for low frequencies $(2.3 \times 10^{-5} - 7.3 \times 10^{-5}$ Hz, red) and high frequencies $(2.3 \times 10^{-4} - 7.3 \times 10^{-4}$ Hz, blue) for the model including reprocessing on the soft excess (solid red and blue) and model including reprocessing on both the soft excess and the disc (dashed red and blue). Black solid line shows total spectrum, dotted lines show model components.

lag-energy spectrum and high frequency covariance spectrum suggest there is some reprocessing on both the disc and the soft excess. In Fig.7.10b, c and d we also show the coherence, lag-frequency and lag-energy spectra for the combined disc and soft excess reprocessing model (circles). This shows that a model with a combination of propagating fluctuations, from the disc through a separate soft excess component to the coronal power law, together with reprocessing on both the disc and the soft excess component, can capture all the main features of the data.

We note that including reprocessing on the disc as well as the soft excess does not exactly match the shape of the high frequency lag-energy spectrum. This is because the disc reprocessing is likely to be concentrated on the inner edge of the disc, which our simple model doesn't account for. In other words, the data suggest the reprocessed emission should have a spectrum similar to that of the soft excess component, but extending to slightly lower energies to include the innermost radii of the disc.

7.4   Comparison with a Reflection Model for the Soft X-ray Excess

Reflection dominated models have a strong soft X-ray excess from a combination of line and continuum. However, the observed soft X-ray excess is smooth, so this requires strong relativistic smearing. In complex NLS1s such as 1H0707–495, this smearing is so extreme as to require strongly centrally concentrated emissivity ($\propto r^{-7}$) onto a high spin black hole (Fabian et al., 2009). This makes all size-scales, and hence reflection/reprocessing lag times shorter.

However, while this can reproduce the soft lags, a single power law which varies only in normalisation and its constant ionisation reflection cannot match the soft lead at low frequencies. Chainakun & Young (2012) showed that the mismatch was made worse by including the radial and time variability of the ionisation of the reflector which should arise from the extremely centrally concentrated and variable illumination. This means that more complex continuum models are required to match the spectral-timing data, in particular to match the soft lead.

7.4.1   Hard Coronal Power Law and its Reflection with Additional Soft Power Law from the Jet

K14 decompose the spectrum of PG1244+026 as a soft power law, together with a harder power law and its reflection (Fig.7.12a, see Table 7.2 for spectral parameters). The soft power law then gives an additional component which could give a source for the soft lead. However, they interpret this soft power law as coming from the jet. Since fluctuations propagate through the accretion flow and then up the jet, the soft emission from the jet should always lag the harder emission from the accretion flow. In this model there is no component of the soft emission which leads the power law, as reflection also always lags, so this cannot give an origin for the soft lead seen in the data at low frequencies. We show this explicitly below.

We assume the intrinsic power law variations have a broad power spectrum consisting of four Lorentzians centred at $f_{visc,p} = 3 \times 10^{-5}$, 1×10^{-4}, 3×10^{-4} and

Component	Parameter	§7.4.1/§7.4.3	§7.4.2	§7.4.4
Galactic absorption	N_h (10^{22} cm^{-2})	0.019	0.019	0.019
Intrinsic absorption	N_h (10^{22} cm^{-2})	0.019	0.029	0.047
Hard power law	Γ	2.26	2.3	2.3
	norm	6.9×10^{-4}	7.3×10^{-4}	4.9×10^{-4}
Soft power law	Γ	3.5	3.5	3.0
	norm	1.6×10^{-3}	1.5×10^{-3}	9.3×10^{-4}
KDBLUR	index	3.4	3.9	4.5
	r_{in} (R_g)	3.2	3.1	3.0
RFXCONV	relative refl norm	-2.6	-3.2	-2.8
	$\log(x_i)$	2.7	2.7	3.0
BBODY	kT (keV)	-	0.027	0.032
	norm	-	4.0×10^{-3}	3.6×10^{-3}

Table 7.2: Parameters for the spectral models shown in Fig.7.12–Fig.7.15, with reflection dominating the soft X-ray excess emission. a). Model used in Sections 7.4.1 and 7.4.3: WABS*ZWABS(POWERLAW + POWERLAW + KDBLUR*RFXCONV*POWERLAW). b). Model used in Section 7.4.2: WABS*ZWABS(BBODY + POWERLAW + POWERLAW + KD-BLUR*RFXCONV*POWERLAW). c). Model used in Section 7.4.4: WABS*ZWABS(BBODY + POWERLAW + POWERLAW + KDBLUR*RFXCONV(POWERLAW + POWERLAW)). For all models Z_{Fe} was fixed at 1.0.

1×10^{-3} Hz, representing the accumulation of fluctuations generated at different radii. This is necessary to replicate the broad range of power observed in the hard band power spectrum. Since there are no disc/soft excess components, all these frequencies must be generated intrinsically in the power law. The jet power law peaks in the soft band. Since the observed soft band power spectrum has less high frequency power than the hard band, we assign the intrinsic jet fluctuations a power spectral Lorentzian at 1×10^{-4} Hz. Since the data require a drop in coherence above 10^4 Hz, we allow fluctuations to propagate from the corona up into the jet with a lag time of $t_{lag,jet} = 10^4$ s. This means any fluctuations faster than $\frac{1}{t_{lag,jet}}$ should be smoothed out, causing a drop in coherence at 10^{-4} Hz. We calculate the disc transfer function between $r_{in} = 1\,R_g$ and $r_{out} = 12\,R_g$, to better match the smaller radii used to calculate the reflection spectrum in K14. We assume the ionisation state of the reflector remains constant with both radius and time (Chainakun & Young, 2012).

Fig.7.12 shows the resulting model power spectra, coherence, covariance, lag-frequency and lag-energy spectra. The model power spectra for the hard and soft bands are nearly identical (Fig.7.12c). This is in disagreement with the observations,

Figure 7.12: Results for model with a hard coronal power law and its reflection with additional soft power law from the jet. $t_{lag,jet} = 10^4$ s for propagation of fluctuations from the corona to the jet. Reflection from $R = 1 - 12 R_g$. a). Spectral decomposition: soft jet power law (green), hard coronal power law (blue) and reflection of the hard power law (magenta), with total shown in black. Data points show time averaged spectrum (OBS ID: 0675320101). b). $4 - 10$ keV covariance spectra for low frequencies ($2.3 \times 10^{-5} - 7.3 \times 10^{-5}$ Hz, red) and high frequencies ($2.3 \times 10^{-4} - 7.3 \times 10^{-4}$ Hz, blue). Black solid line shows total spectrum, dotted lines show model components. c). Power spectra: jet (green), coronal power law (blue), reflection (magenta), soft band ($0.3 - 1.0$ keV, black), hard band ($2 - 10$ keV, grey). d). Coherence between hard ($1.2 - 4$ keV) and soft ($0.3 - 0.7$ keV) bands. e). Lag-frequency spectrum between hard ($1.2 - 4$ keV)and soft ($0.3 - 0.7$ keV) bands. f). Lag-energy spectrum calculated using $1.2 - 4$ keV reference band: low frequency lag ($2.3 \times 10^{-5} - 7.3 \times 10^{-5}$ Hz), red points; high frequency lag ($2.3 \times 10^{-4} - 7.3 \times 10^{-4}$ Hz), blue points. Shaded regions show approximate range of error on values measured by ADV14 for PG1244+026.

which show more high frequency power in the hard band than the soft band (J13). This is because all three components — soft power law, hard power law and reflection — contribute strongly to both the hard and soft energy bands. To achieve a drop in high frequency power in the soft band requires a long lag time between the hard power law and the soft jet, to smooth out the high frequency variability. But because the hard band consists of the same three components this also produces the same effect in the hard band. Hence the incoherence between the hard and soft power laws, required to limit the high frequency power in the soft band, also restricts the amount of high frequency power in the hard band. Reducing the lag between power law and jet to increase high frequency power in the hard band only serves to give a worse match to the soft band power spectrum. Producing different hard and soft band power spectra requires at least some of the components to be confined to one band. This spectral decomposition does not meet this criterion.

The similarity of the model hard and soft band power spectra means the spectra of the correlated variability at low frequencies and high frequencies are both identical to the total spectrum (Fig.7.12b). This is in disagreement with the data, which require the spectrum of the variability correlated at high frequencies to drop off below $1\,\mathrm{keV}$.

In Fig.7.12d we show the model coherence. For this model the coherence remains high up to very high frequencies ($\sim 10^{-3}\,\mathrm{Hz}$). This is despite our attempts to limit the coherence with a long propagation time to the jet, and again is a result of all three components contributing strongly to both hard and soft bands. In particular, the small radii used for reflection allow the reflection component to respond at high frequencies, which are coherent between hard and soft bands because reflection dominates the soft X-ray excess and also the iron line region at $\sim 6.7\,\mathrm{keV}$. Consequently it is the roll-over in the reflected component power spectrum at $\sim 4 \times 10^{-4}\,\mathrm{Hz}$ that determines the drop in model coherence, not our jet lag time. This is in disagreement with the observations, which show a drop in coherence at $\sim 10^{-4}\,\mathrm{Hz}$ (ADV14).

The lag-frequency spectrum shows this model cannot reproduce the observed soft leads at low frequencies (Fig.7.12e). This is shown more clearly in the lag-energy spectrum (Fig.7.12f). At high frequencies the model lag-energy spectrum

picks out the shape of the reflection spectrum (blue points). This roughly matches the observed high frequency hard lags, although the lag times are a little short due to the very small radii required to produce the relativistically smeared reflection. In contrast, the model's low frequency lag-energy spectrum does not match the data at all. There is no slowly varying soft component in this model that leads the hard power law. The only slowly varying component is the soft jet power law which can only lag, since fluctuations go through the accretion flow before travelling up the jet. Hence at low frequencies there is a soft lag which is inconsistent with the soft lead seen in the data. Above $\sim 2\,\mathrm{keV}$, the hard energy bins increasingly lead the hard reference band as energy increases. This is a result of the contribution from the soft jet power law to the total flux decreasing and causing less dilution of the hard power law.

7.4.2 Hard Coronal Power Law and its Reflection with Additional Soft Power Law from the Jet and a Disc Component

We rerun the model of the previous section, this time assigning the lowest frequency fluctuations ($f_{visc} = 3 \times 10^{-5}\,\mathrm{Hz}$) to a disc component (Fig.7.13a, see Table 7.2 for spectral parameters), and allowing them to propagate down to the power law ($t_{lag,p} = 600\,\mathrm{s}$) and on to the jet ($t_{lag,jet} = 10^4\,\mathrm{s}$).

Fig.7.13 shows the resulting model power spectra, coherence, covariance, lag-frequency and lag-energy spectra. In this model, the soft lags of the lowest energy bins ($\lesssim 0.4\,\mathrm{keV}$) have been diluted but only the lowest energy bin has been diluted enough to show a soft lead at low frequencies (Fig.7.13f). Even if the propagation lag time to the jet component is reduced (to minimise dilution of the disc lead), the disc makes little contribution to the spectrum above $0.4\,\mathrm{keV}$, so cannot replicate the observed soft leads up to $\sim 1\,\mathrm{keV}$. The dominance of the jet spectrum together with reflection, both of which lag the power law, prevent there being any soft leads at these energies. Hence extending the model to include a separate very soft component from the disc does not help, because the disc makes too small a contribution to the

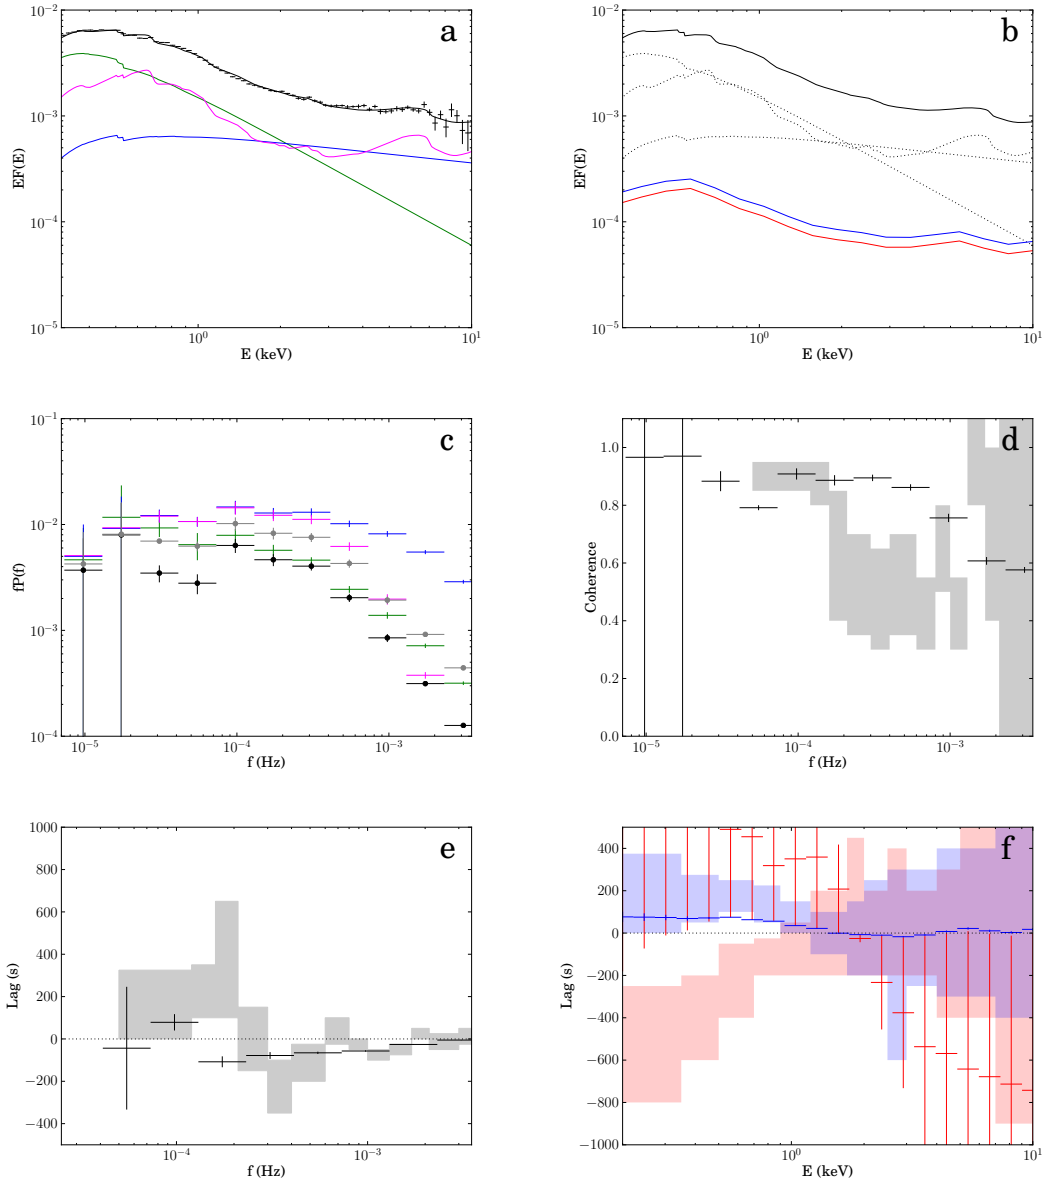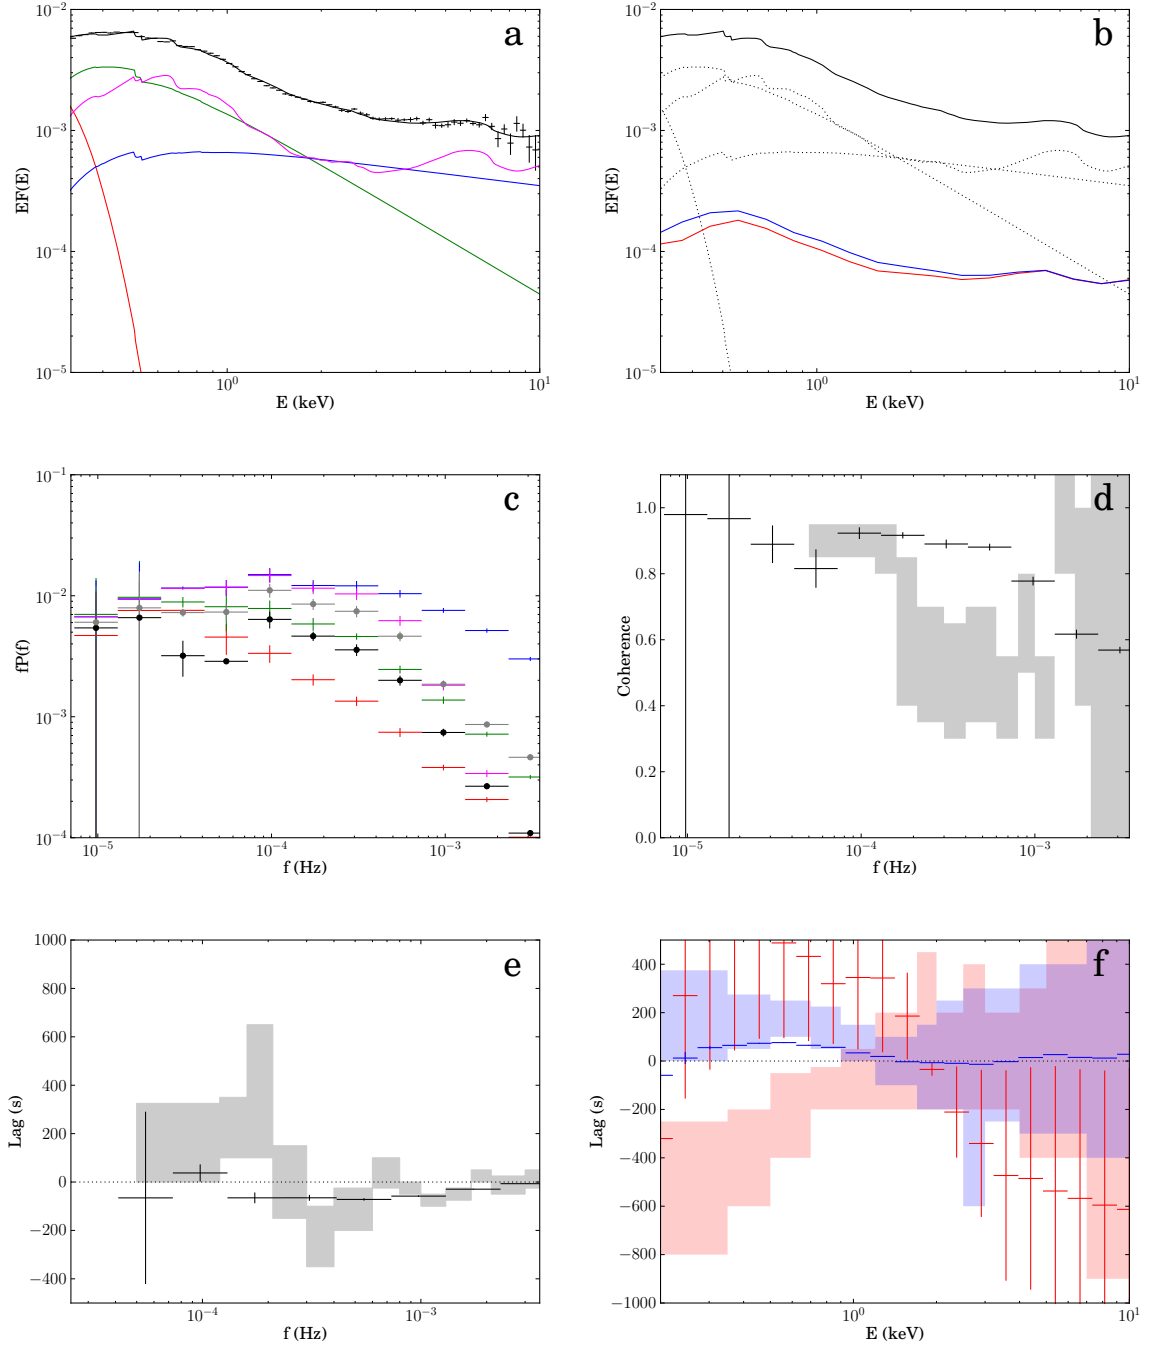**Figure 7.13:** As in Fig.7.12, but for model with a hard coronal power law and its reflection with additional soft power law from the jet and a disc component (red). $t_{lag,p} = 600\,\mathrm{s}$ for propagation of fluctuations from the disc to the corona and $t_{lag,jet} = 10^4\,\mathrm{s}$ for propagation from the corona to the jet. Reflection from $R = 1 - 12\,R_g$.

soft band to give a large enough effect, and its soft lead is swamped even at low frequencies by the soft lag from the reflection component.

7.4.3 Hard Coronal Power Law and its Reflection with Additional Soft Power Law from the Accretion Flow

Instead, the model lags match more easily to the data if the soft power law is produced between the disc and the hard power law region, i.e. it comes from the accretion flow rather than the jet (Fig.7.14a).

We assume the soft power law represents the outer parts of the accretion flow, with a power spectrum consisting of two low frequency Lorentzians ($f_{visc,s} = 3 \times 10^{-5}$ and 1×10^{-4} Hz). We allow these fluctuations to propagate down into the inner harder power law, which generates intrinsic fluctuations at $f_{visc,p} = 3 \times 10^{-4}$ and 1×10^{-3} Hz. The inner harder power law produces the reflection.

Fig.7.14 shows the resulting model power spectra, coherence, covariance, lag-frequency and lag-energy spectra. The model lag-frequency spectrum (Fig.7.14e) now shows a soft lead below 2×10^{-4} Hz and a soft lag above $\sim 3 \times 10^{-4}$ Hz. However, producing this soft lead requires a lag time between the soft and hard power laws double that of the previous separate soft excess model, due to the strong reflection component causing much more dilution. The model lag-energy spectrum taken at the frequency of the soft lead shows the energy bins lead the hard band at energies below $\lesssim 0.6$ keV, where the contribution from reflection decreases (Fig.7.14f, red points). The highest energy bins (> 2 keV) lag the hard reference band at low frequencies, with the lag increasing with increasing energy. This is because at 2 keV the contributions from the hard and soft power laws are equal and above this energy the fraction of total flux contributed by the leading soft power law decreases while the fraction contributed by the lagging hard power law increases. The data do not show a systematically increasing low frequency lag at high energies, despite the large errors. This suggests whatever soft component is leading the hard power law is confined to the soft band and does not contribute significant flux at energies above 2 keV. In other words, the low temperature optically thick soft excess component of Section 7.3, which is confined to the soft band, provides a better match to the data
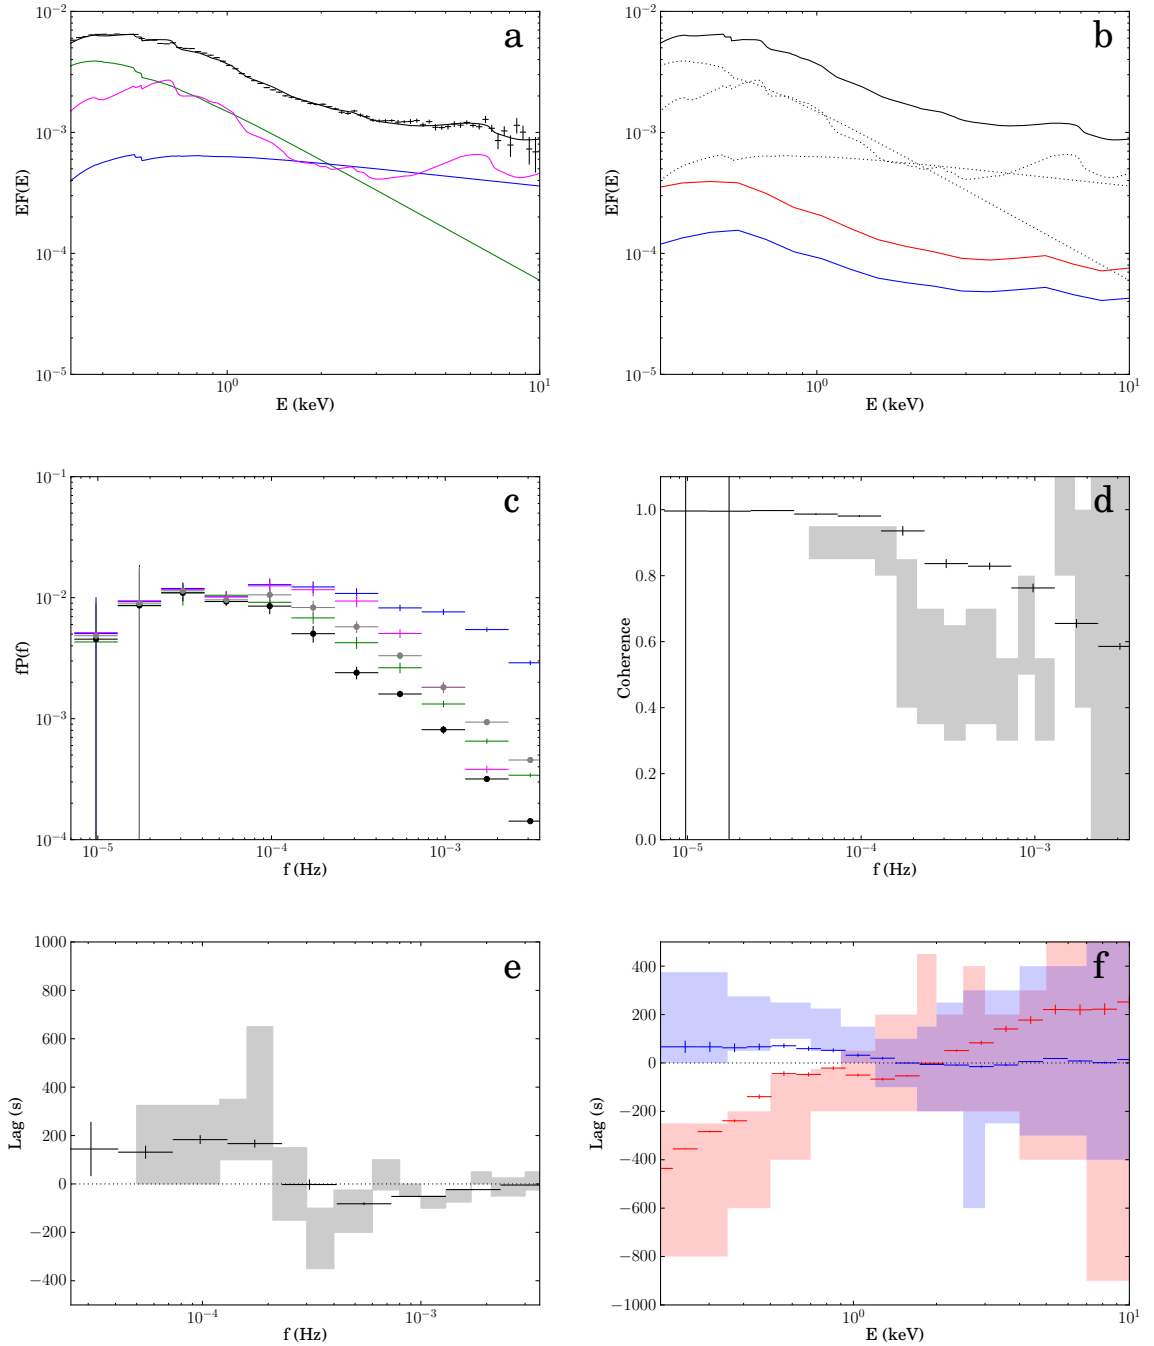
Figure 7.14: As in Fig.7.12, but for model with a hard coronal power law and its reflection with additional soft power law from the accretion flow. $t_{lag,p} = 1200\,\mathrm{s}$ for propagation of fluctuations from the soft power law to the hard. Reflection from $R = 1 - 12\,R_g$. Low frequency lag-energy spectrum now calculated between $7.3 \times 10^{-5} - 2.3 \times 10^{-4}\,\mathrm{Hz}$, high frequency between $4.1 \times 10^{-4} - 1.3 \times 10^{-3}\,\mathrm{Hz}$.

than an extended soft power law.

At high frequencies, the model lag-energy spectrum keeps the shape of the reflection spectrum (Fig.7.14f, blue points). However the small radii required for relativistically smeared reflection mean shorter lag times, which do not match the observed high frequency soft lags as well as the longer lag times produced by the previous soft excess model, where reflection/reprocessing occurred at larger radii (Section 7.3). The roll-over at $0.6\,\mathrm{keV}$ of the reflection spectrum also occurs at too low an energy to reproduce the observed decrease in correlated high frequency variability at $1\,\mathrm{keV}$ in the covariance spectrum (Fig.7.14b).

Whilst changing the order in which fluctuations propagate through the components has given a better match to the lags compared to the original model in Section 7.4.1, the spectral decomposition itself has not changed. Therefore this model has the same problems replicating the power spectra and coherence as the first model did, i.e. the model coherence remains higher than the observations at all frequencies (Fig.7.14d) and the model hard and soft band power spectra are still too similar, simply because both bands contain strong contributions from all three spectral components.

7.4.4   Soft and Hard Coronal Power Laws plus Reflection and a Disc Component

Finally we test the more complex model used to describe the NLS1 1H0707–495, consisting of two power laws from the accretion flow, both with reflection spectra, and a contribution from a BB disc.

As before, we assign the lowest frequency intrinsic fluctuations to the disc ($f_{visc,d} = 3 \times 10^{-5}\,\mathrm{Hz}$) and increase the frequency of the fluctuations from the soft power law ($f_{visc,s} = 1 \times 10^{-4}\,\mathrm{Hz}$) to the hard power law ($f_{visc,p} = 3 \times 10^{-4}$ and $1 \times 10^{-3}\,\mathrm{Hz}$). Since the soft power law represents the outer parts of the accretion flow, we calculate the disc transfer function for its reflection between $12 - 20\,R_g$. We assume the inner hard power law is reflected from the inner regions of the disc ($1 - 12\,R_g$). We set $t_{lag,s} = 1000\,\mathrm{s}$ for propagation of fluctuations from the disc to the soft power law and $t_{lag,p} = 600\,\mathrm{s}$ for propagation of fluctuations from the soft power law and to the

Figure 7.15: As in Fig.7.12, but for model with soft and hard coronal power laws plus their reflection (orange ($R = 12 - 20\,R_g$) and magenta ($R = 1 - 12\,R_g$) respectively) and a disc component (red). $t_{lag,s} = 1000\,\mathrm{s}$ for propagation of fluctuations from the disc to the soft power law and $t_{lag,p} = 600\,\mathrm{s}$ for propagation from the soft power law and to the hard.

hard.

Fig.7.15 shows the resulting model power spectra, coherence, covariance, lag-frequency and lag-energy spectra. Including two reflection spectra swamps the propagation leads at low frequencies (Fig.7.15e). They are only visible in the model lag-energy spectrum in the lowest energy bins dominated by the disc (Fig.7.15f, red points). The model's high frequency lag-energy spectrum (Fig.7.15f, blue points) lacks a strong lag at the lowest energies due to the roll-over of the reflection spectra and the lack of any reprocessing on the disc. The soft lags are also still too short to match the data, despite including reflection from slightly larger size-scales, because there is strong dilution by the other four components.

Moreover, this model suffers from the same problem as the three previous reflection models: the hard and soft bands both contain strong contributions from (nearly) all components, leading to high coherence at all frequencies and very similar power spectra for both bands (Fig.7.15b, c and d). In contrast the data require low frequencies to be generated in the soft band, high frequencies to be generated in the hard band, longer disc transfer functions to limit the amount of high frequency power transmitted back to the soft band and, importantly, for the main reverberation response to be concentrated in the soft band not spread over both bands.

7.5   Discussion and Conclusions

Time variability gives additional information which can break spectral degeneracies. Any successful model must be able to fit both spectral and timing properties of the data. Here we show quite generally that the switch in behaviour from soft leading hard at low frequencies to soft lagging hard at high frequencies favours a model where the majority of the soft X-ray excess is not produced by reflection, but consists of a combination of emission from the accretion flow and reprocessed emission. We show this explicitly using the simple NLS1 PG1244+026 as an example. Reflection dominated models including a soft power law from the jet (K14) do not produce soft leads at low frequencies, as fluctuations propagate to the jet only after they have

gone through the accretion flow.

If the additional soft power law is from the accretion flow rather than the jet, so the fluctuations propagate from the soft power law to the hard, these models can produce soft leads but they do not match the shape of the observed lag-energy spectrum, due to the soft power law extending into the hard band, and they cannot match the other observed timing features of the data. The fact that reflection contributes a large fraction of the flux in both the hard and soft energy bands means these models produce high coherence over a wide frequency range, and very similar hard and soft band power spectra. This is made worse by the small size-scales over which the reflection occurs (necessary to produce such strongly relativistically smeared emission), which allow the reflection component to respond at very high frequencies. All this is at odds with the observations, which show a drop in coherence between the hard and soft bands at high frequencies, and much less high frequency power in the soft band compared to the hard.

Since the argument is quite general, it is likely that the similar model (disc plus soft and hard power laws and their reflection) used for complex NLS1s, e.g. 1H0707–495, cannot explain their data either. All the papers using this model concentrate on how this fits the soft lags at high frequencies (Zoghbi et al., 2011; Kara et al., 2013; K14), but the soft leading at low frequencies also needs to be explained and any spectral decomposition *must* be able to replicate the observed power spectra and coherence.

Instead, the switch in lag behaviour can be explained if the soft X-ray excess is dominated by an intrinsically curved component, such as a low temperature, high optical depth Compton component. Fluctuations start in the disc, propagate down through this intrinsic soft excess component and then into the power law. Reflection alone does not produce enough soft lag in these models as the reflected emission makes too small a contribution to the soft X-ray bandpass. However, the non-reflected flux should be reprocessed, and this reprocessed emission reverberates in a similar way to reflection, as is seen in the BHBs (Uttley et al., 2011). Importantly this response is concentrated in the soft band and high frequencies are filtered out due to the larger radii involved. We explore two models of this reprocessed emission,

one where it thermalises to the disc temperature, and one where it thermalises to produce part of the soft excess. The low temperature of the disc in our spectral models means that it does not contribute enough flux to the soft X-ray band to turn the propagation soft lead into a lag, but thermalising to the soft excess can match the switch in behaviour of the lag as a function of frequency. In reality there is likely to be some reprocessing on both the disc and soft excess. We show that a model including both gives a good match to all the timing observations.

An important additional finding from our simulations is that the observed lag-energy spectra appear to show no evidence of spectral pivoting of the coronal power law; the power law spectral slope does not vary on short timescales, and instead only changes in normalisation. This lack of spectral pivoting could be explained if there are correlated changes in coronal power and seed photons, perhaps through evaporation of material from the soft excess into the corona. Alternatively this could suggest that the coronal emission mechanism in NLS1s is non-thermal, so that the seed photons only experience one scattering before escaping from the corona.

A model for the soft X-ray excess in which it consists of a combination of intrinsic emission from the accretion flow, together with reprocessed emission, can fit all the current spectral and timing properties of PG1244+026. Such models do not require inner disc radii smaller than $\sim 6\,R_g$, so put no constraints on the spin of the black hole. In contrast, reflection dominated models for the soft X-ray excess, which require the black hole to be highly spinning, can only replicate the high frequency soft lags and cannot match the observed power spectra, coherence or covariance. This is a direct consequence of the extremely small radii ($\sim 1\,R_g$) required to produce these spectra and the dominance of the smeared reflection component in both the hard and soft energy bands.

However, we note that neither model can explain the possible detection of very high frequency soft lags by ADV14. We show that the model proposed by ADV14, where the irradiated disc fluctuations propagate down to the soft excess, cannot quantitatively explain this feature, suggesting this feature requires some additional reverberation signal from smaller size-scales.

Chapter 8

Complex Narrow Line Seyfert Type 1s: High Spin or High Inclination?

8.1 Introduction

In the previous chapter we showed that the spectral and timing properties of the simple Narrow Line Seyfert 1 (NLS1) PG1244+026 were best explained by a model where the soft X-ray excess is a combination of optically thick Comptonisation and reprocessing. In contrast, high spin, reflection dominated models did not provide a good match to its behaviour. However some NLS1s have very different spectral and variability properties.

All NLS1s show rapid X-ray variability (Leighly, 1999), but some also show deep dips in the X-ray light curve. These dips coincide with the appearance of high energy complexity in the $2-10\,\mathrm{keV}$ spectra, either gradual curvature or strong features around the Fe Kα line energy. Gallo (2006b) termed these 'complex' NLS1s to distinguish them from the 'simple' NLS1s which do not show dips and have relatively power law like spectra from $2-10\,\mathrm{keV}$. Two different models have been proposed to explain the deep dips and associated spectral complexity: partial covering and relativistic reflection.

In the partial covering model, the dips are caused by low ionisation material

moving into the line of sight, increasing the absorption at the iron edge energy at
7.1 keV. This material can only partially cover the source as some fraction of the
flux at low energies is still seen (e.g. Inoue & Matsumoto, 2003; Turner et al.,
2007; Miyakawa et al., 2012). Conversely, in the relativistic reflection model the
dips are caused by an extremely compact X-ray source on the spin axis of the black
hole approaching the event horizon. The resulting strong light bending focusses the
intrinsic continuum away from the observer (producing the drop in flux) so it instead
strongly illuminates the very inner disc. For high spin black holes the resulting
spectrum can be dominated by highly smeared relativistic reflection, marked by a
strong but extremely broad and skewed Fe Kα line (e.g. Fabian et al., 2004; Miniutti
& Fabian, 2004; Fabian et al., 2009). In both models, the complex and simple NLS1s
are intrinsically similar, and can change from one to the other (as observed: Gallo,
2006b) depending on whether there is absorption along the line of sight, or in the
reflection model, whether the compact X-ray source is close to the horizon.

Both absorption and reflection models can fit the observed $0.3 - 10$ keV spectra,
as fitting complex models over a limited bandpass is highly degenerate. Variability
can be used as an extra diagnostic. Detection of a very short soft lag (~ 30 s) in the
complex NLS1 1H0707–495 is often taken as unequivocal support for the high spin
relativistic reflection picture, as this implies distances of the source from the disc
of $< 2 R_g$ for a $3 \times 10^6 M_\odot$ black hole. In contrast, the simple NLS1 PG1244+026
shows a much longer lag time of ~ 200 s (Alston et al., 2014), consistent with the
source being somewhat further from the horizon if source height is assumed to cause
the difference between simple and complex sources.

Here we investigate the effect of the partial covering model on the timing proper-
ties of the source. NLS1s all accrete close to the Eddington limit, so it seems unlikely
that they are actually described by a flat disc, as assumed in the relativistic rever-
beration models. Global MRI simulations at $\dot{M}_{in} = 20\dot{M}_{Edd}$ show a complex flow
structure, with a large scale-height radiation pressure driven wind from the inner
disc carrying away 30% of the input mass accretion rate (Jiang et al., 2014). This
wind is likely to be less strong in PG1244+026 since this source is only $\dot{M}_{in} \sim \dot{M}_{Edd}$.
UV line driving is unlikely to help since the low mass and high mass accretion rate

of NLS1s mean that their disc is too hot for its photosphere to have the required opacity to UV line transitions (Hagino et al., 2015). Hence strong mass loss is not expected in NLS1s, but some turbulent, clumpy, failed Eddington wind in the inner disc could easily form (Jiang et al., 2014). High inclination angles have a high probability of a clump intersecting the line of sight, while low inclination angles are mostly free of obscuration, giving a potential mapping from complex to simple NLS1 as a function of inclination. The spectral signature of this time dependent absorption should be complex, depending on the amount of source occulted and the ionisation state of the absorber.

Here we couple orbiting clump occultations with the full spectral timing model developed in the previous chapter to explain the properties of the simple NLS1 PG1244+026. We investigate whether introducing occultations can replicate the change in observed spectral and timing properties from a simple to a complex NLS1, without requiring a highly spinning central black hole.

8.2   Linear Occulation

We choose to model the underlying accretion flow using the separate soft excess model found in the previous chapter to well describe the spectral and timing properties of the simple NLS1 PG1244+026 (Gardner & Done, 2014a, hereafter GD14). To briefly recap: we assume the accretion flow consists of three components. The outermost radii form a standard accretion disc. This is truncated at some radius ($R_{cor} \sim 20\,R_g$). The remaining gravitational energy liberated between R_{cor} and the innermost stable circular orbit (R_{isco}) is used to power the remaining two components: the soft excess and the corona. Below R_{cor} material is unable to thermalise completely and form a cool accretion disc, perhaps due to a larger scale-height from the photosphere lifting to form the (failed) wind. Instead some of the electrons emit via optically thick Comptonisation of the cooler disc seed photons. This optically thick Compton emission adds to the spectrum at low energies, producing an excess of soft X-rays, hence we call the physical region producing this emission the 'soft excess'. An optically thin corona extends above the soft excess at the very smallest

Figure 8.1: a). Physical location of accretion flow model components, as viewed at high inclination ($i = 70°$): disc (red), soft excess (green), corona (blue), obscuring cloud (black). For simplicity we have approximated the geometry of the accretion flow as a flattened disc. b). Spectrum emitted by each component, with total spectrum shown in black, and hard and soft bands shaded in blue and red respectively.

radii. This corona contains hot electrons at lower density and higher energy than the soft excess and provides a source of optically thin Comptonisation, using seed photons from the cooler soft excess region (Jin et al., 2013).

Fig.8.1a shows the physical locations of each model component and Fig.8.1b shows their contributions to the total X-ray spectrum. For simplicity we approximate the geometry of the accretion flow as a flattened disc. We model the wind as a series of individual clouds which transit the flow (from left to right in Fig.8.1a, co-rotating with the flow) and obscure the intrinsic emission. We assume identical spherical clouds. The model has five free parameters: cloud transit time, cloud radius, cloud number density, transit latitude and cloud ionisation. In the following sections (§8.2–8.3) we investigate the effect of occultations with no intrinsic variations in luminosity of the accretion flow components, i.e. we assume constant flux from the underlying accretion flow. In all cases we assume an inclination angle of the accretion flow with respect to the line of sight of $70°$ and black hole mass of $10^7 \, M_\odot$ and define a soft band from $0.3 - 1 \, \mathrm{keV}$ and hard band from $2 - 5 \, \mathrm{keV}$ (shaded in red and blue respectively in Fig.8.1b).
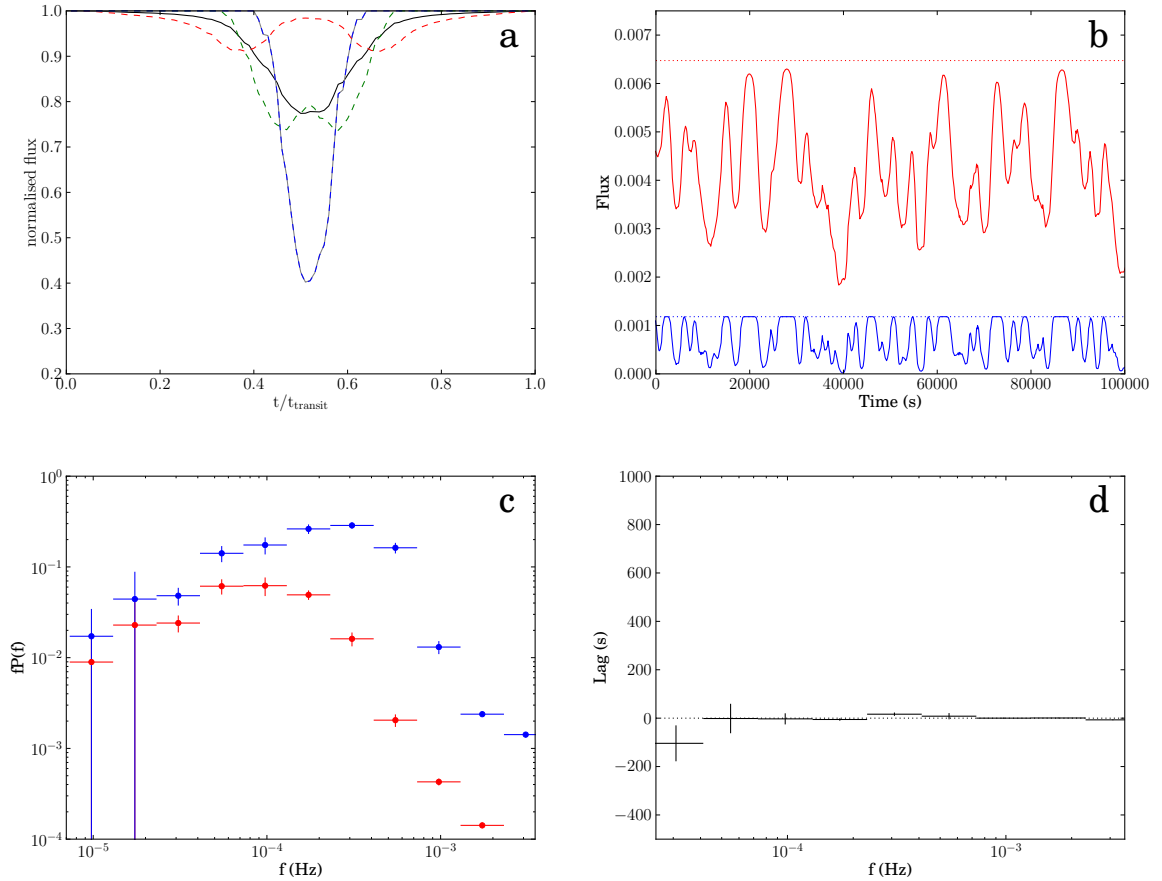
Figure 8.2: Model occulting a static accretion flow. a). Fractional flux drop for a single occultation as a function of total transit time: disc (dashed red), soft excess (dashed green), corona (dashed blue), hard band $(2 - 5 \, \mathrm{keV}$, solid grey; same shape as coronal flux drop), soft band $(0.3 - 1 \, \mathrm{keV}$, solid black). b). Sample hard (blue) and soft (red) model light curves, showing effect of occultations. Dotted lines show unobscured flux level. c). Hard and soft band power spectra (blue and red respectively). d). Lag-frequency spectrum between hard and soft bands.

8.2.1   Effect of Doppler Boosting

Fig.8.2a shows the fractional flux drop as a single cloud passes across a static accretion flow. We assume the cloud reduces the observed flux by $e^{-\tau}$, where the optical depth $\tau = 1$ and does not vary with the energy of the incident radiation, i.e. the material is completely ionised.

The dashed red line shows the drop in disc flux. Since the apparent size of the disc is much larger than the cloud the drop in flux is small ($< 10\%$). We assume a radial emissivity profile for the disc of $\varepsilon(r) \propto r^{-3}$, so when the cloud occults the outer parts of the disc the flux drop is small. The flux drop increases as the transit progresses and the cloud begins to occult the brighter inner disc radii. The disc flux recovers during the middle of the transit, as the cloud passes over the innermost regions occupied by the corona and soft excess, and then drops again as the cloud crosses the far side of the disc. The dashed green line shows the drop in flux from the soft excess. This shows similar behaviour, but more centrally concentrated, since the soft excess region is smaller. The corona shows the biggest drop in flux (dashed blue line), being a similar size to the occulting cloud. The solid grey and black lines show the flux drop in the hard and soft bands respectively. The hard band is dominated by emission from the corona and hence shows a greater flux drop than the soft band, which is dominated by emission from the disc and soft excess.

Fig.8.2b shows sample hard and soft band light curves (in blue and red respectively), showing the effect of multiple occultations. Dotted lines show the unobscured flux levels. The light curves are generated by allocating each cloud a random start time for its transit of the disc. Each cloud has a radius of $5\,R_g$ and takes $T_{tr} = 10^4\,\text{s}$ to cross the accretion flow (from left to right as seen in Fig.8.1a). On careful inspection of the light curves it can be seen that the width of the occultations is narrower in the hard band than the soft band, due to the smaller physical size of the corona, which is the main contributor of hard band flux.

Fig.8.2c shows the corresponding power spectra. The occultations add power to the light curve at a frequency that is related to the transit time. In this simulation all clouds were given a transit time of $T_{tr} = 10^4\,\text{s}$, corresponding to a frequency of $10^{-4}\,\text{Hz}$. The power in the hard band peaks at a slightly higher frequency and

greater amplitude. This is a direct result of the the shape of the flux drops shown in Fig.8.2a. The more compact coronal emission experiences a narrower, deeper flux drop than the more extended soft band components, hence occultations add more power to the hard band and at higher frequencies. The width of the hard band flux drop is $\sim 0.2 T_{tr} \sim 2 \times 10^3\,\mathrm{s}$, which corresponds to a frequency of $5 \times 10^{-4}\,\mathrm{Hz}$. Hence the power drops off sharply above $5 \times 10^{-4}\,\mathrm{Hz}$. Nevertheless there is a low power tail extending to higher frequencies in both bands. A single occultation cannot add power at these frequencies. This power comes from the superposition of occultations. The hard and soft band light curves in Fig.8.2b show that multiple occultations close together can add variability on timescales much shorter than that of an individual transit.

Fig.8.2d shows the lag as a function of frequency between the hard and soft bands (calculated following Nowak et al., 1999). There is no lag at any frequency. This is because, even though the soft band flux drop is wider than the hard band, they are both symmetric around a common centre. Even though the soft band flux drops before the hard, the hard flux then recovers before the soft with the same time delay, cancelling out any net lag. This is the case for a stationary disc.

However the accretion flow is not stationary. Material should be rotating at the Keplerian frequency. As a consequence material on one side of the flow is travelling towards the observer and Doppler boosted, while on the other side the emission is deboosted. Fig.8.3 shows the resulting flux drops, light curves, power spectra and lags now including the effect of this Doppler boosting.

Fig.8.3a shows that the flux drop in each component is now no longer symmetric. Doppler boosting means that the approaching side of the accretion flow appears brighter than the receding side. We assume the occulting clouds are co-rotating with the flow. The approaching side of the flow, which now contributes a greater fraction of the total flux, is occulted first. Hence the first half of the transit shows a much stronger flux drop. The receding side of the flow contributes much less flux to the total spectrum, hence there is a much smaller flux drop during the second half of the transit. This effect is more noticeable in the more extended components — the disc and soft excess — and most noticeable in the soft excess, where the smaller
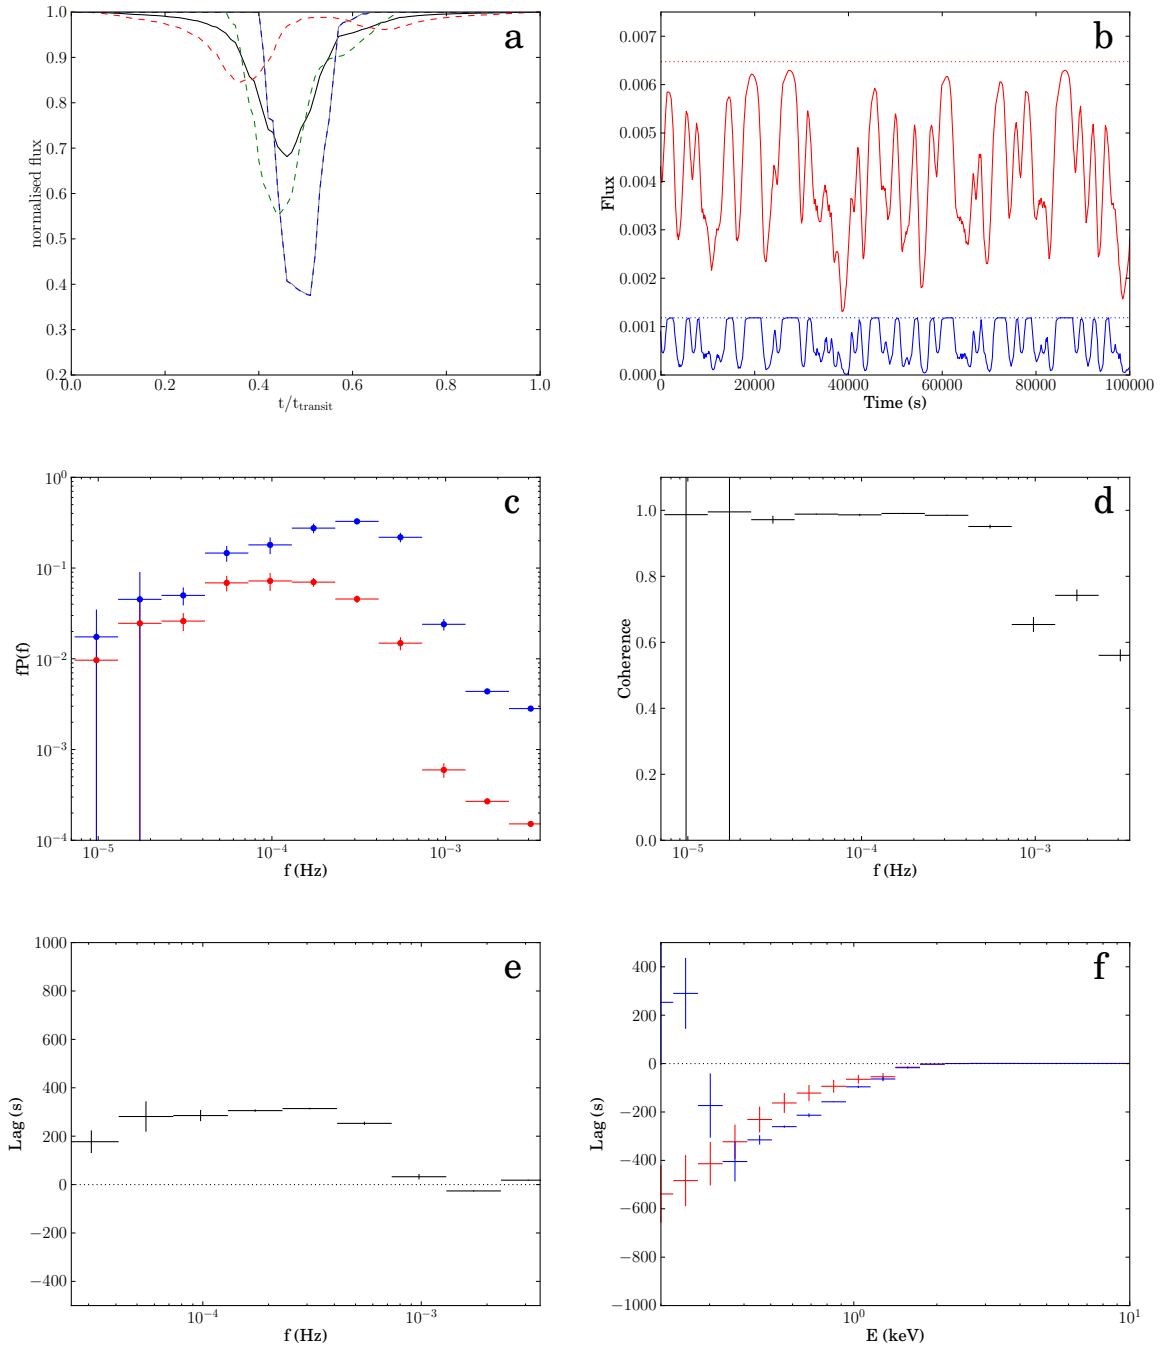
Figure 8.3: Occultation model including effects of Doppler boosting. a). Fractional flux drop for a single occultation as a function of total transit time: disc (dashed red), soft excess (dashed green), corona (dashed blue), hard band ($2-5$ keV, solid grey; same shape as coronal flux drop), soft band ($0.3-1$ keV, solid black). b). Sample hard (blue) and soft (red) model light curves, showing effect of occultations. Dotted lines show unobscured flux level. c). Hard and soft band power spectra (blue and red respectively). d). Coherence between hard and soft energy bands. e). Lag-frequency spectrum between hard and soft bands. f). Lag-energy spectrum calculated using $2-5$ keV reference band. Red points show energy spectrum of lag at low frequencies ($2.3 \times 10^{-5} - 7.3 \times 10^{-5}$ Hz), blue points show the lag-energy spectrum at high frequencies ($2.3 \times 10^{-4} - 7.3 \times 10^{-4}$ Hz).

radii give faster radial velocities and stronger Doppler boosting/deboosting. The flux drops in the hard and soft bands (grey and black solid lines) are consequently skewed towards the first half of the transit, with the soft band being more strongly skewed. On closer inspection of Fig.8.3a, it can be seen that the soft band flux now drops before the hard flux *and recovers* before the hard flux. The occultations themselves have introduced a lag between the hard and soft energy bands.

Fig.8.3e shows the lag as a function of frequency between the hard and soft energy bands. The only source of variability in the light curves is the occultations, we keep the intrinsic flux from the accretion flow constant. By including the effects of Doppler boosting, the occultations have introduced a soft lead of $\sim 300\,\mathrm{s}$ (a positive lag value indicates the soft band leading the hard). This lead remains roughly constant for the frequency range over which the occultations introduce power into the light curves ($\sim 5 \times 10^{-5} - 5 \times 10^{-4}\,\mathrm{Hz}$).

Fig.8.3f shows the lag as a function of energy. This is constructed by choosing a reference band (in this case the hard band, $2-5\,\mathrm{keV}$) and then dividing the spectrum into a series of energy bins. The flux in each energy bin is summed up as a function of time to create a light curve for that energy bin. The light curve of the energy bin is then compared with the reference band light curve and the value of the lag between the two is computed as a function of frequency. In Fig.8.3f we plot the value of the lag from each energy bin for two frequency ranges: low frequency (red points, $2.3 \times 10^{-5} - 7.3 \times 10^{-5}\,\mathrm{Hz}$) and high frequency (blue points, $2.3 \times 10^{-4} - 7.3 \times 10^{-4}\,\mathrm{Hz}$). This gives the energy spectrum of the lag at that frequency. For each energy bin, a negative lag value implies that energy bin leads the hard reference band.

At low frequencies, energy bins below $1\,\mathrm{keV}$ lead the hard reference band, with a lead that increases as the energy of the bin decreases (red points, Fig.8.3f). These are the energies at which the disc and soft excess dominate. The soft excess emission peaks at $\sim 0.5\,\mathrm{keV}$, giving way to the disc at lower energies. Since the clouds occult the outermost components first, crossing first the disc and then the soft excess before passing in front of the corona, the low energy disc emission shows the strongest lead, giving way to a slightly shorter lead from the soft excess at smaller radii and higher energies. Above $1\,\mathrm{keV}$ the emission is dominated by the corona. These are the

energies also covered by the reference band $(2 - 5\,\text{keV})$. Hence the lag of the energy bins with respect to the reference band tends to zero at high energies $(> 1\,\text{keV})$.

This pattern of soft leads (at low frequencies) is a signature normally associated with propagation; low frequency fluctuations are generated in the cooler outer components and propagate down to the hotter smaller radii which produce the high energy emission. Yet in this scenario we have produced soft leads simply by the movement of absorbing clouds, i.e. occultations affect the lag-frequency and lag-energy spectra the same way as propagation, by introducing soft leads.

Soft leads due to propagation of fluctuations are generally confined to low frequencies, since large radii only generate slow fluctuations. The blue points in Fig.8.3f show that occultations can continue producing strong soft leads up to much higher frequencies. This is because the strength and frequency of the soft leads are no longer determined by the properties of the accretion flow but by the properties of the transiting clouds. However, comparison of the red and blue points in Fig.8.3f shows that, while all energy bins below 1keV show a soft lead at low frequencies, at high frequencies the two lowest energy bins $(< 0.3\,\text{keV})$ switch from a soft lead to a soft lag. These two energy bins are dominated by disc emission. High frequency variability results from short timescale features in the flux drops shown in Fig.8.3a, implying this soft lag confined to very low energy comes from the cloud covering the deboosted side of the disc after covering the corona.

For completeness we also show the coherence between hard and soft energy bands (Fig.8.3d), where 1 is perfect coherence between the two bands and 0 is incoherence. The coherence remains high up to high frequencies, since occultation is the only source of variability in the two light curves and is common to both. The coherence drops off above $\sim 5 \times 10^{-4}\,\text{Hz}$, where the variability power introduced by the occultations also drops off. $5 \times 10^{-4}\,\text{Hz}$ corresponds roughly to the width of the hard band flux drops $(\sim 0.2 T_{tr} \sim 2 \times 10^{3}\,\text{s})$.

8.2.2 Effect of Transit Time

Fig.8.4 shows the effect of changing the cloud transit time. We increase the transit time from $T_{tr} = 5 \times 10^{3}\,\text{s}$ (black) to $1.5 \times 10^{4}\,\text{s}$ (magenta). The hard band power
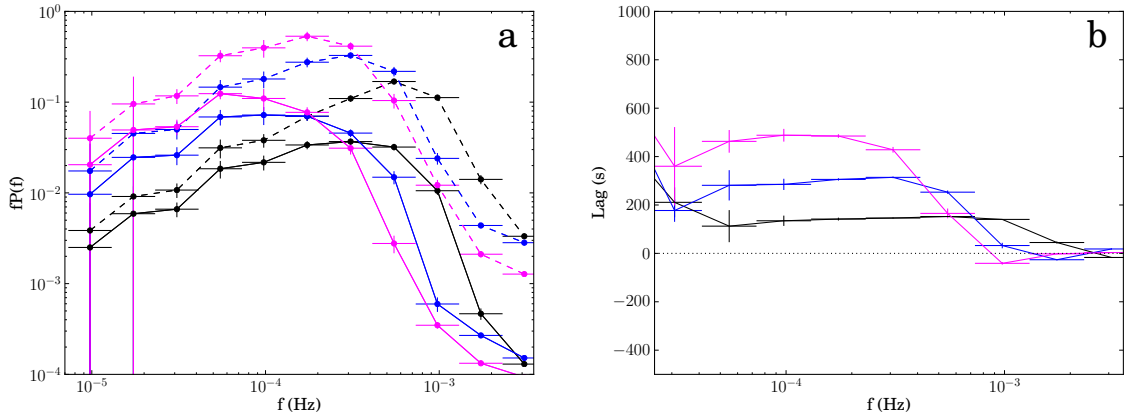
Figure 8.4: Effect of increasing cloud transit time on a). power spectrum and b). lag-frequency spectrum, for $T_{tr} = 5 \times 10^3$ (black), 10^4 (blue) and 1.5×10^4 s (magenta). Solid lines show soft band $(0.3 - 1 \, \mathrm{keV})$ power spectra, dashed lines show hard band $(2 - 5 \, \mathrm{keV})$ power spectra.

spectra (dashed lines, Fig.8.4a) clearly show that as the transit time increases, the peak frequency at which power is added to the light curves decreases. The decrease in peak frequency of approximately half an order of magnitude roughly matches the threefold increase in transit time. The total amount of power added to the light curves also increases by a similar amount in both bands.

Fig.8.4b shows how this affects the lag measured between the hard and soft band light curves. For short transit times (5×10^3 s, black), a short lag (~ 150 s) is measured up to high frequencies (10^{-3} Hz). As the transit time increases, the maximum frequency at which a lag is measured decreases. This is because occultations with a longer transit time cannot add high frequency power to the light curves (as shown by the power spectra in Fig.8.4a). The absolute value of the lag also increases, since for longer transit times the clouds spend longer occulting the outer soft components before they cross and occult the central corona. The measured lag drops from ~ 550 s to 150 s, roughly matching the decrease in transit time of a third from 1.5×10^4 to 5×10^3 s.

We note that increasing black hole mass or increasing the size scales of the individual components has a similar effect on the power spectrum and variability as increasing the transit time.

Figure 8.5: Effect of increasing cloud radius on a). power spectrum and b). lag-frequency spectrum, for $R_{cl} = 5$ (black), 10 (blue) and $15\,R_g$ (magenta). Solid lines show soft band $(0.3 - 1\,\mathrm{keV})$ power spectra, dashed lines show hard band $(2 - 5\,\mathrm{keV})$ power spectra.

8.2.3   Effect of Cloud Radius

We now fix the transit time at $10^4\,\mathrm{s}$ and investigate the effect of changing the cloud radius. Fig.8.5 shows the resulting power spectra and lag-frequency spectra for $R_{cl} = 5$, 10 and $15\,R_g$ (black, blue and magenta lines, respectively).

Increasing the cloud radius increases the amount of power in the light curves. The soft band shows the biggest increase (Fig.8.5a, solid lines), with the amount of power at $10^{-4}\,\mathrm{Hz}$ increasing by nearly one and a half orders of magnitude. The effect is much less in the corona dominated hard band (just under an order of magnitude), because the corona is much smaller, so that it is already completely obscured by a small cloud of $5\,R_g$. Increasing the cloud radius only prolongs the length of time it is obscured. By contrast the much larger disc is never completely obscured by a $5\,R_g$ cloud. Increasing the cloud radius therefore increases the area of the disc that experiences obscuration and hence adds more power to the soft band light curve.

The more noticeable change to the hard band light curve is that the frequency at which the hard band power peaks decreases as the cloud radius increases (Fig.8.5a, dashed lines). This is because, for a larger cloud radius (and fixed transit time), the time taken between covering and uncovering the corona increases. Consequently the transit cannot add as much high frequency power to the light curve. The peak in power drops from $\sim 3 \times 10^{-4}$ to $2 \times 10^{-4}\,\mathrm{Hz}$, as more power is added at low
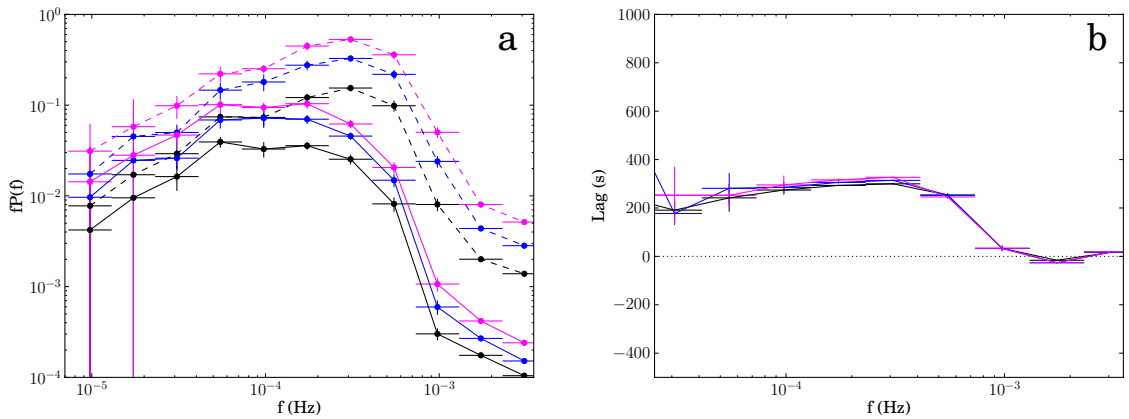
Figure 8.6: Effect of increasing number of clouds on a). power spectrum and b). lag-frequency spectrum, for $n_{cl} \sim 5 \times 10^{-4}$ (black), 10^{-3} (blue) and 1.5×10^{-3} (magenta) clouds per second. Solid lines show soft band $(0.3 - 1\,\text{keV})$ power spectra, dashed lines show hard band $(2 - 5\,\text{keV})$ power spectra.

frequencies and lost at high frequencies.

Fig.8.5b shows the lag as a function of frequency between the hard and soft band light curves. The lag measured actually slightly decreases as cloud radius increases. This is because the larger the cloud the more time it spends obscuring hard and soft components simultaneously. This results in very broad, very similar flux drops in both the hard and soft bands. In contrast, the strongest soft leads are seen when the cloud is small enough to obscure the blue wing of the disc and soft excess and then the corona in turn. This results in much narrower flux drops in the hard and soft bands, where the skew due to Doppler boosting (which causes the soft lead) is much more prominent.

8.2.4   Effect of Cloud Number Density

Fig.8.6 shows the effect of increasing the number of occulting clouds. We fix the cloud radius and transit time at $5\,R_g$ and $10^4\,\text{s}$ and increase the number density of clouds from $n_{cl} \sim 5 \times 10^{-4}$ (black) to $1.5 \times 10^{-3}\,\text{s}^{-1}$ (magenta). Our total simulation time is $1.024\,\text{Ms}$, in practice this corresponds to increasing the total number of occulting clouds from 500 to 1500 clouds, each of which is assigned a random start time for its transit.

Fig.8.6a shows that increasing the number of clouds increases the power spectrum

Figure 8.7: Effect of increasing transit latitude on a). power spectrum and b). lag-frequency spectrum, for $z_{cl} = 0$ (black), 5 (blue) and 8 R_g (magenta), where z_{cl} defines the apparent latitude of the centre of the clouds as shown in Fig.8.1a. Solid lines show soft band $(0.3 - 1 \, \text{keV})$ power spectra, dashed lines show hard band $(2 - 5 \, \text{keV})$ power spectra.

normalisation without affecting its shape. This is because more occultations simply add more power to the light curve. Both the hard and soft bands are affected equally. Fig.8.6b shows that increasing the number of occultations has no effect on the lag measured between the hard and soft band light curves. The value of the lag is determined primarily by the transit time, with a weak dependence on cloud radius.

8.2.5   Effect of Transit Latitude

In Fig.8.7 we show the effect of increasing the latitude of the cloud path so that it no longer aligns exactly with the black hole. In all three cases we fix the cloud radius and transit time to $5 \, R_g$ and $10^4 \, \text{s}$ and the number density of clouds to $10^{-4} \, \text{s}^{-1}$. We increase the apparent latitude of the cloud center from $z_{cl} = 0 \, R_g$ (black) to $z_{cl} = 8 \, R_g$ (magenta). Fig.8.8 shows as an illustration the path taken by a cloud transiting at $z_{cl} = 8 \, R_g$, where the dashed lines bound the region experiencing obscuration.

Fig.8.7a shows the effect on the hard and soft band power spectra. As the latitude increases the cloud only obscures the 'back half' of the accretion flow ($z > 0 \, R_g$ in Fig.8.8). Consequently the power in the light curves decreases. By the time the latitude has increased to $8 \, R_g$ the cloud no longer obscures the central corona, hence

Figure 8.8: Dashed lines show path of cloud across accretion flow for a high transit latitude ($z_{cl} = 8\,R_g$), with disc (red), soft excess (green), corona (blue), obscuring cloud (black) and accretion flow viewed at $i = 70°$.

there is almost no power in the hard band (hence no dashed magenta line), since we have fixed the underlying accretion flow to have constant flux. Consequently high latitude transits that do not occult the corona add power to the soft band that is uncorrelated with the hard band.

The soft band power drops from $\sim 7 \times 10^{-2}$ for a central transit (solid black line) to $\sim 10^{-3}$ at the highest latitude (solid magenta line). Comparison of the black and magenta solid lines shows that, not only has total power been lost, but power has preferentially been lost at high frequencies. Below $5 \times 10^{-5}\,$Hz the two soft band power spectra show the same shape, whilst above $5 \times 10^{-5}\,$Hz the high latitude power spectrum shows a cut off, with the power dropping off sharply above $10^{-4}\,$Hz. This is because the transiting cloud just clips the very edge of the disc and soft excess. These regions of the flow are travelling nearly perpendicular to the line of sight so experience very little Doppler boosting/deboosting. They therefore carry only a moderate fraction of the total soft band flux. The shortest, sharpest dips in the soft band light curve arise through the cloud occulting the innermost parts of the disc/soft excess, which are centrally concentrated and strongly Doppler boosted. These add the highest frequency components to the soft band light curve. Occultations of the outer parts of the flow result in slower more gradual flux drops

and hence add low frequency power. The coronal power law also contributes some flux to the soft band, so with no coronal occultations that removes its additional source of high frequency power.

Fig.8.7b shows the effect of increasing the transit latitude on the lag measured between the hard and soft bands. As high frequency power is lost from both the hard and soft bands, the measured lag begins to tend to zero at lower frequencies ($\sim 2 \times 10^{-4}$ Hz for the high latitude magenta spectrum, compared to 7×10^{-4} Hz for the central transit back spectrum). The lag tends more gradually to zero in the case of the high latitude transit. This is again because the highest frequency components come from the shortest, sharpest flux drops which arise from occulting the brightest central regions. For a high latitude transit this is when the cloud just clips the small part of the soft excess that appears at high latitude (i.e. large z and $y = 0\,R_g$ in Fig.8.8) at the midpoint of the transit. Because the coronal flux is not occulted, occultation of the soft excess is the only source of variability in the hard band. It is also the only source of high frequency variability in the soft band. When the source of variability is the same in both bands, there can be no lag between them. Hence as frequency increases above $\sim 2 \times 10^{-4}$ Hz the lag tends gradually to zero as the only source of variability becomes occultation of the soft excess in both bands.

8.2.6 Effect of Cloud Ionisation

So far we have approximated the absorption of the cloud as $e^{-\tau}$, where the electron scattering optical depth, $\tau = 1$ (equivalent to a pure hydrogen column of 1.5×10^{24} cm^{-2}, which is 1.25×10^{24} cm^{-2} for solar abundance material), is a constant with energy. However this is only appropriate for completely ionised material. In general, the optical depth of the cloud is a function of energy, depending on the ionisation state of the cloud.

In Fig.8.9a we show our original model (dotted black line) compared to transmission spectra for $N_H = 1.5 \times 10^{24}$ cm^{-2} at $\log \xi = 4$ (magenta line) and at $\xi = 0$ (solid black line). The $\log \xi = 4$ transmission spectrum is calculated using the XSPEC model ZXIPCF. This assumes a turbulent velocity of $200\,\mathrm{km\,s^{-1}}$, so the line strength can be enhanced for the same column density of material for higher veloci-

Figure 8.9: Effect of changing cloud ionisation on a). transmission spectrum, b). power spectrum and c). lag-frequency spectrum, for two cloud ionisation states: $\xi = 0$ (black) and $\log \xi = 4$ (magenta). Solid lines show soft band $(0.3 - 1 \, \text{keV})$ power spectra, dashed lines show hard band $(2 - 5 \, \text{keV})$ power spectra. Dotted line shows transmitted flux level for constant $e^{-\tau}$, $\tau = 1$.

ties. The neutral spectrum is calculated using PHABS. As before, we fix $R_{cl} = 5\,R_g$, $T_{tr} = 10^4\,\text{s}$, $n_{cl} = 10^{-4}\,\text{s}^{-1}$ and $z_{cl} = 0\,R_g$. Having fixed both the optical depth and cloud radius, this constrains the cloud density, since $n = \tau/(2R_{cl}\sigma_T) \sim 10^{11}\,\text{cm}^{-3}$. The ionisation state of the clouds is related to their density and the X-ray luminosity as $\xi = L_X/(nD^2)$, where D is the distance of the cloud from the central X-ray source. For $L_X = 10^{42}\,\text{erg s}^{-1}$ and $D = 20\,R_g$, $\xi \sim 10^4$.

Fig.8.9b shows that increasing the ionisation state of the cloud reduces the amount of power added to the hard and soft band light curves. The difference is roughly half an order of magnitude in both bands, since the change in opacity is roughly the same for both bands. For a neutral cloud, the transmitted flux below $5\,\text{keV}$ is zero. In contrast, when the cloud is highly ionised there is very little absorption left, so the fraction of transmitted flux rises to $e^{-\tau} \sim 0.3$. This results in shallower flux drops and hence less power in the hard and soft band light curves.

Fig.8.9c shows that changing the ionisation state of the clouds has no effect on the lag measured between the hard and soft energy bands. The lag depends on the motion of the cloud with time, not on the relative amounts of power carried by the hard and soft bands.

8.3   Circular Occultation

So far we have modelled linear occultations, where the apparent velocity of the clouds remains constant during the transit. However the clouds should be rotating with the accretion flow, in which case their apparent velocity during the transit will vary as a cosine function. During the middle of the transit the component of the cloud's velocity perpendicular to the line of sight is greatest and the cloud appears to move faster. At the beginning and end of the transit the cloud is moving towards/away from the observer, the component of its velocity perpendicular to the line of sight is small and its apparent velocity is much slower. If the orbital radius of the clouds is much larger than the radius of the region being occulted then linear occultation is a reasonable approximation. As the orbital radius of the clouds becomes similar to the occulting region size the effect becomes more important. A transit time of

Figure 8.10: Change in a). power spectrum and b). lag-frequency spectrum when circular motion of clouds is taken into account (magenta), compared with linear occultations (black). Solid lines show soft band $(0.3 - 1 \, \text{keV})$ power spectra, dashed lines show hard band $(2 - 5 \, \text{keV})$ power spectra.

$10^4 \, \text{s}$ corresponds to a Keplerian orbital velocity at $\sim 20 \, R_g$ for a $10^7 \, M_\odot$ black hole, implying this effect should be taken into account.

Fig.8.10 shows the effect on the power spectra and lag-frequency spectrum. We fix $R_{cl} = 5 \, R_g$, $T_{tr} = 10^4 \, \text{s}$, $n_{cl} = 10^{-4} s^{-1}$ and $z_{cl} = 0 \, R_g$ and show the result of linear occultations in black and accounting for circular motion in magenta. In the circular case the cloud moves faster while it is occulting the brightest central region of the accretion flow. As a result the flux drops are narrower. This adds more power at higher frequencies, hence both the hard and soft band power spectra are shifted to slightly higher frequencies for the case of circular occultations (Fig.8.10a).

Similarly the lag-frequency spectrum extends to slightly higher frequencies in the circular motion case, showing a non-zero lag up to $10^{-3} \, \text{Hz}$ compared to $7 \times 10^{-4} \, \text{Hz}$ for linear occultations (Fig.8.10b). The measured lag is also shorter ($\sim 200 \, \text{s}$ compared to $300 \, \text{s}$ for linear occultations). Again this is a consequence of the cloud moving faster during the central part of the occultation. The lag predominantly arises from the delay between occulting soft excess and then coronal emission on the Doppler boosted side of the flow and when the cloud is moving faster the delay is shorter.

Figure 8.11: a). Schematic of the simple NLS1 model of GD14. Slow fluctuations propagate inwards from the outer components and are modulated by the faster fluctuations generated at smaller radii. The high energy coronal emission then reflects off and is reprocessed by the soft excess component. b). Scenario for the transition from simple to complex NLS1 as a function of inclination, where the soft excess is a turbulent region of rotating clouds which partially obscures the line of sight to the central regions in complex NLS1s. The clouds are then responsible for the bulk of the reflected/reprocessed emission, while partially obscuring the intrinsic emission.

8.4 Transition from Simple to Complex NLS1 by Including Occulting Clouds

We now investigate whether the addition of occultations can change the timing properties of a simple NLS1 so that they appear more typical of complex NLS1s. That is, can the effect of occultations reduce the maximum measured reverberation lag from $\sim 200\,\mathrm{s}$ to nearer $50\,\mathrm{s}$ and shift it to higher frequencies?

In the previous sections we assumed constant flux from the underlying accretion flow. We now replace this static model with the time dependent model of the previous chapter shown in Fig.8.11a (GD14). The disc, which is at the largest radii, generates the slowest fluctuations. These propagate down to the soft excess, which is at smaller radii and generates its own slightly faster fluctuations. The fluctuations in soft excess emission therefore consist of the slow fluctuations from the disc, delayed by some lag related to the propagation time, modulated by the faster fluctuations generated in the soft excess. These fluctuations then propagate down to the corona, which generates even faster fluctuations. The hard coronal emission therefore shows fluctuations on a whole range of timescales, as it responds to mass accretion rate fluctuations propagating down from all radii. A fraction of these central hard X-rays illuminate the cooler soft excess and disc components. Some of this illuminating flux will be reflected, the rest will thermalise and be reprocessed. Fig.8.12 shows our spectral decomposition now including these reflected and reprocessed components. These come from the fit to the time averaged spectrum of the simple NLS1 PG1244+026 (OBS ID: 0675320101, shown in black data points) used in the previous chapter. For simplicity we assume all reflection/reprocessing occurs on the soft excess ($6 - 12\,R_g$). Thus the fluctuations in the reflected/reprocessed emission follow the coronal fluctuations (Fig.8.11a), except for the very fastest fluctuations which are smoothed out by the range of light travel time delays. Hence the reflected and reprocessed fluctuations are a lagged and smoothed version of the hard coronal fluctuations. The soft excess therefore consists of intrinsic emission from the accretion flow (dashed green line, Fig.8.12), which varies slowly due to intrinsic mass accretion rate fluctuations in the soft excess and those that have propagated inwards

Figure 8.12: Spectral decomposition for PG1244+026: disc (red), soft excess (green, with dashed line for intrinsic emission, dotted line for reprocessed emission, solid line for total), corona (blue), reflection (magenta), total (black). Data points show time averaged spectrum (OBS ID: 0675320101). Grey line shows an example of the spectrum after introducing absorption by intervening clouds.

from the disc, and reprocessed emission (dotted green line, Fig.8.12), which follows the faster coronal fluctuations. In the previous chapter we showed that this model can reproduce all the observed timing properties of the simple NLS1 PG1244+026 (GD14).

We use this model to describe the emission from the underlying accretion flow as a function of time, and now add the effect of occulting clouds. We fix the cloud parameters to $R_{cl} = 5\,R_g$, $T_{tr} = 10^4\,\mathrm{s}$ and $z_{cl} = 0\,R_g$. A transit time of $10^4\,\mathrm{s}$ implies an orbital radius of $20\,R_g$. Hence we reduce our transit radius from $40\,R_g$ in Fig.8.1a to $20\,R_g$ and take into account the circular motion of the clouds. Since the clouds are launched so close to the central X-ray source we allow them to be highly ionised and use the magenta transmission spectrum shown in Fig.8.9a ($\log\xi = 4$, $N_h = 1.5 \times 10^{24}\,\mathrm{cm}^{-2}$). $20\,R_g$ is consistent with the clouds being launched as part of a failed Eddington (radiation pressure driven) wind from the soft excess region as sketched in Fig.8.11b. Radiation pressure lifts material from the accretion flow, which forms clumps as it rises (Takeuchi et al., 2014). As soon as the optical depth of the clumps becomes $\tau > 1$, some of the material is self-shielded from X-ray photons. The mass of the clump is still the same but the radiation pressure on it

is now less. If the source is not strongly super-Eddington, the radiation pressure is not strong enough to expel the material so it falls back to the disc, resulting in a failed rather than outflowing wind. The whole turbulent large scale-height region is the source of the soft excess. Propagation of fluctuations occurs through the lower disc-like regions, which are the source of the intrinsic emission, while the bulk of the reflected and reprocessed emission comes from the turbulent clouds. We assume the turbulent velocity is less than the orbital velocity ($v_{turb} < v_{Kepl}$) and that the clouds remain largely intact on the timescale of a single transit (although v_{Kepl} may be sufficient to shred them on longer timescales, stripping off material before what remains falls back to the disc). For a source at high inclination, these clouds will intercept the line of sight to the central regions. As the clouds transit the line of sight, we assume they obscure the intrinsic disc, intrinsic soft excess and coronal emission. We do not obscure the reflected or reprocessed emission, since we assume these are predominantly from the clouds.

8.4.1   Fourier Timing Properties

Fig.8.13 shows the resulting soft and hard band light curves (left and right respectively). The top panels show the original simple NLS1 model with no occultations. These light curves have power spectra that match the hard and soft band power spectra of the simple NLS1 PG1244+026. In the subsequent panels we increase the number of occulting clouds ($n_{cl} = 10^{-4}$ and $10^{-3}\,\mathrm{s}^{-1}$). The occultations are most obvious in the hard band, where the flux drops are conspicuously narrower than in the soft band, due to the smaller physical size of the corona compared to the more extended soft band components. These occultations add power to the light curve. The most heavily occulted hard band light curve (bottom right panel) shows peaks and deep troughs more typical of a complex NLS1.

Fig.8.14 shows the power spectra, lag-frequency spectra and coherence between hard and soft bands for the same three simulations. Comparing the hard band power spectra (Fig.8.14a, dashed lines) of the original model (black) with the most heavily obscured model (magenta) shows that occultations have increased the power at $\sim 3 \times 10^{-4}\,\mathrm{Hz}$ by almost one and a half orders of magnitude. Hard band power

Figure 8.13: Soft band $(0.3 - 1\,\mathrm{keV}$, red) and hard band $(2 - 5\,\mathrm{keV}$, blue) light curves showing effect of adding obscuring clouds to the simple NLS1 model of GD14. Top panels: simple NLS1 model with no clouds. Middle panels: $n_{cl} = 10^{-4}\,\mathrm{s}^{-1}$. Bottom panels: $n_{cl} = 10^{-3}\,\mathrm{s}^{-1}$. For both cases we fix the cloud parameters to $R_{cl} = 5\,R_g$, $T_{tr} = 10^4\,\mathrm{s}$, $z_{cl} = 0$, $\log \xi = 4$ and $N_H = 1.5 \times 10^{24}\,\mathrm{cm}^{-2}$ and assume the clouds are launched from $\sim 20\,R_g$.

Figure 8.14: Effect of adding occultations to the simple NLS1 model of GD14 on a). hard band ($2 - 5\,\mathrm{keV}$, dashed lines) and soft band ($0.3 - 1\,\mathrm{keV}$, solid lines) power spectra, b). lag frequency spectrum and c). coherence between hard and soft bands. Black lines show simple NLS1 model with no clouds, blue $n_{cl} = 10^{-4}$ and magenta $n_{cl} = 10^{-3}\,\mathrm{s}^{-1}$, with cloud parameters as in Fig.8.13.

spectra of complex NLS1s routinely show similarly high power at these frequencies. Thus occultations are more than capable of increasing the hard band high frequency power from the $fP(f) \sim 10^{-2}$ typical of a simple NLS1 to ~ 0.1, as is typical of a complex NLS1. The power increase in the soft band (solid lines) is much smaller (roughly half an order of magnitude). However we have only included occultations at $z_{cl} = 0 \, R_g$. Higher latitude occultations would add additional power to the soft band, however this power would be uncorrelated with the hard band.

Fig.8.14b shows the lag as a function of frequency between the hard and soft bands. The black points show the original simple NLS1 model, with a strong reverberation lag of $\sim 200\,\mathrm{s}$ at $\sim 5 \times 10^{-4}\,\mathrm{Hz}$, matching that seen in PG1244+026. As the number of occultations increases (blue to magenta), the maximum measured reverberation lag decreases from $200\,\mathrm{s}$ to $50\,\mathrm{s}$ and increases in frequency from $\sim 5 \times 10^{-4}$ to $10^{-3}\,\mathrm{Hz}$. This much shorter reverberation lag, at higher frequency, is much more typical of those seen in complex NLS1s such as 1H0707–495. In our model, this is a direct result of the soft leads introduced by the occultations at low frequencies, diluting the reverberation lag and shifting its minimum to higher frequencies. This is also in good agreement with the findings of Kara et al. (2013), who showed that the reverberation lag of the complex NLS1 IRAS 13224–3809 is much shorter and at higher frequency during low flux periods (when in this scenario it would be more obscured) than high flux periods.

Fig.8.14c shows the coherence between hard and soft bands. As the number of occultations increases, the coherence drops slightly, particularly at low frequencies (from ~ 1 to ~ 0.8). This is due to the slightly differently shaped flux drops in the hard and soft bands. The sudden drop in coherence at $10^{-4}\,\mathrm{Hz}$, characteristic of the simple NLS1 PG1244+026, becomes less obvious in the more obscured simulations. The coherence functions of complex NLS1s do show a more gradual drop in coherence with increasing frequency, so this is not in disagreement with the data.

8.4.2 Spectral Changes

In Fig.8.12, in grey, we show an example of the total spectrum when the source is highly absorbed. We assume the clouds are the source of the reflected and repro-

cessed emission, hence these components are not absorbed. Consequently the low flux spectrum is dominated by reprocessed emission at low energies ($< 1\,\text{keV}$) and reflected emission at high energies. This gives a total spectrum that is no longer power law like but instead shows strong curvature, with a large soft excess and a strong iron emission line. But the tell tale signature of highly ionised occultation is the presence of highly ionised Fe Kα absorption lines at 6.7 and 6.95 keV. The strength of these features is probably underestimated in our model as the ZXIPCF model used assumes a turbulent velocity of only $200\,\text{km}\,\text{s}^{-1}$, which is probably much smaller than expected in a failed wind structure from the inner disc.

Strong highly ionised absorption lines are detected in the deep dip states of several complex NLS1s such as MRK 766 (Miller et al., 2007), MCG 6-30-15 (Miller et al., 2008) and MRK 335 (Gallo et al., 2013). It is possible that even higher turbulent velocities in the failed wind could merge the 6.7 and 6.95 keV absorption lines into each other, and into the absorption edge, which could produce the more dramatic drop at Kα in 1H0707–495 (Hagino et al., 2015).

Our assumption that the clouds do not occult the reflected and reprocessed emission sets a limit on the amount of absorption present in our model. Complex NLS1s often show rather stronger drops at low energy, which in this scenario would require that the clouds do also occult part of the failed wind structure or that our reprocessed emission is overestimated due to our models taking only the model flux rather than weighting this by the instrument response. Nonetheless, this model demonstrates the potential of this scenario to match the spectral variability, given a more sophisticated prescription for the reflected/reprocessed flux.

8.5 Discussion

We have shown that occultation of the accretion flow can introduce lags between the hard and soft energy bands. In particular the occultations generate soft leads at a frequency related to the transit time. For a transit time of $10^4\,\text{s}$, corresponding to an orbital radius of $20\,R_g$, these lags predominate at low frequencies ($\sim 10^{-4}\,\text{Hz}$ for a $10^7\,M_\odot$ black hole). When combined with a model for the intrinsic variability of

the accretion flow, these low frequency soft leads act to dilute the negative reverberation lag, reducing the maximum measured lag and shifting it to higher frequencies, reproducing the trend seen in the data from simple to complex NLS1. This can plausibly be produced by increasing inclination with respect to a clumpy, turbulent structure above the inner disc. Since NLS1s are high accretion rate sources, it is likely that the inner regions of the disc will become 'puffed up' to a large scale-height where the local accretion rate exceeds the Eddington limit (e.g. Jiang et al., 2014), producing a turbulent structure capable of obscuring the innermost regions. Low inclination sources are rarely occulted (simple NLS1s), while high inclination ones have multiple occultations (complex NLS1s). The assumed higher inclination for complex NLS1s also explains the observed association of extreme (deep dip) spectra with warm absorbers. This is not causal, in that the warm absorbers are not distorting our view of the extreme spectra (e.g. Chiang & Fabian, 2011), but is instead a consequence of high inclination so that the line of sight is more likely to intercept a wind driven from the torus/flattened broad line region (BLR). Importantly, the occultations superimpose highly ionised absorption lines at FeKα in the dips, as are seen in the data.

We show the evolution of the lag-frequency spectrum with increased occultations for a fixed black hole mass of $10^7 \, M_\odot$, but this should also depend on black hole mass. The intrinsic lags/leads in the spectral components should scale simply with mass, as should the occultation timescale. However, mass for NLS1s is hard to determine accurately, as they are accreting close to Eddington. Masses estimated from line widths assume the BLR clouds are virialised. However the effective gravity experienced by the clouds will be reduced due to radiation pressure from the central source, leading to an underestimate for masses of NLS1s (Marconi et al., 2008). Inclination is also another uncertainty, as the BLR velocity field is not completely virialised, but contains a clear equatorial component (e.g. Collin et al., 2006; Kollatschny & Zetzl, 2013; Pancoast et al., 2014). Any equatorial component to the velocity field will be suppressed in low inclination (simple) NLS1s, so their masses will be systematically biased towards higher values compared to high inclination (complex) NLS1s.

Leighly (1999) gives the FWHM for the Hβ line widths in PG1244+026 and 1H0707–495 as 830 and 1050 km s^{-1}, respectively. These are corrected for FeII and have the narrow Hβ component subtracted assuming that this is $0.1\times$ the [OIII] line intensity (see also Leighly & Moore, 2004). Both sources have intrinsic optical luminosities which are very similar, so these give masses which are 2.4×10^6 and $3.7 \times 10^6 \, M_\odot$, respectively (Nikołajuk et al., 2009). The two object masses may be even closer if the inclination dependence discussed above is important. An Eddington correction to the mass of PG1244+026 increases the mass estimate to $10^7 \, M_\odot$ (Done et al., 2013), but this should also be similar for 1H0707–495. Hence, while there are large uncertainties on masses for NLS1s, these two objects should be very similar. At this larger mass, PG1244+026 is at the Eddington limit for a low spin black hole (Done et al., 2013). A lower mass and/or higher spin pushes the system to higher Eddington fractions, so making it even less likely that the disc is flat.

Our model is more of a pilot study than a complete description. Obvious improvements are to include general relativistic effects of light bending on the disc image (Fig.8.1a), e.g. Miniutti et al. (2014), Middleton & Ingram (2015). This would be most important for the central coronal regions, as the far side of this small source always has a small impact parameter with the black hole. This would make the corona appear larger, so a $5\,R_g$ cloud may not occult the entire corona. As a consequence the coronal flux drops would not be quite as narrow and deep and this would reduce the difference between the total amount of power added to the hard and soft bands. Complex NLS1s do show more power in the hard band than the soft band (Zoghbi et al., 2011) at all frequencies, unlike the simple NLS1 PG1244+026, which shows comparable power at low frequencies in the hard and soft bands. Our occultation model replicates this, since occultations add power to both bands. However the small size of the corona necessarily adds much more power to the hard band. Including light-bending (or a smaller cloud size) would slightly lessen this difference, in better agreement with the data.

The model presented here is in some way a composite between the previous extreme relativistic reflection models and partial covering models. It follows the

partial covering model in identifying absorption (as opposed to light bending) events as the origin for the deep dips, but has the occulting material be closer to the source (ten gravitational radii rather than a few tens–hundreds), and be more highly ionised. Reflection does make an important contribution to the spectrum during the dips in our model, but it is not extremely smeared by relativistic effects. Instead, and in a step beyond what is modelled here, we envisage the reflector as a clumpy, turbulent, failed wind rather than a flat Keplerian disc (see also Miller et al., 2008). The clumps may be only marginally optically thick, so their reflected/scattered emission is not quite the same shape as from $\tau \gg 1$ (Miller & Turner, 2013), and they may be embedded in hotter material which Comptonises the reflected emission. Additionally, the cloud itself could have a complex structure due to the ionisation instability of X-ray illuminated material (Krolik et al., 1981). The illuminated face of the cloud will be heated to the local Compton temperature ($\sim 10^6$ K). Temperature decreases at larger depths into the cloud, where scattering reduces the heating, so the density must increase to keep in pressure balance. This lower ionisation state material has more line cooling, so the temperature drops abruptly, giving a sharp transition between a highly ionised skin and a nearly neutral core (Chevallier et al., 2006). Reflection from such structures, especially with a turbulent velocity field, may be a feasible way to reproduce the observed $2 - 10$ keV spectra in the dips.

Another difference between this model and standard relativistic reflection is that we include thermalisation of the illuminating flux. Hard X-rays which are not reflected can either heat the disc and be re-emitted as (mostly) thermalised radiation, or the energy can be released as lines/recombination continua. The relative importance of these two processes depends strongly on the vertical structure of the disc. Thermalisation is more important for discs in hydrostatic equilibrium (Nayakshin & Kallman, 2001) but current reflection models are calculated for constant density discs (e.g. Ross & Fabian, 2005; García et al., 2013). The high disc temperatures expected in NLS1s means that this component must be important at some level in contributing to the soft X-ray excess in these objects, and since it is predominantly thermal then it has no strong soft X-rays lines which require high spin to smear them into the observed smooth continuum.

However, probably the most important effect which should be included in matching to real data is that the light curves in the soft and hard bands are weighted by the detector response rather than being simply flux integrated over the energy band as used here. This is a key requirement to fit the model to real data in future work.

8.6   Conclusions

We have constructed a simple occultation model to investigate whether the change in spectral and timing properties between simple and complex NLS1s can be explained by a difference in inclination with respect to a failed wind. In this scenario, clumps of material lifted from the inner parts of the accretion disc by radiation pressure obscure the X-ray emission for sources seen at high inclination angles, resulting in more extreme variability and more complex spectra. Associating the deep dips with occultation superimposes absorption features from FeKα on the dip spectra. This is seen in complex NLS1s such as MRK 776 (Miller et al., 2007) and is a clear indication of the presence of these occultation events.

We model the obscuration as a series of individual clouds of constant ionisation parameter which transit the inner accretion flow, co-rotating with the flow and obscure the underlying emission. The underlying accretion flow emission is radially stratified, with the softest X-rays (disc) from the largest radii, and then the soft X-ray excess and corona at progressively smaller radii. We find that occultations add power to the X-ray light curves over a range of frequencies related to the transit time. Occultations also introduce a lag between the hard and soft energy bands when Doppler boosting of the underlying accretion flow emission is taken into account; specifically occultations introduce a soft lead, with the hard band lagging the soft band.

We then combined our occultation model with the full spectral-timing model of GD14 which describes the accretion flow emission of the simple NLS1 PG1244+026. This model also includes reprocessed emission as part of the soft X-ray excess, as well as reflection from it. This reproduces the timing properties of PG1244+026 by assuming slow fluctuations are generated in the outer components and these prop-

agate down to the corona, producing low frequency hard lags. The high frequency soft lags (reverberation lags) are produced predominantly by fast coronal fluctuations being reprocessed in the soft excess wind material rather than by reflection from it.

The effect of the occultations is to dilute the negative reverberation lag and shift it to higher frequencies. By increasing the rate of occultations we can match the change in reverberation lag from the $200\,\mathrm{s}$ at $5 \times 10^{-4}\,\mathrm{Hz}$ seen in the simple NLS1 PG1244+026, to the $\sim 50\,\mathrm{s}$ lag at $10^{-3}\,\mathrm{Hz}$ seen in the complex NLS1 1H0707–495. The lag times and light travel times put into the model are the same in both the obscured and unobscured cases. The only difference is the presence of occultations. It is the soft leads caused by the broken symmetry of the flux drops, due to Doppler boosting of the underlying occulted disc emission, which result in a shorter net reverberation lag in the obscured case. The occultations also change the energy spectrum from a simple NLS1, with a strong soft excess and a steep power law above $2\,\mathrm{keV}$, to something resembling a complex NLS1, with a prominent iron line from reflection off the soft excess, highly ionised Fe Kα absorption lines and strong spectral curvature.

The short $\sim 50\,\mathrm{s}$ reverberation lags have been taken as evidence for extreme relativistic reflection in complex NLS1s. If $50\,\mathrm{s}$ really is a light travel time this requires reflection from the innermost radii of an accretion disc around a highly spinning black hole. Our occultation model shows that this need not be the case. A short reverberation lag can be the result of a much longer light travel time, diluted and shifted by the soft leads introduced by occulting clouds. In our model, reflection and reprocessing occurs between $6 - 12\,R_g$ and puts no constraints on the spin of the black hole. Changing inclination then naturally explains the change from smooth spectra to complex spectra, and long, lower frequency reverberation lags to shorter, higher frequency reverberation lags in simple and complex NLS1s. Given that NLS1s are high accretion rate sources, it is quite natural to expect that the disc is not flat, that radiation pressure can lift material from the disc which will obscure the central emission for high inclination lines of sight. This is a promising geometry to explore further.

Chapter 9

Concluding Remarks

In this thesis I have presented five papers on accretion and ejection around astrophysical black holes. In Chapter 4 I showed that the truncated disc model, developed to explain the dramatic hard to soft spectral state transition in stellar mass black hole binaries (BHBs), continues to well describe the behaviour of low/hard state (LHS) BHBs as they drop down into quiescence. The switch in seed photon source for Comptonisation, from the truncating disc to cyclo-synchrotron radiation generated within the hot flow that replaces it, provides a natural explanation for the hardening and then softening of the X-ray spectral index with decreasing luminosity. By coupling a synchrotron jet to the truncated disc model I showed that a transition from flow dominated X-ray emission to jet dominated X-ray emission causes a break in the radio–X-ray correlation. Since the observed radio–X-ray correlation remains unbroken down to low luminosities, jet emission taking over from accretion flow emission cannot be the cause of the X-ray spectral softening at low accretion rates. The X-ray emission of LHS BHBs therefore remains dominated by emission from the accretion flow down into quiescence, while the low bulk Lorentz factor (BLF) jet provides the low energy radio emission.

By contrast, the X-ray spectra from blazars are dominated by jet emission. These are Active Galactic Nuclei (AGN) with high BLF jets, where the jet is orientated directly towards us so that its emission is maximally Doppler boosted and dominates the spectrum of the AGN from radio up to gamma-ray energies. BL Lacs are low accretion rate blazars, in the same radiatively inefficient accretion regime as LHS

BHBs, while flat spectrum radio quasars (FSRQs) represent the high accretion rate end of the blazar sequence, corresponding to high/soft state BHBs with bright, radiatively efficient accretion discs. In Chapters 5 and 6 I combined spectral models of blazar jet emission with prescriptions for how they should scale with black hole mass and accretion rate and the number densities of black holes from cosmological simulations to predict the observed populations of BL Lacs and FSRQs, given certain assumptions for which black holes should produce highly relativistic blazar jets. I showed that assuming every low accretion rate black hole ($\dot{m} < 0.01$) produces a BL Lac jet and every high accretion rate black hole ($\dot{m} > 0.01$) produces a FSRQ jet vastly overpredicts the observed numbers of BL Lacs and FSRQs. I showed that assuming only high spin black holes ($a > 0.8$) produce BL Lac jets gives a much better match to the observed number of BL Lacs and their mass, accretion rate and redshift distributions, providing supermassive black hole accretion is chaotic. High spin black holes are then only produced by late, gas poor mergers and so correspond to the most massive black holes in the local Universe ($z < 2$). Matching the observed number of FSRQs likewise requires some limiting condition for FSRQ jet production, presumably also high spin. However chaotic accretion models for black hole spin evolution struggle to produce enough high spin black holes at $z = 2 - 3$ (where quasar activity and the FSRQ population peaks), since the rapid accretion at these times spins down any black holes that were previously spun up by the first generation of galaxy and black hole mergers at early times. This suggests accretion may have been more ordered at early times, allowing some black holes to maintain high spins despite high accretion rates. If FSRQs are due to high spin black holes then their redshift distribution provides powerful constraints for future, more sophisticated cosmological simulations of black hole spin evolution.

Fanaroff-Riley Type I and Type II (FRI and FRII) radio galaxies are the misaligned counterparts of BL Lacs and FSRQs, respectively. If BL Lacs and FSRQs host highly spinning black holes then so should FRI and FRII AGN. Broderick & Fender (2011) showed that FRIs, FRIIs and broad line radio galaxies (BLRGs; the lower mass, lower accretion rate counterparts of FRIIs) are all slightly offset from the fundamental plane that relates the X-ray and radio luminosity of both stellar

mass and supermassive black holes; specifically these three classes of AGN are more radio loud than the radio quiet quasars (RQQs), Seyferts, Low Ionisation Emission Region galaxies (LINERs) and BHBs which do follow the fundamental plane. If FRIs/BL Lacs, FRIIs/FSRQs and BLRGs are radio louder because their jets can tap the rotational spin energy of the black hole in addition to energy from the accretion flow, then this implies the RQQs, Seyferts, LINERs and BHBs should host low spin black holes. Directly measuring the black hole spins of AGN is both difficult and controversial. Narrow Line Seyfert 1s (NLS1s) provide the best opportunity as, unlike other AGN, their small masses and high accretion rates mean the disc emission in these objects can extend into the soft X-ray regime rather than peaking in the unobservable UV.

The X-ray spectra of NLS1s are often fitted with reflection dominated models, where the strong soft X-ray excess in these objects is fitted by an extremely relativistically smeared reflection component. Producing sufficient relativistic smearing requires the hard coronal X-rays are reflected off the innermost radii of a flat accretion disc extending down to $\sim 1 \, R_g$, implying the black hole must be highly spinning. The alternative to these reflection dominated models is the separate soft excess model, where the soft excess emission is produced by optically thick Comptonisation in the inner regions of the disc, which do not completely thermalise perhaps due to a larger scale-height caused by the turbulent inner regions being close to Eddington in these sources. This soft excess component is then a combination of continuum emission and reprocessed emission due to illumination by the coronal power law and there is only moderate reflection in these models to explain the iron line at $\sim 6.7 \, \mathrm{keV}$. In Chapter 7 I built a fully self-consistent spectral-timing model to investigate which of these two spectral models best matches the timing properties of the NLS1 PG1244+026 and found that only the separate soft excess model could reproduce all the observed timing properties of the source — power spectra, coherence, covariance, lag-frequency and lag-energy spectra. In contrast, reflection dominated models are too coherent over a wide energy and frequency range due to relativistic reflection contributing strongly at both hard and soft energies and responding to coronal fluctuations up to high frequencies due to the small reflection

radii involved. In addition, the observed soft leads at low frequencies require that at least some of the soft excess must be produced by a continuum component — it cannot be solely due to a reverberating component — ruling out the simplest reflection dominated models where this is the case. Since the disc is truncated by the optically thick Comptonised region in the separate soft excess model, fitting this model puts no constraints on the spin of the black hole, so is consistent with radio quiet NLS1s having low spin.

However, PG1244+026 is a simple NLS1 with only moderate variability. Complex NLS1s, such as 1H0707–495, MCG 6-30-15, MRK 766 and IRAS 13224–3809, have much more extreme variability with deep dip spectra showing even stronger soft excess and iron line features and shorter, higher frequency reverberation lags. These features are often taken as evidence for an extreme reflection dominated scenario, where the dim spectra appear more reflection dominated as the height of the corona drops and more coronal flux is gravitationally bent back towards the disc, and the short $\sim 30\,\mathrm{s}$ reverberation lags are taken as the light travel time from corona to disc. In Chapter 8 I showed that complex NLS1s similarly need not be explained by reflection dominated models. By coupling a simple occultation model to the previous separate soft excess spectral-timing model for simple NLS1s, I showed that the switch in spectral-timing properties from a simple NLS1 to a complex NLS1 can be understood in terms of changing inclination with respect to a clumpy turbulent region in the inner disc. For low inclination sources there is a clean line of sight down to the inner regions of the accretion flow and the source appears as a simple NLS1, while for high inclinations our line of sight intercepts the turbulent failed wind structure (which I identify with the soft excess) resulting in variable absorption and enhancing the variability of complex NLS1s. Occultations of the central regions introduce additional hard lags when Doppler boosting of the underlying accretion flow emission is taken into account. These dilute the soft reverberation lag and shift it to higher frequencies, so giving a smooth transition from the longer lags in simple NLS1s to the shorter lags of complex NLS1s, without requiring the black holes of either to be highly spinning.

If the bulk of NLS1s, which are radio quiet, host low spin black holes, then this

fits in with the theory that the subset of radio and gamma-ray loud NLS1s may host high spin black holes, like their higher mass FSRQ counterparts. If there really is a one to one correlation between highly spinning black holes and production of the most relativistic ($\Gamma \sim 15$) jets, this provides enormous potential for the inclusion of feedback in cosmological simulations. High spin black holes with $\Gamma \sim 15$ jets would feed significantly more power back into their surroundings over much larger scales than lower spin black holes limited to feeding back via slower $\Gamma \sim 2$ jets and winds, where the strength and type of wind will depend on the mass and accretion rate of the black hole. Combining jet and accretion flow spectral models with new simulations and observational constraints provides the way forward for testing such theories, using population statistics to continually refine the predicted spin evolution and feedback prescriptions.

Scaling jets with mass and accretion rate is a useful tool to compare the jet emission from different objects and test the scaling assumptions themselves. BL Lacs and LHS BHBs are in the same accretion regime. Scaling the blazar jet model down to stellar mass black holes provides a way to test whether BHBs could host $\Gamma \sim 15$ spines to their jets that are just not visible due to Doppler deboosting, or whether they really do lack a highly relativistic spine to their jet, as their position on the fundamental place, and presumably moderate spin, suggests.

I have shown that the soft excess of NLS1s is well modelled by optically thick Comptonisation and reprocessing in a turbulent region in the inner disc, where radiation pressure lifts material to a large scale-height, but where does that leave the soft excess in broad line Seyferts and quasars? These objects are at lower accretion rates so a failed Eddington wind cannot be the cause. The spectra of these objects are UV weak and X-ray stronger. Instead the large scale-height may be due to UV radiation lifting material, the luminous hard X-rays over-ionising it and it falling back down. Full radiative transfer models are needed to show whether this process can transform the disc into an optically thick Comptonising region in these objects.

Importantly, any model of the accretion flow must be able to replicate both the spectral and timing properties of a source. Including timing information provides a

powerful extra diagnostic for distinguishing between spectrally degenerate models, with power spectra, coherence, covariance, lag-frequency and lag-energy spectra all providing model constraints. However caution must be exercised when interpreting spectral lags. If one energy band contains several spectral components, each of which fluctuate in normalisation with time, then the net lag that is measured will be the sum of the lags of each component with respect to the reference band, weighted by the contribution of each component to the total flux in the band. Lags due to propagation of mass accretion rate fluctuations tend to dominate at low frequencies, with the hard emission progressively lagging the soft. Reverberation lags due to reflection and reprocessing, with soft emission and line emission lagging the hard illumination, tend to dominate at high frequencies. However these are not the only sources of lags. A single power law component on its own can produce lags if it pivots as a function of time due to fluctuations in its seed photons. These lags then occur on the timescale of the seed photon fluctuations. In AGN, the highest frequency power law fluctuations are approaching the light crossing time of the corona, in which case lags between individual Compton scattering orders may even be affecting the lag-energy spectrum. Spectral components which change shape as a function of time, such as reflection spectra responding to changes in the ionisation state of the reflector, can further complicate the lags, as the reflector ionisation depends not only on the instantaneous illuminating flux but also the flux history of the source (Silva & Uttley, in prep.). Finally, I have shown that occultations by clouds orbiting at small radii can affect the lag-frequency spectrum by adding soft leads on the transit timescale when Doppler boosting of the underlying accretion flow emission is taken into account. Combining spectral and timing information is clearly a powerful tool, but also a complex area. Understanding these effects is key to making the most out of these techniques, and better understanding accretion and ejection around both stellar mass and supermassive black holes.

Appendix A

A.1 Black Hole Binary Accretion Flow Model

The model consists of an outer black body (BB) disc, truncated at some radius (R_t), with an inner hot flow of radius R_{hot}, where $R_{hot} = 20\,R_g$. The hot flow is taken to be radiatively inefficient, such that:

$$L_{hot} = L_{BBdisc}(R < R_t) \left(\frac{\dot{m}}{\dot{m}_c} \right) \tag{1.1.1}$$

where $\dot{m} = \dot{M}/\dot{M}_E$, $\dot{m}_c = 0.1$ and $L_{BBdisc}(R < R_t)$ is the luminosity of a BB disc extending from the truncation radius down to the last stable orbit.

We scale the truncation radius with accretion rate, such that:

$$R_t = 20\,R_g \left(\frac{\dot{m}}{\dot{m}_c} \right)^{-1/2} \tag{1.1.2}$$

The optical depth (τ) of the Comptonising region is fixed at 2 for $\dot{m} = \dot{m}_c$, and scales with $\dot{m}$ as:

$$\tau = 2 \left(\frac{\dot{m}}{\dot{m}_c} \right) \tag{1.1.3}$$

The unabsorbed cyclo-synchrotron emission from the hot flow is calculated following Di Matteo et al. (1997):

$$L_{cyclo}(\nu) = 5.57 \times 10^{-29} \frac{n\nu I(x)V}{K_2(1/\theta_e)} \tag{1.1.4}$$

where $V = 2/3\pi R_{hot}^3$ is the volume of the Comptonising hot flow, n is the number

density of electrons calculated from the optical depth, $\theta = kT/m_e c^2$ is the dimensionless electron temperature, $K_2(1/\theta)$ is the modified Bessel function, $x = 2\nu/3\nu_B\theta^2$, $\nu_B = eB/2\pi m_e c$ is the Larmor frequency, and the function $I(x)$ is given by:

$$I(x) = \frac{4.050}{x^{1/6}} \left(1 + \frac{0.40}{x^{1/4}} + \frac{0.532}{x^{1/2}}\right) \exp(-1.8899x^{1/3}) \tag{1.1.5}$$

The magnetic field (B) of the hot flow is calculated from the density by assuming the ions are at the virial temperature and the magnetic field is 10% of the gas pressure, giving:

$$B = \sqrt{0.1 n m_p c^2 \frac{8\pi}{r_{hot}}} \tag{1.1.6}$$

where $r_{hot} = R_{hot}/R_g$. The cyclo-synchrotron self-absorption frequency is given by:

$$\nu_{csa} = \frac{3}{2}\nu_B\theta^2 x_m \tag{1.1.7}$$

where x_m is found by solving for x when the cyclo-synchrotron and BB emission are set equal:

$$L_{cyclo}(\nu_{csa}) = 8\pi^2 m_e \nu_{csa}^2 \theta_e R_{hot}^2 \tag{1.1.8}$$

Below the self-absorption frequency the absorbed emission is calculated as:

$$L(\nu < \nu_{csa}) = \left(\frac{\nu}{\nu_{csa}}\right)^{5/2} L(\nu_{csa}) \tag{1.1.9}$$

Thermal Comptonisation is modelled using EQPAIR (Coppi, 1999), with seed photons from both the disc and cyclo-synchrotron emission, where the fraction of disc photons from a given radius (R) intercepted by the hot flow is given by:

$$\frac{L_{seed,disc}}{L_{disc}} = \left(\frac{R_{hot}}{R}\right) \frac{\arcsin(R_{hot}/R)}{\pi} \tag{1.1.10}$$

And we calculate the mean seed photon temperature:

$$kT_{seed} = \frac{k(L_{seed,disc}T_{disc} + L_{cyclo}T_{cyclo})}{L_{seed,disc} + L_{cyclo}} \tag{1.1.11}$$

The electron temperature, a parameter in both the cyclo-synchrotron equations and EQPAIR, is calculated self consistently.

A.1.1   Jet Model

We construct a conical jet, where opening angle ($\phi = 0.1$) relates jet radius (R_j) to distance along the jet (z):

$$R_j(z) = \phi Z = \phi z R_g \qquad (1.1.12)$$

We assume a fraction of the accreting material (f_j) is diverted up the jet. The energy density in relativistic particles at the jet base is set to be some fraction ($f_{rel} = 0.1$) of the magnetic energy density:

$$m_e c^2 \int_{\gamma_{min}}^{\gamma_{max}} \gamma N(\gamma) d\gamma = U_{rel,0} = f_{rel} U_{B,0} \qquad (1.1.13)$$

We conserve magnetic energy and particle number along the jet such that:

$$B(z) = B_0 \left(\frac{z}{z_0} \right)^{-1} \qquad (1.1.14)$$

$$K(z) = K_0 \left(\frac{z}{z_0} \right)^{-2} \qquad (1.1.15)$$

and allow B_0 and K_0 to scale with accretion rate as:

$$B_0(\dot{m}) = B_0(\dot{m}_c) \left(\frac{\dot{m}}{\dot{m}_c} \right)^{1/2} \qquad (1.1.16)$$

$$K_0(\dot{m}) = K_0(\dot{m}_c) \left(\frac{\dot{m}}{\dot{m}_c} \right) \qquad (1.1.17)$$

where $\dot{m}_c = 0.1$, and $B_0(\dot{m}_c)$, $K_0(\dot{m}_c)$ and z_0 are fixed by requiring the radio luminosity and the optically thick-optically thin synchrotron break match observations of GX 339-4 (Gandhi et al., 2011).

We assume electrons in the jet are continually accelerated into a power law distribution of the form:

$$N(\gamma) = K\gamma^{-p} \qquad (1.1.18)$$

where $p = 2.4$, for electron Lorentz factors ranging from $\gamma = 1.0 - 1 \times 10^5$.

We split the jet into conical sections and calculate the synchrotron emission from electrons in each section:

$$L_{sync}(\nu) = \frac{\sigma_T c}{8\pi\nu_B} U_B \gamma N(\gamma) V \delta^3 \qquad (1.1.19)$$

where V is the volume of the conic section, $\delta = 1/(\Gamma - \cos\psi\sqrt{\Gamma^2 - 1})$ is the boosting factor of the jet, ψ is the angle of the jet with respect to the observer, and the electron Lorentz factor and synchrotron photon frequency are related by $\gamma = \sqrt{3\nu/4\nu_B}$.

The synchrotron self-absorption frequency (ν_{ssa}) in each section is given by (Ghisellini et al., 1985):

$$\nu_{ssa} = \left(4.62 \times 10^{14} K B^{2.5} \frac{R_j}{0.7}\right)^{2/7} \qquad (1.1.20)$$

The frequency of the observed radiation is boosted by a factor $\nu_{obs} = \nu\delta$. We neglect synchrotron self-Comptonisation.

A.2 Blazar Jet Model

Table A1 lists the blazar jet model parameters. The first four set the parameters of the black hole and the distance. Parameters 5–8 set the physical parameters of the jet: inclination to the line of sight, BLF, jet opening angle and distance of the emission region from the black hole. When combined together the last two of these set the radius of the emission region, since the code assumes a conical jet. Parameters 9 and 10 set the jet magnetic field and the power injected into relativistic electrons. The remaining parameters 11–15 determine the shape of the injected electron distribution, and parameter 16 sets the redshift.

The code can be used to model both FSRQs and BL Lacs. If $\log\dot{m} < -2$ (parameter 3) the code assumes the accretion flow regime is radiatively inefficient, corresponding to a BL Lac. In this case the external seed photon energy density is

	Parameter	Description
1	M	Black hole mass in solar masses
2	R_{co}	Comoving distance in Mpc
3	$\log \dot{m}$	Mass accretion rate in units of L/L_{Edd}
		(if $\log \dot{m} < -2$, does SSC with no external seed photons)
4	$a_{\star}$	Dimensionless black hole spin
5	θ_{obs}	Inclination of jet axis from line of sight in degrees
6	Γ	Jet bulk Lorentz factor
7	ϕ	Jet opening angle in radians
8	z_{diss}	Distance of dissipation region from the black hole in Rg
		(radius of dissipation region, $R_{diss} = \phi Z_{diss} = \phi z_{diss} R_g$)
9	B	Magnetic field in Gauss
10	$\log P_{rel}$	Power injected into relativistic electrons in erg s^{-1}
11	$\gamma_{min,inj}$	Minimum Lorentz factor of injected electron distribution
12	γ_b	Lorentz factor of break in injected electron distribution
13	γ_{max}	Maximum Lorentz factor of injected electron distribution
14	n_1	Index of electron distribution below the break
15	n_2	Index above the break
16	z	Redshift

Table A.1: Summary of blazar jet model parameters.

set to zero and the model calculates only synchrotron and synchrotron self-Compton (SSC) emission. If $\log \dot{m} \geqslant -2$ then the model assumes a radiatively efficient accretion disc is present and includes external Compton (EC) emission from external sources of seed photons. The radiatively efficient disc is assumed to illuminate the broad line region (BLR) and torus, both of which reprocess a fraction of the disc emission. The code calculates the energy density of seed photons from direct disc and coronal emission, BLR emission, reflection of coronal X-rays off the BLR and emission from the torus, following the method of Ghisellini & Tavecchio (2009).

The code prints to screen which type of jet is calculated (SSC or SSC+EC), along with the logarithm of the power in radiation, magnetic fields, electrons, protons and total jet power (P_r, P_B, P_e, P_p and P_j). For SSC+EC jets, the code flags if $Z_{diss} > R_{BLR}$ and $Z_{diss} > R_{IR}$. It also writes out into files the steady state electron distribution ($\gamma N(\gamma)$) and the energy density of seed photons in the jet frame ($U'(\nu')$).

Since the inclination of the jet is a parameter, the code can in practice be used to model any jet, not just highly aligned blazars. However the code assumes a single emission zone, so it is best suited for modelling the high energy jet base emission.

Although the FORTRAN subroutine can be easily modified to be called multiple times with increasing emission region size to model more extended structures.

A.2.1   Jet Emission Calculation

The emission comes from a single spherical zone of radius R_{diss}. We assume the jet has a constant opening angle (ϕ), such that the distance of the emission region from the central black hole ($Z_{diss} = z_{diss} R_g$) is related to the radius of the emission region by: $R_{diss} = \phi Z_{diss}$. We assume material in the jet moves at a constant bulk Lorentz factor (Γ) and that some fraction of the transported electrons are accelerated into a power law distribution between minimum and maximum Lorentz factors $\gamma_{min,inj}$ and γ_{max}, of the form:

$$Q(\gamma) = Q_0 \left(\frac{\gamma}{\gamma_b}\right)^{-n_1} \left(1 + \frac{\gamma}{\gamma_b}\right)^{n1-n2} = Q_0 q(\gamma) \text{ for } \gamma_{min,inj} < \gamma < \gamma_{max} \quad (1.2.21)$$

γ_b is the Lorentz factor at which the electron distribution changes in slope from n_1 to n_2. We calculate the normalisation Q_0 from the power injected into the accelerated electrons (P_{rel}):

$$P_{rel} = \frac{4\pi}{3} R_{diss}^3 m_e c^2 Q_0 \int_{\gamma_{min,inj}}^{\gamma_{max}} \gamma q(\gamma) d\gamma \quad (1.2.22)$$

We calculate γ_{cool} after a light crossing time $t_{cross} = R_{diss}/c = \gamma_{cool}/\dot{\gamma}_{cool}$, as:

$$\gamma_{cool} = \frac{3m_e c^2}{4\sigma_T R U_{seed}} \quad (1.2.23)$$

where $U_{seed} = U_B + U_{sync} + U_{ex}$ is the sum of the energy density in magnetic fields, synchrotron emission and external emission which provides the seed photons for cooling.

We solve the continuity equation to find the self consistent steady state electron distribution:

$$N(\gamma, t_{cross}) = Kn(\gamma)$$

$$
= \begin{cases}
AQ_0 q(\gamma) & \text{for } \gamma_{min,inj} < \gamma < \gamma_{cool} \\
\frac{3m_e c^2}{4\sigma_T c U_{seed}} \frac{Q_0}{\gamma^2} \int_\gamma^{\gamma_{max}} q(\gamma) d\gamma & \text{for } \gamma_{cool} < \gamma < \gamma_{max}
\end{cases}
\tag{1.2.24}
$$

where A is found by matching at γ_{cool}.

We use the delta function approximation and calculate the synchrotron emissivity as:

$$j_{sync}(\nu) = \frac{\sigma_T c}{6\pi\nu_B} U_B \gamma N(\gamma) \tag{1.2.25}$$

where the electron Lorentz factor and synchrotron photon frequency are related by $\gamma = \sqrt{3\nu/4\nu_B}$ and we calculate the synchrotron self-absorption frequency (ν_{ssa}) as given by (Ghisellini et al., 1985):

$$\nu_{ssa} = \left(4.62 \times 10^{14} K B^{2.5} \frac{R_{diss}}{0.7}\right)^{2/7} \tag{1.2.26}$$

We calculate Compton emission including the Klein-Nishina cross section using the delta approximation:

$$j_{comp}(\nu) = \frac{\sigma_T c}{6\pi} \int_{\gamma_{min}}^{\gamma_{max}} \int_{\nu_{seed,min}}^{\nu_{seed,max}} \frac{U_{seed}(\nu_{seed})}{\nu_{seed}} \gamma N(\gamma) d\nu_{seed} d\gamma \tag{1.2.27}$$

where electron Lorentz factor and Compton photon frequency are related by $\gamma = \sqrt{3\nu/4\nu_{seed}}$ and $\gamma_{min} = \gamma_{min,inj}$, unless $\gamma_{cool} < \gamma_{min,inj}$ and then $\gamma_{min} = \gamma_{cool}$.

Bulk motion of the jet boosts and blue shifts the emission. We calculate the observed flux as:

$$F(\nu\delta/(1+z)) = \frac{(j_{sync}(\nu) + j_{comp}(\nu))}{R_{co}^2} \frac{4\pi}{3} R_{diss}^3 \delta^3 \tag{1.2.28}$$

where $\delta = (\Gamma - \cos\theta\sqrt{\Gamma^2 - 1})^{-1}$ is the Doppler factor and R_{co} is the comoving distance to the object at redshift z.

We neglect photon-photon pair production. However the code calculates the source compactness and flags a warning if $l' \geqslant 30$. This corresponds to $\tau_{\gamma\gamma} \sim 1$, i.e.

when the source starts to become optically thick to photon-photon pair production and this effect becomes important. For most blazar jets the compactness is typically < 3.

A.2.2 External Seed Photons

If $\log \dot{m} \geqslant -2$ then the model assumes a radiatively efficient accretion disc is present and includes EC emission from external sources of seed photons, calculating the energy density of seed photons following the method of Ghisellini & Tavecchio (2009). The model includes direct disc and coronal emission, BLR emission, reflection of coronal X-rays off the BLR and emission from the torus.

The accretion disc luminosity (L_d) is calculated from M and $\dot{m}$ (parameters 1 and 3). Each annulus of the disc is seen at a different angle with respect to the jet emission region so receives a difference amount of Doppler deboosting (b_d). We approximate the energy density of disc seed photons from each annulus in the jet frame as:

$$U_d'(\nu') = \frac{4\pi h b_d}{c^3} \frac{(\nu/b_d)^3}{\exp\left[\frac{h\nu/b_d}{kT}\right] - 1} d\mu_d \qquad (1.2.29)$$

where $b_d = \Gamma(1 - \beta\mu_d)$, $\mu_d = \cos\eta$ and η is the angle of the annulus with respect to the jet axis. μ_d therefore varies between $\mu_{max} = 1$, for the innermost radii which are directly behind the jet and experience most deboosting, to $\mu_{min} = [1 + (R_{d,max}/Z_{diss})^2]^{-1/2}$ for the outermost radius $R_{d,max} = 1000\,R_g$. We calculate the temperature of each disc radius from the mass, accretion rate and spin input in parameters 1, 3 and 4.

We assume the luminosity of coronal X-rays is $L_X = f_X L_d = 0.1 L_d$ and the corona extends to $R_X = 60\,R_g$. Its emission is therefore deboosted by a factor $b_X = \Gamma(1 - \beta\mu_X)$, where $\mu_X = [1 + (R_X/Z_{diss})^2]^{-1/2}$. The total energy density of coronal seed photons in the jet frame is therefore:

$$U_X' = \frac{f_X L_d \Gamma^2}{\pi R_X^2 c} \left[1 - \mu_X - \beta(1 - \mu_X^2) + \frac{\beta^2}{3}(1 - \mu_X^3)\right] \qquad (1.2.30)$$

We assume the spectrum of this emission is a cut-off power law starting from $b_X \nu_{d,peak}$, where $\nu_{d,peak} = 4kT_{max}/h$ is the frequency at which the unboosted disc spectrum peaks. We assume the power law cut-off $v_c = 150 \times 10^3 e/h$, so that:

$$U'_X(\nu') \propto \nu'^{-\alpha_X} \exp\left[-\frac{\nu'}{b_X \nu_c}\right] \tag{1.2.31}$$

where $\alpha = 1$.

We assume a fraction $f_{BLR} = 0.1$ of the disc luminosity is reprocessed by the BLR. This emission takes the form of a BB centred on the frequency of the Lyman α line ($\nu_{Ly\alpha} = \frac{c}{4(1216 \times 10^{-8})}$ Hz), so that:

$$U'_{BLR}(\nu') \propto \frac{\nu'^3}{\exp\left[\frac{\nu'}{b_{BLR}\nu_{Ly\alpha}}\right] - 1} \tag{1.2.32}$$

The total energy density in the jet frame (U'_{BLR}) and boosting factor (b_{BLR}) depend on the radius of the BLR (R_{BLR}) compared to Z_{diss}. The radius of the BLR scales with L_d as:

$$R_{BLR} = 10^{17} \left(\frac{L_d}{10^{45} \text{ erg s}^{-1}}\right)^{1/2} \text{cm} \tag{1.2.33}$$

If $Z_{diss} < R_{BLR}$:

$$U'_{BLR} = \frac{17\Gamma^2}{12} \frac{f_{BLR}L_d}{4\pi c R_{BLR}^2} \tag{1.2.34}$$

$$b_{BLR} = \Gamma \tag{1.2.35}$$

If $Z_{diss} > 3R_{BLR}$:

$$U'_{BLR} = \frac{f_{BLR}L_d}{4\pi c R_{BLR}^2} \frac{\Gamma^2}{3\beta} \left[2(1-\beta\mu_1)^3 - (1-\beta\mu_2)^3 - (1-\beta)^3\right] \tag{1.2.36}$$

$$\mu_1 = \left[1 + (R_{BLR}/Z_{diss})^2\right]^{-1/2} \tag{1.2.37}$$

$$\mu_2 = \left[1 - (R_{BLR}/Z_{diss})^2\right]^{1/2} \tag{1.2.38}$$

$$b_{BLR} = \Gamma(1 - \beta\mu_1) \tag{1.2.39}$$

If $R_{BLR} \leqslant Z_{diss} \leqslant 3R_{BLR}$, we use a power law interpolation between the two regimes for U'_{BLR} and assume $b_{BLR} = \Gamma(1 - \beta\mu_1)$ for $Z_{diss} = 3R_{BLR}$.

We assume a fraction $f_{XBLR} = 0.01$ of the coronal X-rays are reflected by the BLR clouds. We assume the reflected emission has the same cut-off power law shape as the direct coronal emission. Both b_{XBLR} and U'_{XBLR} vary as b_{BLR} and U'_{BLR}, with $f_{BLR}L_d$ replaced by $f_{XBLR}f_X L_d$.

We assume a fraction $f_{IR} = 0.3$ of the disc luminosity is reprocessed by the torus. This emission takes the form of a BB at ~ 370 K (i.e. $\nu_{IR} = 370k/h$), so that:

$$U'_{IR}(\nu') \propto \frac{\nu'^3}{\exp\left[\frac{\nu'}{b_{IR}\nu_{IR}}\right] - 1} \tag{1.2.40}$$

As in the case of the BLR seed photons, U'_{IR} and b_{IR} depend on the radius of the torus (R_{IR}) compared to Z_{diss}. R_{IR} scales with L_d as:

$$R_{IR} = 2.5 \times 10^{18} \left(\frac{L_d}{10^{45}\,\text{erg s}^{-1}}\right)^{1/2} \text{cm} \tag{1.2.41}$$

and again we consider three regimes. If $Z_{diss} < R_{IR}$:

$$U'_{IR} = \frac{f_{IR}L_d\Gamma^2}{4\pi cR_{IR}^2} \tag{1.2.42}$$

If $Z_{diss} > 3R_{IR}$:

$$U'_{IR} = \frac{f_{IR}L_d}{4\pi cR_{IR}^2}\frac{\Gamma^2}{3\beta}\left[2(1 - \beta\mu_1)^3 - (1 - \beta\mu_2)^3 - (1 - \beta)^3\right] \tag{1.2.43}$$

$$\mu_1 = \left[1 + (R_{IR}/Z_{diss})^2\right]^{-1/2} \tag{1.2.44}$$

$$\mu_2 = \left[1 - (R_{IR}/Z_{diss})^2\right]^{1/2} \tag{1.2.45}$$

If $R_{IR} \leqslant Z_{diss} \leqslant 3R_{IR}$, we use a power law interpolation between the two

regimes for U'_{IR}. In all three cases we use $b_{IR} = \Gamma(1 - \beta\mu_{IR})$, where $\mu_{IR} = \cos(\arctan(R_{IR}/Z_{diss}))$.

A.3   Time Series Analysis Techniques

For two light curves, $s(t)$ and $h(t)$, with Fourier transforms $S(f)$ and $H(f)$, we define the complex valued cross spectrum as (Vaughan & Nowak, 1997; Nowak et al., 1999):

$$C(f) = S^*(f)H(f) \tag{1.3.46}$$

where S^* denotes the complex conjugate of S.

The coherence is then:

$$\mathrm{coh}(f) = \frac{|\langle C(f)\rangle|^2}{\langle|S(f)|^2\rangle\langle|H(f)|^2\rangle} \tag{1.3.47}$$

where angle brackets denote an average over multiple light curve segments.

The lag is calculated as:

$$\mathrm{lag}(f) = \frac{\arg[C(f)]}{2\pi f} \tag{1.3.48}$$

The covariance between the two lightcurves, if $h(t)$ is the reference band, is calculated in the time domain as (Wilkinson & Uttley, 2009):

$$\mathrm{cov} = \frac{\sum(s(t) - \bar{s})(h(t) - \bar{h})}{(N - 1)\sqrt{\sigma_h^2}} \tag{1.3.49}$$

where N is the number of time steps in the light curve, $\bar{s}$ and $\bar{h}$ are the mean fluxes of $s(t)$ and $h(t)$ respectively, and the factor of $\sqrt{\sigma_h^2}$ is required for normalisation. σ_h^2 is the excess variance of the reference band light curve, defined as:

$$\sigma_h^2 = \frac{\sum(h(t) - \bar{h})^2}{(N - 1)} \tag{1.3.50}$$

We normalise all power spectra $(P(f) = |H(f)|^2)$, such that:

$$\int_0^\infty P(f)df = \frac{\sigma^2}{I^2} \tag{1.3.51}$$

where σ^2 is the excess variance of the light curve and I is its mean flux.

Bibliography

Abdo A. A., Ackermann M., Ajello M., Allafort A., Antolini E., Atwood W. B., Axelsson M., Baldini L., Ballet J., Barbiellini G., et al. 2010, ApJ, 715, 429

Abdo A. A., Ackermann M., Ajello M., Axelsson M., Baldini L., Ballet J., Barbiellini G., Bastieri D., Baughman B. M., Bechtol K., et al. 2009, ApJ, 707, 727

Abdo A. A., Ackermann M., Ajello M., Baldini L., Ballet J., Barbiellini G., Bastieri D., Bechtol K., Bellazzini R., Berenji B., et al. 2009, ApJ, 707, L142

Alston W. N., Done C., Vaughan S., 2014, MNRAS, 439, 1548

Antonucci R., 1993, ARA&A, 31, 473

Antonucci R. R. J., Miller J. S., 1985, ApJ, 297, 621

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

Arévalo P., Uttley P., Kaspi S., Breedt E., Lira P., McHardy I. M., 2008, MNRAS, 389, 1479

Axelsson M., Borgonovo L., Larsson S., 2005, A&A, 438, 999

Balbus S. A., Hawley J. F., 1991, ApJ, 376, 214

Beckmann V., Shrader C. R., 2012, Active Galactic Nuclei

Begelman M. C., McKee C. F., Shields G. A., 1983, ApJ, 271, 70

Best P. N., Heckman T. M., 2012, MNRAS, 421, 1569

Bîrzan L., Rafferty D. A., McNamara B. R., Wise M. W., Nulsen P. E. J., 2004, ApJ, 607, 800

Blandford R. D., Königl A., 1979, ApJ, 232, 34

Blandford R. D., Znajek R. L., 1977, MNRAS, 179, 433

Boller T., Brandt W. N., Fink H., 1996, A&A, 305, 53

Boroson T. A., 2002, ApJ, 565, 78

Bower R. G., Benson A. J., Malbon R., Helly J. C., Frenk C. S., Baugh C. M., Cole
S., Lacey C. G., 2006, MNRAS, 370, 645

Bozza V., 2010, General Relativity and Gravitation, 42, 2269

Broderick J. W., Fender R. P., 2011, MNRAS, 417, 184

Bromm V., Larson R. B., 2004, ARA&A, 42, 79

Cackett E. M., Fabian A. C., Zogbhi A., Kara E., Reynolds C., Uttley P., 2013,
ApJ, 764, L9

Cappi M., 2006, Astronomische Nachrichten, 327, 1012

Castor J. I., Abbott D. C., Klein R. I., 1975, ApJ, 195, 157

Chadwick P. M., Lyons K., McComb T. J. L., Orford K. J., Osborne J. L., Rayner
S. M., Shaw S. E., Turver K. E., Wieczorek G. J., 1999, ApJ, 513, 161

Chainakun P., Young A. J., 2012, MNRAS, 420, 1145

Chandrasekhar S., 1931, ApJ, 74, 81

Chevallier L., Collin S., Dumont A.-M., Czerny B., Mouchet M., Gonçalves A. C.,
Goosmann R., 2006, A&A, 449, 493

Chiaberge M., Celotti A., Capetti A., Ghisellini G., 2000, A&A, 358, 104

Chiang C.-Y., Fabian A. C., 2011, MNRAS, 414, 2345

Chitre D. M., Hartle J. B., 1976, ApJ, 207, 592

Churazov E., Gilfanov M., Revnivtsev M., 2001, MNRAS, 321, 759

Collin S., Kawaguchi T., Peterson B. M., Vestergaard M., 2006, A&A, 456, 75

Collinson J. S., Ward M. J., Done C., Landt H., Elvis M., McDowell J. C., 2015, MNRAS, 449, 2174

Coppi P. S., 1999, in Poutanen J., Svensson R., eds, High Energy Processes in Accreting Black Holes Vol. 161 of Astronomical Society of the Pacific Conference Series, The Physics of Hybrid Thermal/Non-Thermal Plasmas. p. 375

Corbel S., Coriat M., Brocksopp C., Tzioumis A. K., Fender R. P., Tomsick J. A., Buxton M. M., Bailyn C. D., 2013, MNRAS, 428, 2500

Corbel S., Koerding E., Kaaret P., 2008, MNRAS, 389, 1697

Corbel S., Nowak M. A., Fender R. P., Tzioumis A. K., Markoff S., 2003, A&A, 400, 1007

Corbel S., Tomsick J. A., Kaaret P., 2006, ApJ, 636, 971

Corbett E. A., Croom S. M., Boyle B. J., Netzer H., Miller L., Outram P. J., Shanks T., Smith R. J., Rhook K., 2003, MNRAS, 343, 705

Crummy J., Fabian A. C., Gallo L., Ross R. R., 2006, MNRAS, 365, 1067

Czerny B., Hryniewicz K., 2011, A&A, 525, L8

Czerny B., Modzelewska J., Petrogalli F., Pych W., Adhikari T. P., Życki P. T., Hryniewicz K., Krupa M., Świętoń A., Nikołajuk M., 2015, Advances in Space Research, 55, 1806

Czerny B., Nikołajuk M., Różańska A., Dumont A.-M., Loska Z., Zycki P. T., 2003, A&A, 412, 317

Czerny B., Rózańska A., Kuraszkiewicz J., 2004, A&A, 428, 39

Das S., Becker P. A., Le T., 2009, ApJ, 702, 649

De Marco B., Ponti G., Cappi M., Dadina M., Uttley P., Cackett E. M., Fabian A. C., Miniutti G., 2013, MNRAS, 431, 2441

Dermer C. D., Menon G., 2009, High Energy Radiation from Black Holes: Gamma Rays, Cosmic Rays, and Neutrinos

Dermer C. D., Schlickeiser R., Mastichiadis A., 1992, A&A, 256, L27

Di Matteo T., Celotti A., Fabian A. C., 1997, MNRAS, 291, 805

Ding J., Yuan F., Liang E., 2010, ApJ, 708, 1545

Done C., 2010, ArXiv e-prints

Done C., Davis S. W., Jin C., Blaes O., Ward M., 2012, MNRAS, 420, 1848

Done C., Gierliński M., 2005, MNRAS, 364, 208

Done C., Gierliński M., Kubota A., 2007, A&A Rev., 15, 1

Done C., Jin C., Middleton M., Ward M., 2013, MNRAS, 434, 1955

Done C., Nayakshin S., 2007, MNRAS, 377, L59

Donoso E., Li C., Kauffmann G., Best P. N., Heckman T. M., 2010, MNRAS, 407, 1078

Elvis M., 2000, ApJ, 545, 63

Elvis M., Wilkes B. J., McDowell J. C., Green R. F., Bechtold J., Willner S. P., Oey M. S., Polomski E., Cutri R., 1994, ApJS, 95, 1

Emmanoulopoulos D., McHardy I. M., Papadakis I. E., 2011, MNRAS, 416, L94

Esin A. A., McClintock J. E., Narayan R., 1997, ApJ, 489, 865

Fabian A. C., 2012, ARA&A, 50, 455

Fabian A. C., Ballantyne D. R., Merloni A., Vaughan S., Iwasawa K., Boller T., 2002, MNRAS, 331, L35

Fabian A. C., Iwasawa K., Reynolds C. S., Young A. J., 2000, PASP, 112, 1145

Fabian A. C., Miniutti G., Gallo L., Boller T., Tanaka Y., Vaughan S., Ross R. R., 2004, MNRAS, 353, 1071

Fabian A. C., Rees M. J., Stella L., White N. E., 1989, MNRAS, 238, 729

Fabian A. C., Zoghbi A., Ross R. R., Uttley P., Gallo L. C., Brandt W. N., Blustin A. J., Boller T., Caballero-Garcia M. D., Larsson J., Miller J. M., Miniutti G., Ponti G., Reis R. C., Reynolds C. S., Tanaka Y., Young A. J., 2009, Nature, 459, 540

Falcke H., Körding E., Markoff S., 2004, A&A, 414, 895

Fanaroff B. L., Riley J. M., 1974, MNRAS, 167, 31P

Fanidakis N., Baugh C. M., Benson A. J., Bower R. G., Cole S., Done C., Frenk C. S., 2011, MNRAS, 410, 53

Fanidakis N., Baugh C. M., Benson A. J., Bower R. G., Cole S., Done C., Frenk C. S., Hickox R. C., Lacey C., Del P. Lagos C., 2012, MNRAS, 419, 2797

Fender R. P., Belloni T. M., Gallo E., 2004, MNRAS, 355, 1105

Fender R. P., Gallo E., Jonker P. G., 2003, MNRAS, 343, L99

Ferrarese L., Merritt D., 2000, ApJ, 539, L9

Foschini L., Angelakis E., Fuhrmann L., Ghisellini G., Hovatta T., Lahteenmaki A., Lister M. L., Braito V., Gallo L., Hamilton T. S., et al. 2012, A&A, 548, A106

Foschini L., Berton M., Caccianiga A., Ciroi S., Cracco V., Peterson B. M., Angelakis E., Braito V., Fuhrmann L., Gallo L., et al. 2015, A&A, 575, A13

Fragile P. C., Blaes O. M., 2008, ApJ, 687, 757

Gallo E., 2010, in Belloni T., ed., Lecture Notes in Physics, Berlin Springer Verlag Vol. 794 of Lecture Notes in Physics, Berlin Springer Verlag, Radio Emission and Jets from Microquasars. p. 85

Gallo E., Fender R., Kaiser C., Russell D., Morganti R., Oosterloo T., Heinz S., 2005, Nature, 436, 819

Gallo E., Fender R. P., Miller-Jones J. C. A., Merloni A., Jonker P. G., Heinz S., Maccarone T. J., van der Klis M., 2006, MNRAS, 370, 1351

Gallo E., Fender R. P., Pooley G. G., 2003, MNRAS, 344, 60

Gallo L. C., 2006a, MNRAS, 368, 479

Gallo L. C., 2006b, MNRAS, 368, 479

Gallo L. C., Fabian A. C., Grupe D., Bonson K., Komossa S., Longinotti A. L., Miniutti G., Walton D. J., Zoghbi A., Mathur S., 2013, MNRAS, 428, 1191

Gandhi P., Blain A. W., Russell D. M., Casella P., Malzac J., Corbel S., D'Avanzo P., Lewis F. W., Markoff S., Cadolle Bel M., Goldoni P., Wachter S., Khangulyan D., Mainzer A., 2011, ApJ, 740, L13

Gandhi P., Makishima K., Durant M., Fabian A. C., Dhillon V. S., Marsh T. R., Miller J. M., Shahbaz T., Spruit H. C., 2008, MNRAS, 390, L29

García J., Dauser T., Reynolds C. S., Kallman T. R., McClintock J. E., Wilms J., Eikmann W., 2013, ApJ, 768, 146

Gardner E., Done C., 2014a, MNRAS, 442, 2456

Gardner E., Done C., 2014b, MNRAS, 438, 779

Gebhardt K., Bender R., Bower G., Dressler A., Faber S. M., Filippenko A. V., Green R., Grillmair C., Ho L. C., Kormendy J., Lauer T. R., Magorrian J., Pinkney J., Richstone D., Tremaine S., 2000, ApJ, 539, L13

George I. M., Fabian A. C., 1991, MNRAS, 249, 352

Ghisellini G., Celotti A., Fossati G., Maraschi L., Comastri A., 1998, MNRAS, 301, 451

Ghisellini G., Maraschi L., Treves A., 1985, A&A, 146, 204

Ghisellini G., Tavecchio F., 2008a, MNRAS, 386, L28

Ghisellini G., Tavecchio F., 2008b, MNRAS, 387, 1669

Ghisellini G., Tavecchio F., 2009, MNRAS, 397, 985

Ghisellini G., Tavecchio F., Foschini L., Ghirlanda G., 2011, MNRAS, 414, 2674

Ghisellini G., Tavecchio F., Foschini L., Ghirlanda G., Maraschi L., Celotti A., 2010, MNRAS, 402, 497

Gierliński M., Done C., 2004, MNRAS, 349, L7

Gierliński M., Done C., Page K., 2008, MNRAS, 388, 753

Gierliński M., Done C., Page K., 2009, MNRAS, 392, 1106

Gierliński M., Zdziarski A. A., Poutanen J., Coppi P. S., Ebisawa K., Johnson W. N., 1999, MNRAS, 309, 496

Gonzalez-Martin O., Masegosa J., Marquez I., Rodríguez-Espinosa J. M., Acosta-Pulido J. A., Ramos-Almeida C., Dultzin D., Hernandez-Garcia L., Ruschel-Dutra D., Alonso-Herrero A., 2015, ArXiv e-prints

Grove J. E., Johnson W. N., Kroeger R. A., McNaron-Brown K., Skibo J. G., Phlips B. F., 1998, ApJ, 500, 899

Haardt F., Maraschi L., 1991, ApJ, 380, L51

Hagino K., Odaka H., Done C., Gandhi P., Watanabe S., Sako M., Takahashi T., 2015, MNRAS, 446, 663

Hannikainen D. C., Hunstead R. W., Campbell-Wilson D., Sood R. K., 1998, A&A, 337, 460

Hardcastle M. J., 2005, Royal Society of London Philosophical Transactions Series A, 363, 2711

Hardcastle M. J., Krause M. G. H., 2013, MNRAS, 430, 174

Heckman T. M., 1980, A&A, 87, 152

Heinz S., 2004, MNRAS, 355, 835

Heinz S., Sunyaev R. A., 2003, MNRAS, 343, L59

Henisey K. B., Blaes O. M., Fragile P. C., 2012, ApJ, 761, 18

Higginbottom N., Proga D., Knigge C., Long K. S., Matthews J. H., Sim S. A., 2014, ApJ, 789, 19

Hoshino M., 2013, ApJ, 773, 118

Ibragimov A., Poutanen J., Gilfanov M., Zdziarski A. A., Shrader C. R., 2005, MNRAS, 362, 1435

Ingram A., Done C., 2011, MNRAS, 415, 2323

Ingram A., Done C., 2012a, MNRAS, 419, 2369

Ingram A., Done C., 2012b, MNRAS, 427, 934

Ingram A., Done C., Fragile P. C., 2009, MNRAS, 397, L101

Inoue H., Matsumoto C., 2003, PASJ, 55, 625

Jiang Y.-F., Stone J. M., Davis S. W., 2014, ApJ, 796, 106

Jin C., Done C., Middleton M., Ward M., 2013, MNRAS, 436, 3173

Jin C., Ward M., Done C., 2012, MNRAS, 425, 907

Jin C., Ward M., Done C., Gelbord J., 2012, MNRAS, 420, 1825

Kalogera V., Baym G., 1996, ApJ, 470, L61

Kanbach G., Straubmeier C., Spruit H. C., Belloni T., 2001, Nature, 414, 180

Kara E., Cackett E. M., Fabian A. C., Reynolds C., Uttley P., 2014, MNRAS, 439, L26

Kara E., Fabian A. C., Cackett E. M., Miniutti G., Uttley P., 2013, MNRAS, 430, 1408

Kara E., Fabian A. C., Cackett E. M., Steiner J. F., Uttley P., Wilkins D. R., Zoghbi A., 2013, MNRAS, 428, 2795

Kato S., Fukue J., Mineshige S., eds, 1998, Black-hole accretion disks

Kellermann K. I., Sramek R., Schmidt M., Shaffer D. B., Green R., 1989, AJ, 98, 1195

Khachikian E. Y., Weedman D. W., 1974, ApJ, 192, 581

King A. R., Pounds K. A., 2003, MNRAS, 345, 657

King A. R., Pringle J. E., Hofmann J. A., 2008, MNRAS, 385, 1621

Kneiske T. M., Dole H., 2010, A&A, 515, A19

Kolehmainen M., Done C., Díaz Trigo M., 2014, MNRAS, 437, 316

Kollatschny W., Zetzl M., 2013, A&A, 558, A26

Komossa S., Voges W., Xu D., Mathur S., Adorf H.-M., Lemson G., Duschl W. J., Grupe D., 2006, AJ, 132, 531

Kotov O., Churazov E., Gilfanov M., 2001, MNRAS, 327, 799

Krause M., Alexander P., Riley J., Hopton D., 2012, MNRAS, 427, 3196

Krolik J. H., McKee C. F., Tarter C. B., 1981, ApJ, 249, 422

Kubota A., Done C., 2004, MNRAS, 353, 980

Laor A., Netzer H., 1989, MNRAS, 238, 897

Leahy J. P., Perley R. A., 1991, AJ, 102, 537

Leighly K. M., 1999, ApJS, 125, 317

Leighly K. M., Moore J. R., 2004, ApJ, 611, 107

Liu B. F., Done C., Taam R. E., 2011, ApJ, 726, 10

Liu B. F., Meyer F., Meyer-Hofmeister E., 2006, A&A, 454, L9

Liu B. F., Mineshige S., Meyer F., Meyer-Hofmeister E., Kawaguchi T., 2002, ApJ, 575, 117

Lyubarskii Y. E., 1997, MNRAS, 292, 679

Maccarone T. J., Coppi P. S., 2003, A&A, 399, 1151

Madau P., Rees M. J., 2001, ApJ, 551, L27

Maiolino R., Salvati M., Bassani L., Dadina M., della Ceca R., Matt G., Risaliti G., Zamorani G., 1998, A&A, 338, 781

Malzac J., Belloni T., Spruit H. C., Kanbach G., 2004, Nuclear Physics B Proceedings Supplements, 132, 400

Malzac J., Belmont R., 2009, MNRAS, 392, 570

Malzac J., Belmont R., Fabian A. C., 2009, MNRAS, 400, 1512

Malzac J., Dumont A. M., Mouchet M., 2005, A&A, 430, 761

Maraschi L., Colpi M., Ghisellini G., Perego A., Tavecchio F., 2012, Journal of Physics Conference Series, 355, 012016

Marconi A., Axon D. J., Maiolino R., Nagao T., Pastorini G., Pietrini P., Robinson A., Torricelli G., 2008, ApJ, 678, 693

Markoff S., Nowak M. A., Wilms J., 2005, ApJ, 635, 1203

Matt G., Marinucci A., Guainazzi M., Brenneman L. W., Elvis M., Lohfink A., Arèvalo P., Boggs S. E., Cappi M., Christensen F. E., et al. 2014, MNRAS, 439, 3016

Mayer M., Pringle J. E., 2007, MNRAS, 376, 435

McConnell M. L., Zdziarski A. A., Bennett K., Bloemen H., Collmar W., Hermsen W., Kuiper L., Paciesas W., Phlips B. F., Poutanen J., Ryan J. M., Schönfelder V., Steinle H., Strong A. W., 2002, ApJ, 572, 984

McConnell N. J., Ma C.-P., 2013, ApJ, 764, 184

McHardy I. M., Arévalo P., Uttley P., Papadakis I. E., Summons D. P., Brinkmann W., Page M. J., 2007, MNRAS, 382, 985

McHardy I. M., Koerding E., Knigge C., Uttley P., Fender R. P., 2006, Nature, 444, 730

McHardy I. M., Papadakis I. E., Uttley P., Page M. J., Mason K. O., 2004, MNRAS, 348, 783

Mehdipour M., Branduardi-Raymont G., Kaastra J. S., Petrucci P. O., Kriss G. A., Ponti G., Blustin A. J., Paltani S., Cappi M., Detmers R. G., Steenbrugge K. C., 2011, A&A, 534, A39

Merloni A., Heinz S., di Matteo T., 2003, MNRAS, 345, 1057

Merritt D., Milosavljević M., 2005, Living Reviews in Relativity, 8, 8

Meyer F., Meyer-Hofmeister E., 1994, A&A, 288, 175

Middleton M. J., Ingram A. R., 2015, MNRAS, 446, 1312

Miller L., Turner T. J., 2013, ApJ, 773, L5

Miller L., Turner T. J., Reeves J. N., 2008, A&A, 483, 437

Miller L., Turner T. J., Reeves J. N., George I. M., Kraemer S. B., Wingert B., 2007, A&A, 463, 131

Miniutti G., Agís González B., Sanfrutos Carreras M., 2014, in The X-ray Universe 2014 The X-ray continuum and soft excess emitting region sizes from occultation events. p. 141

Miniutti G., Fabian A. C., 2004, MNRAS, 349, 1435

Mirabel I. F., Rodríguez L. F., 1994, Nature, 371, 46

Miyakawa T., Ebisawa K., Inoue H., 2012, PASJ, 64, 140

Miyamoto S., Kimura K., Kitamoto S., Dotani T., Ebisawa K., 1991, ApJ, 383, 784

Motta S., Belloni T., Homan J., 2009, MNRAS, 400, 1603

Murray N., Chiang J., 1996, Nature, 382, 789

Nandra K., O'Neill P. M., George I. M., Reeves J. N., 2007, MNRAS, 382, 194

Narayan R., Barret D., McClintock J. E., 1997, ApJ, 482, 448

Narayan R., McClintock J. E., 2012, MNRAS, 419, L69

Narayan R., Yi I., 1995, ApJ, 452, 710

Nayakshin S., Kallman T. R., 2001, ApJ, 546, 406

Netzer H., 2015, ArXiv e-prints

Netzer H., Laor A., 1993, ApJ, 404, L51

Nikołajuk M., Czerny B., Gurynowicz P., 2009, MNRAS, 394, 2141

Nolan P. L., Abdo A. A., Ackermann M., Ajello M., Allafort A., Antolini E., Atwood
W. B., Axelsson M., Baldini L., Ballet J., et al. 2012, ApJS, 199, 31

Novikov I. D., Thorne K. S., 1973, in Dewitt C., Dewitt B. S., eds, Black Holes (Les
Astres Occlus) Astrophysics of black holes.. pp 343–450

Nowak M. A., Vaughan B. A., Wilms J., Dove J. B., Begelman M. C., 1999, ApJ,
510, 874

Ohsuga K., Mineshige S., 2011, ApJ, 736, 2

Orosz J. A., Groot P. J., van der Klis M., McClintock J. E., Garcia M. R., Zhao P.,
Jain R. K., Bailyn C. D., Remillard R. A., 2002, ApJ, 568, 845

Osterbrock D. E., Pogge R. W., 1985, ApJ, 297, 166

Padovani P., Giommi P., 1995, MNRAS, 277, 1477

Padovani P., Urry C. M., 1990, ApJ, 356, 75

Padovani P., Urry C. M., 1991, ApJ, 368, 373

Padovani P., Urry C. M., 1992, ApJ, 387, 449

Pancoast A., Brewer B. J., Treu T., 2014, MNRAS, 445, 3055

Papadakis I. E., Nandra K., Kazanas D., 2001, ApJ, 554, L133

Perley R. A., Dreher J. W., Cowan J. J., 1984, ApJ, 285, L35

Peterson B. M., 1997, An Introduction to Active Galactic Nuclei

Peterson B. M., Horne K., 2004, Astronomische Nachrichten, 325, 248

Plant D. S., Fender R. P., Ponti G., Muñoz-Darias T., Coriat M., 2014, MNRAS, 442, 1767

Porquet D., Reeves J. N., O'Brien P., Brinkmann W., 2004, A&A, 422, 85

Pounds K. A., Reeves J. N., Page K. L., O'Brien P. T., 2004, ApJ, 605, 670

Poutanen J., Fabian A. C., 1999, MNRAS, 306, L31

Poutanen J., Veledina A., 2014, Space Sci. Rev., 183, 61

Poutanen J., Vurm I., 2009, ApJ, 690, L97

Proga D., Kallman T. R., 2002, ApJ, 565, 455

Proga D., Kallman T. R., 2004, ApJ, 616, 688

Qiao E., Liu B. F., 2013, ApJ, 764, 2

Reeves J. N., Turner M. J. L., Pounds K. A., O'Brien P. T., Boller T., Ferrando P., Kendziorra E., Vercellone S., 2001, A&A, 365, L134

Reis R. C., Fabian A. C., Ross R. R., Miller J. M., 2009, MNRAS, 395, 1257

Remillard R. A., McClintock J. E., 2006, ARA&A, 44, 49

Rezzolla L., Barausse E., Dorband E. N., Pollney D., Reisswig C., Seiler J., Husa S., 2008, Phys. Rev. D, 78, 044002

Riquelme M. A., Quataert E., Sharma P., Spitkovsky A., 2012, ApJ, 755, 50

Risaliti G., Elvis M., 2010, A&A, 516, A89

Ross R. R., Fabian A. C., 2005, MNRAS, 358, 211

Russell D. M., Fender R. P., Gallo E., Kaiser C. R., 2007, MNRAS, 376, 1341

Russell D. M., Gallo E., Fender R. P., 2013, MNRAS, 431, 405

Russell D. M., Maitra D., Dunn R. J. H., Markoff S., 2010, MNRAS, 405, 1759

Rybicki G. B., Lightman A. P., 1979, Radiative processes in astrophysics

Shakura N. I., Sunyaev R. A., 1973, A&A, 24, 337

Shaw M. S., Romani R. W., Cotter G., Healey S. E., Michelson P. F., Readhead
    A. C. S., Richards J. L., Max-Moerbeck W., King O. G., Potter W. J., 2012, ApJ,
    748, 49

Shaw M. S., Romani R. W., Cotter G., Healey S. E., Michelson P. F., Readhead
    A. C. S., Richards J. L., Max-Moerbeck W., King O. G., Potter W. J., 2013, ApJ,
    764, 135

Shimura T., Takahara F., 1995, ApJ, 440, 610

Shlosman I., Vitello P. A., Shaviv G., 1985, ApJ, 294, 96

Sikora M., Begelman M. C., 2013, ApJ, 764, L24

Sikora M., Begelman M. C., Rees M. J., 1994, ApJ, 421, 153

Sikora M., Stawarz Ł., Lasota J.-P., 2007, ApJ, 658, 815

Silk J., Rees M. J., 1998, A&A, 331, L1

Sobolewska M. A., Papadakis I. E., Done C., Malzac J., 2011, MNRAS, 417, 280

Socrates A., Davis S. W., Blaes O., 2004, ApJ, 601, 405

Springel V., White S. D. M., Jenkins A., Frenk C. S., Yoshida N., Gao L., Navarro
    J., Thacker R., Croton D., Helly J., Peacock J. A., Cole S., Thomas P., Couchman
    H., Evrard A., Colberg J., Pearce F., 2005, Nature, 435, 629

Stepney S., 1983, MNRAS, 202, 467

Stirling A. M., Spencer R. E., de la Force C. J., Garrett M. A., Fender R. P., Ogley R. N., 2001, MNRAS, 327, 1273

Suzuki K., Matsuoka M., Inoue H., Mitsuda K., Tanaka Y., Ohashi T., Hirano T., Miyamoto S., 1984, PASJ, 36, 761

Takahashi T., Mitsuda K., Kelley R., Aarts H., Aharonian F., Akamatsu H., Akimoto F., Allen S., Anabuki N., Angelini L., et al. 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 8443 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, The ASTRO-H X-ray Observatory. p. 1

Takeuchi S., Ohsuga K., Mineshige S., 2014, PASJ, 66, 48

Tavecchio F., Ghisellini G., Ghirlanda G., Foschini L., Maraschi L., 2010, MNRAS, 401, 1570

Timmer J., Koenig M., 1995, A&A, 300, 707

Tombesi F., Cappi M., Reeves J. N., Palumbo G. G. C., Yaqoob T., Braito V., Dadina M., 2010, A&A, 521, A57

Torii S., Yamada S., Makishima K., Sakurai S., Nakazawa K., Noda H., Done C., Takahashi H., Gandhi P., 2011, PASJ, 63, 771

Turner T. J., Miller L., Reeves J. N., Kraemer S. B., 2007, A&A, 475, 121

Unger S. W., Pedlar A., Axon D. J., Whittle M., Meurs E. J. A., Ward M. J., 1987, MNRAS, 228, 671

Urry C. M., Padovani P., 1995, PASP, 107, 803

Urry C. M., Padovani P., Stickel M., 1991, ApJ, 382, 501

Uttley P., Cackett E. M., Fabian A. C., Kara E., Wilkins D. R., 2014, A&A Rev., 22, 72

Uttley P., McHardy I. M., 2001, MNRAS, 323, L26

Uttley P., McHardy I. M., 2005, MNRAS, 363, 586

Uttley P., Wilkinson T., Cassatella P., Wilms J., Pottschmidt K., Hanke M., Böck M., 2011, MNRAS, 414, L60

van der Klis M., 1989, ARA&A, 27, 517

van Paradijs J., McClintock J. E., 1995, X-ray Binaries, pp 58–125

Vaughan B. A., Nowak M. A., 1997, ApJ, 474, L43

Vaughan S., Fabian A. C., 2004, MNRAS, 348, 1415

Vaughan S., Fabian A. C., Nandra K., 2003, MNRAS, 339, 1237

Vaughan S., Uttley P., Pounds K. A., Nandra K., Strohmayer T. E., 2011, MNRAS, 413, 2489

Veledina A., Poutanen J., Vurm I., 2011, ApJ, 737, L17

Veledina A., Poutanen J., Vurm I., 2013, MNRAS, 430, 3196

Veledina A., Vurm I., Poutanen J., 2011, MNRAS, 414, 3330

Volonteri M., 2012, Science, 337, 544

Volonteri M., Madau P., Quataert E., Rees M. J., 2005, ApJ, 620, 69

Volonteri M., Sikora M., Lasota J.-P., 2007, ApJ, 667, 704

Vurm I., Poutanen J., 2009, ApJ, 698, 293

Wardziński G., Zdziarski A. A., 2001, MNRAS, 325, 963

Welsh W. F., Horne K., 1991, ApJ, 379, 586

Wilkinson T., Uttley P., 2009, MNRAS, 397, 666

Wilms J., Reynolds C. S., Begelman M. C., Reeves J., Molendi S., Staubert R., Kendziorra E., 2001, MNRAS, 328, L27

Woo J.-H., Urry C. M., 2002, ApJ, 579, 530

Xie F.-G., Yuan F., 2012, MNRAS, 427, 1580

Yamada S., Makishima K., Done C., Torii S., Noda H., Sakurai S., 2013, PASJ, 65, 80

Yu Z., Yuan F., Ho L. C., 2011, ApJ, 726, 87

Yuan F., Cui W., 2005, ApJ, 629, 408

Yuan F., Cui W., Narayan R., 2005, ApJ, 620, 905

Yuan F., Yu Z., Ho L. C., 2009, ApJ, 703, 1034

Yuan F., Zdziarski A. A., 2004, MNRAS, 354, 953

Yuan F., Zdziarski A. A., Xue Y., Wu X.-B., 2007, ApJ, 659, 541

Zdziarski A. A., Gierliński M., Rao A. R., Vadawale S. V., Mikołajewska J., 2005, MNRAS, 360, 825

Zdziarski A. A., Grove J. E., Poutanen J., Rao A. R., Vadawale S. V., 2001, ApJ, 554, L45

Zdziarski A. A., Lubiński P., Gilfanov M., Revnivtsev M., 2003, MNRAS, 342, 355

Zdziarski A. A., Lubiński P., Sikora M., 2012, MNRAS, 423, 663

Zoghbi A., Uttley P., Fabian A. C., 2011, MNRAS, 412, 59